Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

Burkhard Stiller, Jan Gerke,
Peter Reichl, Placi Flury

The Cumulus Pricing Scheme and its Integration into a Generic and Modular Internet Charging System for Differentiated Services

TIK-Report
Nr. 96, September 2000

Institut für Technische Informatik und Kommunikationsnetze
Computer Engineering and Networks Laboratory

Burkhard Stiller, Jan Gerke, Peter Reichl, Placi Flury:
The Cumulus Pricing Scheme and its Integration into a Generic and Modular
Internet Charging System for Differentiated Services
September 2000
Version 1.0
TIK-Report Nr. 96

Computer Engineering and Networks Laboratory,
Swiss Federal Institute of Technology (ETH) Zurich

Institut für Technische Informatik und Kommunikationsnetze,
Eidgenössische Technische Hochschule Zürich

Gloriastrasse 35, ETH-Zentrum, CH-8092 Zürich, Switzerland

The Cumulus Pricing Scheme and its Integration into a Generic and Modular Internet Charging System for Differentiated Services

Burkhard Stiller, Jan Gerke, Peter Reichl, Placi Flury

Computer Engineering and Networks Laboratory TIK, Swiss Federal Institute of Technology, ETH Zürich, Switzerland
E-Mail: [stiller | gerke | reichl | flury]@tik.ee.ethz.ch

Abstract

Pricing and charging are the most important management functionalities future commercial networks need to offer. Since the Internet is on the move to provide differentiated services, for the backbone based on the Differentiated Services Architecture (DiffServ), suitable and scalable management mechanisms are required. Based on a new view of pricing, considered as important management information, on one hand, the Cumulus Pricing Scheme (CPS) proposed targets particularly at DiffServ technology and it is the only approach known so far defining a clear relation between different time-scales of accounting periods, measurement periods, and charging periods. Prices in this scheme are based on flat fees and hence predictable and transparent. On the other hand the scheme is flexible enough to allow network management according to the actual forces of the market.
CPS is backed by the design of a generic and modular Internet Charging System. It offers a service-independent architecture and integrates economically-controlled network management functions of charge calculation and pricing. It has been instantiated to an Internet Charge Calculation and Accounting System (ICCAS) to offer user support functions and is utilized for different Internet services. In particular, it has been applied explicitly in this work to DiffServ and the new Cumulus Pricing Scheme.

Keywords: *Charging Systems, Internet Pricing, Differentiated Services, Service Level Agreements.*

1 Introduction

Despite of the huge amount of bandwidth provided by today's Internet, bottlenecks are still far from having disappeared, *e.g.*, as dropped packets or the long delays experienced by certain packets easily demonstrate. But since a number of applications require at least a statistical guarantee of service characteristics as generally expressed in Quality-of-Service (QoS) parameters, the underlying networking technology needs to offer appropriate mechanisms for managing, providing, and supervisioning a specified set of QoS.

Therefore, pricing and charging will become highly relevant for offering more than a single service class, *e.g.*, a best-effort and a statistically guaranteed class, since financial incentives are the most effective stimulators for triggering user behavior according to users' real communication demands. These tasks are considered an essential part of the business and technology management aspects as well as services management, as defined in the management building [24]. However, the selection of a pricing scheme for a single-service network is easier than for a multi-service network [18]. Therefore, the issue becomes even more complicated for multiple QoS levels. Since Differentiated Services (DiffServ) defines an emerging technology for ease-of-use of multiple services in the Internet's backbone [42], an appropriate solution for pricing and charging these services is essential within a commercialized Internet services market. Internet Service Providers (ISP) are the key players in the communication market and they will offer differentiated Internet services in the near future. Flat fee pricing schemes seem to be favored by users and usage-based charging mechanisms seem to be only viable, if they can be implemented efficiently. A variety of investigations and technology developments in these areas have taken place, such as INDEX [8] and CATI [51], Therefore, a suitable combination of basic flat fees and time-depending or volume-dependent add-on charges and feedbacks for the utilization of differentiated services could offer an appropriate mixture of both worlds. Such schemes are considered as essential management tasks on different time-scales. They effect directly change management and operations management as well as performance management aspects [34].

There exists already a variety of pricing schemes for Internet access (cf. Section 3.1). This paper establishes a new view on them by considering the flow of management information to be as important as their content. In this sense, there are three main types of pricing schemes to be discriminated. Usage-sensitive or volume-based schemes rely on information on the user behavior that may be gathered at the interface between customer and ISP. Congestion-based pricing schemes, on the contrary, depend on the network status that depends only indirectly on the behavior of the single user, i.e. through her externalities. Flat rate schemes, finally, are generally independent of any metering information.

Similarly, the characterization of service classes may take place in three ways: either the user specifies her requirements, the ISP offers pre-defined service classes, or the quality of service classes constitutes itself automatically as a result of certain user and ISP independent dynamic processes, *e.g.*, as it is the case with Paris Metro Pricing. Moreover, this new view significantly contributes to the approach for classifying pricing schemes which will be presented.

With respect to the DiffServ networking approach, the future Interne backbone will consist of many different DiffServ domains, each of them operating under a different administration. Therefore, two principally distinct locations for applying the pricing and charging mechanisms exist. This is (a) the interface between a user and its DiffServ provider and (b) the interface between two DiffServ providers. For this reason, the view termed "Edge Pricing" by [12] has to be extended to include all the edges between neighboring providers as well. Initial ideas of these extensions have been already expressed in [51] and [5].

To overcome the lack of suitable solutions, a pricing scheme for multiple service classes and a charging and accounting system for the Internet, specifically targeted towards DiffServ service classes, has been developed. The Cumulus Pricing Scheme (CPS) is the only approach known so far, which defines a clear relation between different time-scales of accounting periods, measurement periods, and charging periods. Prices in this scheme are predictable and transparent, since they are based on flat fees. In addition, this scheme is technically efficient, because the ISP defines itself appropriate policies for measurement periods or thresholds. CPS provides an early warning system for users in terms of indications that the currently observed user traffic pattern will eventually result in contract changes. However, CPS allows at the same time for the fluctuation of user behavior within certain boundaries. So-called Cumulus Points (CP) are applied as an ISP-to-user feedback mechanism.

The CPS is technologically backed by the design of a generic and modular Internet Charge Calculation and Accounting System (ICCAS), which offers a service-independent architecture. It has been developed to be utilized for different Internet services, to perform multi-service accounting, and to support transport, service and content charging. The ICCAS provides modular functionalities, which can be instantiated according to particular ISP requirements. For the work presented here, it has been applied explicitly to the DiffServ technology and the new Cumulus Pricing Scheme.

1.1 Related Work

The number of projects concerned with pricing, charging, and accounting tasks in the Internet has increased quite significantly over the last few years. Therefore, only a small number of recent and charging-centric work of system's design and modeling is summarized below. While related pricing models are discussed in Section 3.1 in closer detail, another broad overview is to be found in [50] and [53].

Many projects dealing with charging and accounting functionality on the network level try to achieve a high independence from pricing models [54]. However, it has been argued that pricing in general and usage-based pricing in particular can impose a high overhead on telecommunication systems [35], [47]. Any form of usage-based pricing for various telecommunication services is interesting, because underlying resources (such as satellites, frequencies, cables, routers/switches, and most notable operating personnel) are scarce and very costly. The traditional Internet pricing model has been critiqued constantly in the past years for its economic draw-backs of not being incentive-compatible [47], [10], and [23]. Furthermore, it is inflexible — for example, it does not allow for combined sender/receiver payments — and does not provide economic signals which are needed for network planning and expansion. But most importantly, the current model is based on the assumption of a single service best-effort network that provides a similar service to all customers. Therefore, the multi-service paradigm needs to be investigated with respect to heterogeneous networking infrastructures and technologies of the Internet. An early per-flow billing system for TCP (Transmission Control Protocol) flows and initial ideas on a billing service design is presented in [15] and [49] respectively. Advanced per-flow charging and accounting approaches based on reservations have been tackled in [17] and [28].

The objectives of the Swiss National Science Foundation project CATI (Charging and Accounting Technology for the Internet) [51] included the design, implementation, and evaluation of charging and accounting mechanisms for Internet services and Virtual Private Networks [22]. This covered the enabling technology support for open, Internet-based Electronic Commerce platforms in terms of usage-based transport service charging as well as high-quality Internet transport services and its advanced and flexible configurations for VPNs. In addition, security-relevant and trust-related issues in charging, per-flow accounting [14], and billing processes have been investigated. Important application scenarios, such as an Internet telephony application [52], demonstrated the applicability and efficiency of the developed approaches [17]. This work was complemented by investigations of cost recovery for Internet Service Providers, including various investigations of suitable usage-sensitive pricing models for end-to-end communications based on reservations as well as Service Level Agreements in-between service providers.

The main assumption of the work on "Lightweight Policing and Charging" is that a multi-service packet network may be achieved by adding classification and scheduling to routers, but not policing [5]. Therefore, a lightweight, packet-granularity charging system has been investigated emulating a highly open policing function, which is separated from the data path. The amount of charging functions required depends on the customer's selection of services and is operated on the customer's platform. The proposed architecture includes a set of functions distributed to customers, which may include metering, accounting, and billing as well as per-packet or per-flow policing and admission control. The proposal concludes that lower cost is achieved through simplicity without sacrificing commercial flexibility or security. Different reasons for charging, such as inter-provider charging, multicast charging, and open bundling of network charges with those for higher class services, are all catered for within the same design.

Another highly important question concerns the issue of user acceptance of pricing schemes. In early 1998 the INDEX project (Index Demand Experiment) started in order to investigate user reaction when exposed to various pricing schemes for different qualities of Internet access. It turned out that the users were not disinclined to flexible pricing models. Moreover, the widespread flat rate model, at least in its pure form, proved to tend towards waste of resources, unfairness among users and revenue losses for the ISPs. Therefore, an "alternative" ISP has been proposed, offering differentiated services with dynamic volume-based pricing and suitable feedback mechanisms to inform the user on her own pattern of consumption. These project results have become a major stimulus for the efforts of shifting Internet pricing schemes away from the simple flat rate model [8], [16].

The project M3I (Market-Managed Multi-service Internet) [41] started just recently and aims to design and implement a next-generation system that will enable Internet resource management through market forces, specifically by enabling differential charging for multiple levels of service. This approach will increase the value of Internet services to customers through greater choice over price and quality, and a better service quality through reduced congestion. For the Internet Service Provider, flexibility will be improved, management complexity reduced, and the potential for increasing revenues is lead. Price-based resource management pushes intelligence and, hence, complexity to the edges of the Internet, ensuring the same scalability and simplicity of the current known network. It is intended to design a trial system, which will enable players in the Internet services market to explore sophisticated charging options and business models with their customers.

1.2 Terminology

Based on the discussion of related work and the necessity to obtain a uncontradicted terminology, the following basic definitions will be applied within this work.

- **Accounting** defines the summarized information (accounting records) in relation to a customer's service utilization. It is expressed in metered resource consumption, *e.g.*, for the end-system, applications, middleware, calls, or any type of connections.
- **Billing** defines the collection of charging records, summarizing their charging content, and delivering a bill or invoice including an optional list of detailed charges to a user.
- **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. Therefore, charge calculation applies a given tariff to the data accounted for.
- **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, however, it includes the full handling of invoices and customer data tasks, which are of second priority for the technological aspects of charging.
- **Mediation** is intended to filter, aggregate, and correlate raw technical data which in most cases has been collected by metering. Mediation transforms these data into a form which can be used for storing and further processing.
- **Metering** determines the particular usage of resources within end-systems (hosts) or intermediate systems (routers) on a technical level, including Quality-of-Service (QoS), management, and networking parameters.
- **Pricing** covers the specification and setting of prices for goods, specifically networking resources and services in an open market situation. This process may combine technical considerations, *e.g.*, resource consumption, and economical ones, *e.g.*, applying tariffing theory or marketing methods. Prices may be calculated on a cost/profit base or on the current market situation.
- **Tariff** defines the algorithm used to determine a charge for a service usage. It is applied in the charge calculation for a given customer to the service he utilizes. To calculate the charges the tariff may contain, *e.g.*, discount strategies or rebate schemes and marketing information.

The reminder of this paper is organized as follows. Section 2 introduces the basics of the Differentiated Services (DiffServ) approach and identifies the "hooks" where pricing and charging can be applied to. The proposed Cumulus Pricing Scheme (CPS) is outlined, defined, and discussed in Section 3. Afterwards in Section 4, a service-independent generic and modular Internet Charging and Accounting System (ICAS) is introduced, designed, and its particular utilization for CPS and DiffServ is discussed. Finally, Section 6 summarizes the work, draws conclusions, and depicts future work.

2 Differentiated Services in the Internet

At the beginning of the Internet the conveyance of data has been designed as a best effort service. With the raising Quality-of-Service (QoS) requirements of applications several proposals have been submitted to meliorate or surmount this heritage. The most illustrious are the Integrated Services Architecture (IntServ) [3] and the Differentiated Services Architecture (DiffServ) [42] approaches.

IntServ, and especially its offspring, the Resource Reservation Protocol (RSVP) [4], allow highly tailored QoS specifications and precise resource reservations to enforce the specified QoS. The consequences are that every application flow needs to be managed individually and as a result that IntServ lacks of proper scalability for large-scale backbone networks.

The DiffServ approach avoids these disadvantages of IntServ by shifting complexity to network edges and by the discretization of QoS specifications into some few classes. The complexity which is shifted to the borders comprises especially admission control, traffic conditioning (traffic shaping), and last but not least charging and accounting facilities.

2.1 Basic DiffServ Concepts

In DiffServ single application flows are assigned and accumulated to Behavior Aggregates (BAs) at the edges of the DiffServ domain. Packets belonging to a specific BA are identified by DiffServ Code Points (DSCPs). The association and enforcement of the BA's QoS is fulfilled by specific forwarding policies, the so-called Per Hop Behaviors (PHBs). Services in terms of network delivery characteristics can, thus, be described by BAs and their associated PHBs.

Figure 1 depicts a situation, where four different flows (Flow 1 to 4) are accumulated into two different BAs. Since the DiffServ Border Router and the network dispose of only a limited amount of resources, flow set up has to be requested, and accepted flows have to comply to specified traffic characteristics. The access request, implying the commonly known admission control and the conditioning of the traffic, necessarily according to an allegation, i.e. a contract, demand for signaling and for appropriate contracts negotiations.

The status quo of today's DiffServ investigations [1] targets at a central entity, called Bandwidth Broker (BB) [55], which is responsible for signaling and admission control and which configures furthermore the traffic conditioning entity at the involved border router(s). The criteria for flow acceptance or rejection and for setting traffic conditioning parameters are deduced from agreements between the customer and the service provider.

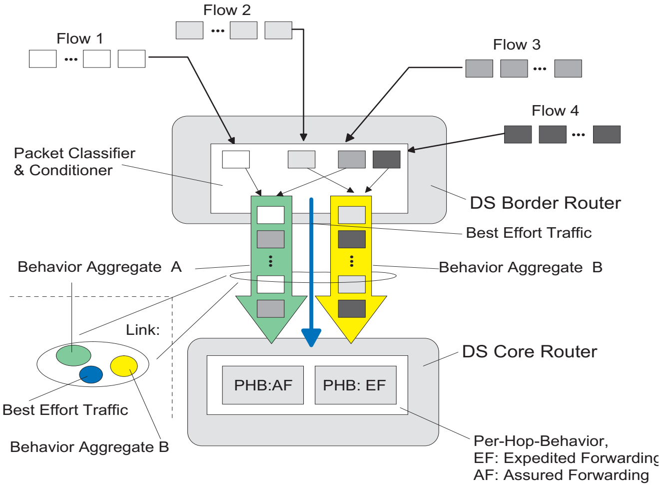

Figure 1: Flow Conditioning, Aggregation and Per-Hop Behavior.

2.2 DiffServ Service Classes

As mentioned DiffServ quantifies services to a small number - at least compared to IntServ - of different classes. Applications need to meet their QoS requirements by choosing from this pool of service classes.

The different service classes either result from the combination of BAs and their PHBs or are simply best effort services. In the momentary DiffServ specification two PHBs, i.e. Expedited Forwarding (EF[1]) [26] and Assured Forwarding (AF) [26], are distinguished besides of best effort. EF provides facilities to create services with low time constraints, *e.g.*, low loss, low latency, low jitter, and with assured bandwidth, so that EF is often said to have the features of a *leased line*, i.e. a virtual leased line. Nevertheless, EF does not optimally satisfy all fields of application, since it provides only a rather weak support of bursts. The AF PHB instead is in fact able to cope with almost any bursty behavior, however, at the expense of the delivery of the IP packets, which can only be guaranteed with certain high probabilities. Furthermore several AF classes can be distinguished to allow for a fine-grained handling of different application requirements. Services and applications with ordinary QoS requirements preferably continue to deploy best effort forwarding. Note that best effort in DiffServ subtly differs from best effort of the early days of the Internet, since the congestion control mechanisms of the Transmission Control Protocol (TCP) are undermined by traffic shaping mechanisms of DiffServ (e.g. AF packets may be degraded to best effort packet in case of congestion). Thus best effort generated from DiffServ congestion policies preempts ordinary application best effort.

1. EF has formerly been named Premium Service.

2.3 Inter-provider Agreements

The relation between the customer, i.e. service user[2], and the ISP, i.e. service provider, requires to be expressed by one or more contracts called Service Level Agreements (SLAs [2]). An SLA comprises technical, *e.g.*, expected bandwidth requirements or certain QoS parameters, and legal and financial aspect, *e.g.*, tariff structure or usage charges, for service delivery and in case of service rejection for possible recoupments. Note that the specific form of SLAs applicable to the Internet is still rather open and subject to heavy discussions. Traditional telecommunication SLAs have been utilized in telephone networks for years, however, they lack the description of Internet traffic characteristics.

From an SLA, which is defined and set up by human parties, Service Level Specifications (SLS) are deduced [55]. Each SLS contains the technical details for service provision. SLSs have the advantage that they can be negotiated and set up automatically, i.e. between communication parties without any user intervention, what is especially advantageous for dial-up and mobile customers. From the SLS, which represents only the commitment for resource reservation, the Traffic Conditioning Specifications (TCS) are derived. A TCS specifies classifier, metering, marking, shaping and dropping rules which are to be applied to the selected BA.

A typical SLA setup and enforcement comprises, that a customer contacts her ISP for specification of her admitted requirements. In this primary user point of view, the customer estimates the QoS she needs for her applications. According to her knowledge the parameters can either be very precise, e.g. max delay, jitter etc. or inaccurate, e.g. ‘a download server has to be supported’. Here the point of view changes and it is up to the ISP to map the customer’s QoS formulation to its supported DiffServ classes and to its available bandwidths. This mapping is transparent to the customer, nevertheless an SLA is a dialog, where customer and ISP have to agree upon. Besides technical aspects financial and legal ones have to be discussed and affirmed. The customer-ISP relation will be described in more detail in Section 3.2.6.

3 Cumulus Pricing Scheme — CPS

Over the last couple of years, the issue of pricing and tariffing Internet services has become an important and widespread research issue. However, there is still a lack of definite solutions, as the proposed schemes either have turned out to be too inflexible or too complex for a satisfying technical solution. Hence, there is still an urgent need for new pricing schemes that are able to deal with the various requirements, such as technical and economic ones, demanded of them.

As it has been shown above, the term “pricing” generally concerns the process of assigning a price to services, products or contents. For the scope of this paper, the focus is on two main types of pricing schemes, i.e. flat rate vs. volume/usage-based ones, and forget about the third type, i.e.

2. The customer can also be an entire domain, *e.g.*, enterprise’s private LAN, or another ISP.

congestion pricing as mentioned in the introduction. A flat rate scheme in its most simple version consists of a constant rate, i.e. charge per month, which upon payment includes free usage of the chosen service without restrictions. These schemes are very popular in today's Internet as their technical demands are rather limited whereas their user acceptance is high (cf. the America On-line flat rate scheme [38]). Volume-based pricing schemes, on the other hand, usually require relatively high technical expenses, but obviously charge what is used and in this sense are much fairer to the user. Within volume/usage-based schemes, there are again two possible solutions: (1) the price per unit of volume may be fixed (terming these schemes "static") or (2) variable, i.e. depending on current market conditions. A good example for the latter, so-called "dynamic" schemes are Generalized Vickrey Auctions [37] with their well-known complexity in terms of implementing them [36].

One main reason for introducing dynamic pricing schemes comes from the desire to use them as an important network management tool [41]. If it is assumed that users react to changing prices by suitably adapting their network activities, it appears natural to express, *e.g.*, network congestion or service quality in terms of prices. The higher such a price will be, the more cautious the user will use the respective service in the respective situation. Thus, dynamic pricing offers a simple way for the ISP to influence user behavior and hence to optimize the global network status with respect to, *e.g.*, utilization, fairness, or revenue.

But there is a second way for pricing models to be used as network management tools. This alternative is based on eliciting user information about expected usage pattern. With the help of this information, the ISP is enabled to optimize, *e.g.*, network configuration or admission control, with respect to objectives like maximization of utilization or revenue. The role of pricing models in this context consists of pushing the user towards delivering her information correctly. There are several well-known pricing schemes proceeding like this, among them the "incentive compatible" Generalized Vickrey Auctions [37] or Kelly's effective bandwidth approach for tariffing ATM traffic [30].

The rest of the section is concerned with introducing a new pricing model that appears to mix the respective positive properties of the different general approaches mentioned so far in a well-balanced way. After presenting a short overview on current pricing models for the Internet and further related work, in Section 3.2 the basic idea of the new "Cumulus Pricing Scheme" CPS is introduced, and a short example is demonstrated. Afterwards, the influence of ISP policies is investigated and the example is continued. The next step consists of generalizing the basic scheme with respect to its application within the DiffServ framework, and finally the major properties of CPS are reviewed in more detail.

3.1 Pricing Models in Theory and Practice

The sobering thing about pricing models is that almost every proposed architecture trying to satisfy QoS demands develops its own pricing models. Thus, different and usually incompatible pricing models exist, reflecting only rather modest standardization efforts. However, from the pool of existing pricing models basic classification and distinction may be deduced.

Therefore, a classification based on five attributes and a number of well-defined parameter each has been developed. Amongst other possibilities, pricing models can be classified according to the attributes time, space, quality (i.e. class quality characterization), technological requirements, and volume where each of these attributes shows a set of different parameters as summarized in Table 1.

For the attribute time, parameter semantics have been applied:

- *Duration:* defines the elapsed time between start and end of service usage, *e.g.*, duration of a video conference.
- *Period:* determines the committed length measured in time, which is per se independent, i.e. decoupled, of the service appliance. This commitment is usually set up in advance, *e.g.*, a leasing period.
- *Time of day*: defines the sensitivity of service usage to a given time of day. The influence of the time of day may be known in advance by the customer, *e.g.*, weekend tariffs, or they as well may dynamically change, *e.g.*, based on congestion in an auction system.
- *Not applicable*: means that the attribute *Time, Space, Class Quality Characterization, Technological Requirement,* or *Volume* is not relevant, *e.g.*, for the time attribute *not applicable* means, that the time has no significance at all, which can be reasonable, *e.g.*, in a volume based system.

The attribute spaces allows for the following parameters:

- *Distance:* is the length of the (virtual circuit) from the sender to the receiver, which is passed by messages. Its length in meters is not relevant, but rather how much infrastructure has been used to enforce service provision.
- *Route/path*: in contrast to distance the attribute route/path describes the relevance of where the message flow passes through, i.e. through which, how many, and what kind of routers. The route/path attribute plays in important role, when particular associations are made between the chosen circuit and the service.
- *Location*: *distance* and *route/path* parameters are not sufficient alone to describe all occurring cases for pricing models. Suppose edge-pricing has to be expressed. Saying that the distance and the route/path are not relevant implies a transparent network from the point of view of space. Indeed a transparent network (cloud) does not imply local importance of service provisioning. Thus *location* allows to consider places/entities in the network, which have particular importance for the pricing model, *e.g.*, just like edge pricing.

The Class Quality Characterization attribute describes the sensitiveness of pricing models for quality classes. It mainly explores how a differentiation of quality is made and who is taking influence on the selection or creation of quality classes, i.e. if only the ISP, the customer, or both. It has to be noted that a differentiation of quality does not imply that only a fix number of classes exists.

Table 1: Pricing Model Attributes and Parameters

Time	Space	Class Quality Characterization	Technological Requirements	Volume
duration	distance	ISP	flow-based	linear cumulation
period	route/path	customer	class specified	non-linear cumulation
time of day	location	self adjusting	not applicable	not applicable
not applicable	not applicable	indifferent		
		not applicable		

- *ISP*: sets up quality classes. Often the ISP will have a limited set of quality classes, which it may slightly adapt and distribute among customers.
- *Customer*: initiates and defines quality class specification, *e.g.,* with a signaling protocol such as RSVP.
- *Self adjusting*: the class quality may change with network state, where the correction towards the new stable state is performed in a system inherent manner, *e.g.*, as with PMP [43].
- *Indifferent*: no different quality classes available.

For the attribute Technological Requirements the following parameters exist:

- *Flow-based*: the supporting network technology offers a clear technology for maintaining flows within the network, such as with the Integrated Services Architecture.
- *Class-based*: the network supplies a set of discrete classes, where the classes are not necessarily associated to particular technologies or QoS commitments.

Finally, the volume attribute distinguishes:

- *Linear cumulation*: the amount of data is accumulated linearly, determining that every single data unit measured has the same weight.
- *Non-linear cumulation*: covers all other cases, where the volume is taken into account of a pricing model.

It is obvious that the combination of all parameters gives raise to a large amount of different pricing models. It is up to the designers to agree upon the most reasonable ones. As already seen by the parameter *not applicable,* pricing models do not require to precise themselves on all attributes. In case that a pricing model has to chose just one parameter per attribute is inappropriate. Therefore a supplementary notation is introduced. The following two variables *x* and *X* describe following alternatives of the importance of a given parameter per attribute:

- x: exactly one, but an arbitrary parameter of an attribute needs to be set, *e.g.,* for the attribute *Volume*: x = [*linear cumulation* | *nonlinear cumulatio*n | not applicable]
- X: at least one, but an arbitrary number of parmeter(s) of the attribute needs to be set. This is required, if a combination of parameters is utilized to precisely define the scope of the pricing model, *e.g.,* for the attribute *Space*: X= *distance* & *route/path*.

Consider the example of a flat rate pricing scheme. Over a fixed time, i.e. described by the attribute and parameter *time* and *period*, respectively, a customer can send as many packets as she likes, i.e. the attribute volume is not applicable. No metering and charging entities are needed, since the volume is irrelevant in this classical flat rate example, the *space* attribute is set to *no relevance* as well. The *quality* attribute instead may have some influence to the initial price set for flat rate, but it is not a necessity for flat rate pricing, so it can be set to an arbitrary parameter. In Table 2 the latter fact is expressed by an ‘*x*’. In addition, the volume-based (static and edge) pricing, as well as the Paris Metro Pricing scheme are described.

Looking at the goals at which pricing models target at, a clear polarization on two concurrent topics can be recognized. The first set of models targets at congestion in networks, i.e. congestion control and avoidance. It is a global approach, i.e. including the entire scope of the ISP’s domain, representing ISP desires in the management of limited network resources by facultative (*e.g.,* PMP) deployment of network technologies (*e.g.,* ECN) and categorical disposal of pricing models [19]. The second goal to be achieved, in contrast to the first one, aims at individual QoS provisioning. It will meet dedicated customer requirements, where pricing models are likewise necessary.

Although the pricing models look like the common denominator of these concurrent goals, they are strongly influenced by them and, thus, can not completely be decoupled from the objectives of the network provider, i.e. if the provider tries to satisfy customer QoS requirements or if it tries to avoid and control congestion. The following paragraphs contain an overview of important Internet pricing models, which have been investigated over the last few years and have turned out to be of special importance from a practical and economic point of view.

3.1.1 Edge Pricing

The fundamental idea of edge pricing is that all pricing decisions are made at the edge of the ISP locally [12], [47]. The advantage is that no uniform pricing standards need to be developed since ISP interconnection involves only bilateral agreements, *e.g.,* in DiffServ this would be part of SLAs between ISPs. So it doesn’t require the whole Internet to support edge pricing schemes at once. Furthermore the transparency of edge pricing enables ISPs to use adapt and evolve own pricing policies independently.

In a basic approach, the customer 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 [17], [28].

Table 2: Example Pricing Models

Example	Time	Space	Class Quality Characterization	Technological Requirements	Volume
Flat Rate	period	not applicable	ISP[a]	x	not applicable
Volume-based (static and edge pricing)	duration	location	X	not applicable	[non]linear cumulation
PMP	duration	not applicable	self adjusting	class specified	x
Vickrey Auction	time of day	X	self adjusting	x	x
CPS	period	location	X (ISP&customer)	not applicable	[non-]linear cumulation

a. For classical flat rates, whereas for CPS, the customer takes the decision on class quality.

3.1.2 Profiles and Classes

[11] 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 dropping scheme is used to preferentially drop “out” packets. Note that this approach prevents traffic being separated at routers into different queues. The same scheme is seen within DiffServ’s AF classes [42], [1].

In contrast, [32] 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 [43], 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.

3.1.3 Volume-based Schemes

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 [6] and Great Britain [46] are already well-known. Moreover there are approaches to model accurately the relationship between the current utilization of a resource and the price to be paid for using it (cf. [44]).

3.1.4 Auction Mechanisms

A congestion pricing scheme (where packets are charged if and only if the network is congested) can 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 [36]. Each packet header contains a bid field, and the packet is admitted, if the bid exceeds the current marginal cost of transportation. In the Generalized Vickrey Auction the highest bidders win, but pay only the market-clearing price. The Vickrey auction approach gives all competing customers an incentive to disclose their true evaluation of the good [37]. The mechanism guarantees only relative priority, no absolute QoS. 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. Improvements of the Generalized Vickrey Auction can be found in Delta Auctions [17] and the CHiPS approach [45].

3.2 CPS in Detail

As discussed above, there exists already a considerable pool of related work on pricing models. In the rest of the chapter the new pricing scheme is presented that turns out to comprise most of the advantages of the different existing schemes while avoiding most of their problems.

3.2.1 Requirements

The new Cumulus Pricing Scheme (CPS) has been developed with respect to the following three main requirements as illustrated in Figure 2.

- **RT-1: Customer.** Over the last couple of years, there has been an extensive discussion on preferences customers show towards dynamic tariff schemes for Internet services, *e.g.*, INDEX [16], CATI [51], and M3I [41]. It has turned out that in this context flat pricing is still regarded as one of the most important approaches, if not in fact the most important one. All investigations performed within the work presented here have taken this fact into consideration. Therefore, the idea of CPS uses flat rate pricing as a conceptual starting point by enhancing it and bringing a degree of dynamicity into play.

- **RT-2: ISP - Economic.** The economic side of ISPs is interested in maximizing either network utilization, total revenue, or similar objectives. To this end, charging and pricing schemes represent an important interface to the customer: Prices may, *e.g.*, be used to indicate well-behavior or misbehavior of the customer or to signal the congestion state of the network. The idea is to use market forces (which reveal themselves in terms of prices) for network management. But this is only possible, if there is some a feedback mechanism, which allows the ISP to communicate the relationship between current customer behavior and overall system status.
- **RT-3: ISP - Technical.** The mentioned feedback mechanism depends on the existence of respective mechanisms for technical accounting. This requirement opens a vast field of possibilities as to which detail data about the network status are to be obtained. It appears almost impossible to commit ISPs to the availability of a certain minimal measurement requirement as the technical conditions use to vary enormously. Therefore, the approach proposed here pre-supposes the availability of measurements, but leaves the level of detail and granularity deliberately to the respective ISP. This openness will turn out to be one of the major advantages of the new scheme.

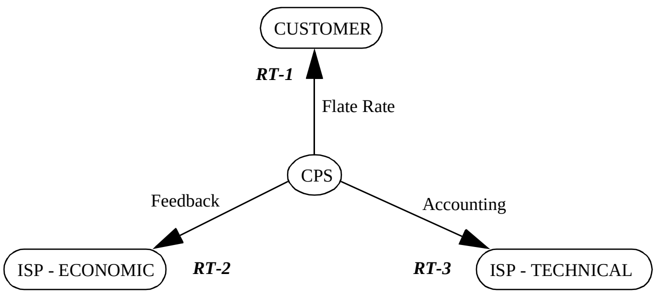

Figure 2: CPS Requirements

3.2.2 General Idea of Cumulus Pricing

Common Internet pricing schemes as introduced in Section 3.1 usually have trouble with at least one of the three requirements presented above. The fundamental decision between static and dynamic schemes touches immediately the customer's desires concerning price stability, *e.g.*, highly fluctuating auctions, whereas orienting a pricing scheme strictly according to the forces of the market readily induces technical infeasibility. A prominent example is provided by Kelly's effective bandwidth approach for ATM pricing, where practical considerations forced the implementation to be based on bounding approximations to the effective bandwidth function instead of the latter function itself, cf. [48], p. 24f. In this situation, the Cumulus Pricing Scheme is an approach to reconcile all three requirements.

CPS is basically a flat rate scheme (but rates may vary over long time-scales), it provides a feedback mechanism to bring market forces into play (where this feedback is not an immediate one, but requires the accumulation of a sufficient number of discrete "flags" indicating user behavior), and it allows a huge flexibility in terms of the technical prerequisites, especially concerning the measuring and accounting mechanisms of the required data records.

The key to the new solution proposed lies in building the contract between customer and ISP upon suitable information about the expected usage pattern of the service plus influencing the actual customer behavior by a new type of feedback mechanism that is specific in terms of its relation to different time-scales. *Measurements* take place *over a short time-scale* and allow evidence about *user behavior on a medium time-scale*. This evidence is expressed in terms of discrete flags (the so-called "Cumulus Points", cf. Section 3.2.5), yet not triggering some sort of *reaction* by themselves, but only as a result of their accumulation over a *long time-scale*.

3.2.3 Time-scales

With almost every pricing scheme for Internet services, the following time-scales ought to be discriminated. In this context, four different time-scales are relevant and introduced. In Section 4.2.3 this classification will be put into a larger framework, mainly used for management purposes as well as the charging and accounting system dimensioning. The four time-scales as mentioned are:

- **Atomic (communication-relevant):** This involves sending packets, round-trip times, and managing feedback between sender(s) and receiver(s).
- **Short-term (application-relevant):** This time-scale is concerned with the usual duration of applications like file-transfer, video-conferencing, or IP phone calls. The tasks of accounting and measuring are closely related to these activities.
- **Medium-term (billing-oriented):** The time-scale for performing billing actions is somewhat arbitrary, since it is influenced strongly by the usual lifestyle habits of humans, which includes, *e.g.*, monthly payments of rents, phone charges, or newspapers.
- **Long-term (contract-specific):** The largest time-scale in this context is the duration of contracts between customers and ISPs, which usually varies from several months to years. Sometimes, especially for technic freaks, it may be shorter.

Figure 3[3] sketches the relationship between these time-scales, the respective management tasks and communication contents (cf. Section 4.2.3) as well as pricing schemes.

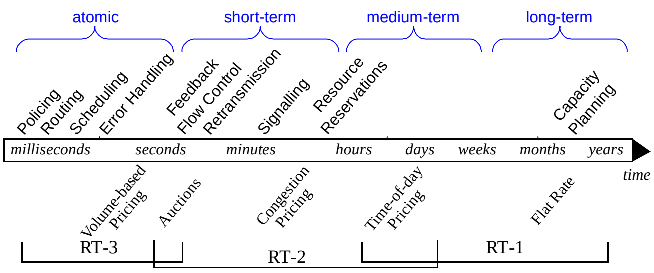

Figure 3: Time-scales

3.2.4 New Methodology

Moreover, these three requirement types may be assigned to appropriate time-scales. In doing so, it is noticeable that

3. This figure is based on [31].

they are placed apparently at transition points between time-scales: Transparency and predictability are relevant on medium- to long-term scales, economic efficiency refers to short- and medium-term scales, and the accounting technology deals with atomic and short-term events. Thus, the problem of balancing these requirements turns out to be an issue of reconciling these different time-scales. Therefore, the analysis of how these scales may be used for constructing a solution to the feasibility problem is important.

Usual flat rate schemes are basically intended to have no a priori time limitations, i.e. the rate is supposed to remain constant more or less forever. In order to introduce market forces, the first necessary step is to break with this supposition by introducing the possibility of price changes on the level of contracts (long time-scale). I.e. the tariff no longer guarantees the flat rate to be constant in the long run, but explicitly allows renegotiation of the contract (and here especially of the charges) in case of, e.g., user misbehavior or fundamental system changes.

The second important time-scale is the one connected with billing periods (medium time-scale). On this level, *e.g.*, in connection with her monthly bill, the user receives feedback about her current behavior and its consequences for the future validity of the contract. If the user has been detected by the ISP to strongly overuse capacities or to misbehave in some sense, she will usually receive some sort of warning that in case of unchanged behavior the contract may sooner or later may be finished by the ISP. On the other hand, if the user is underutilizing her capacities, this behavior may be rewarded by some sort of bonus system.

This feedback system strongly depends on measurement activities that are supposed to take place on the application time-scale, but are deliberately left open to the ISP. Hence, one could think of (1) an ISP monitoring each packet or each connection in the one extreme, (2) ISPs undertaking systematic monitoring, (3) ISPs measuring every now and then (maybe in some sort of statistical framework), or (4) at the other extreme ISPs not measuring at all.

The basic communication process, however, still has to happen on the atomic level. As these activities have been identified as causing the essential feasibility problems, it is proposed to prohibit them and instead express them by activities on the three remaining time-scales. Therefore, the resulting pricing scheme is characterized by a mapping of the atomic scale onto the short-, medium- and long-range scale.

This paradigm shift has immediate consequences, especially w.r.t. the design process itself, which is reduced to finding useful combinations of possible mechanisms on the three larger time-scales. This is demonstrated in the following sections by means of the Cumulus Pricing Scheme CPS.

3.2.5 Cumulus Points as Feedback Mechanism

Current dynamic pricing approaches, *e.g.*, congestion-based pricing or ECN pricing flags, usually rely on using charges as immediate tool for network management tasks. *E.g.*, if congestion is threatening, the prices for using the network are supposed to increase immediately, in this way indicating the congestion and leading users towards a behavior that allows the provider to cope with the critical situation (e.g. by some users postponing, flattening, re-dimensioning etc. their activities). These approaches require a high flexibility on the side of customers as well as ISPs. The usual arguments against them are twofold: On one hand, the technical complexity of monitoring and accounting necessary details of the (real-time) data is huge, on the other hand, the customer appears generally not to be interested in too heavy a fluctuation of charges to be paid.

Therefore, a feedback mechanism is proposed, which is based on so-called "Cumulus Points" (CPs). First of all, customer and ISP are supposed to agree on a contract specifying the expected user requirements in terms of bandwidth, delay etc. as well as a flat rate to be paid for this type of service. Following this agreement, the factual usage may not match the prediction given by the user (for whatever reason, *e.g.*, be it an incorrect statement, changing habits, or new applications). As soon as these discrepancies exceed some threshold, the user receives feedback in terms of the mentioned CPs. They exist as red and green flags: a red CP indicates that the user has been overusing her capacities, a green one indicates the opposite, i.e. that the user might have been allowed to use more resources than she actually did. The larger the discrepancy between contract (i.e. user statement) and reality, the more CPs may be delivered. CPs remain valid for a dedicated number of consecutive billing periods, and it is their accumulation that finally triggers to be defined consequences. I.e., receiving one or more CPs usually requires no immediate reaction. However, over consecutive billing periods CPs are accumulated and, thus, finally their sum may exceed a CP threshold. Only this latter event may have consequences for the user, depending on ISP policies.

3.2.6 Mathematical Description

Suppose that ISP *I* offers only one service, and initially customer *C* has stated her expected bandwidth requirements according to a contract (the SLA) to be *x* MB/s, whereupon ISP *I* has offered a flat rate tariff of *a* $/month for this service which customer *C* has accepted. In reality, the volume consumed by *C* is described by a function *V(t)* of time, which naturally may differ arbitrarily from the stated expected requirement *x*.

Let $\Delta_i = \Delta(t_i)$ describe the monthly over- or underutilize, respectively, of the customer with respect to her statement *x*, i.e.

$$\Delta_i = \int_{t_{i-1}}^{t_i} (V(t) - x)dt = \int_{t_{i-1}}^{t_i} V(t)dt - x(t_i - t_{i-1}), \quad (1)$$

where t_i describes the end of measurement period *i*, *e.g.*, the end of month, *i* = 0,1,2... (note that t_0 describes the start of the contract between ISP and customer).

Cumulus Points are assigned by the ISP *I* according to a rule (the so-called "CP Rule") whose content is up to the ISP, but typically might look like the following:

CP Rule: *Define* θ_n, $n = -N,\ldots,-1,0,1,2,\ldots,N$, *to be the CP thresholds,* $\theta_0 = 0$ *and* $\theta_{\pm(N+1)} = \pm\infty$ *where N describes the maximal number of CPs that could possibly be assigned for one measurement period. Then for measurement period i, the customer is assigned* c_i *cumulus points iff*

$$0 \le \theta_{c_i} \le \Delta_i < \theta_{c_i + 1} \text{ or} \quad (2)$$

$$\theta_{c_i - 1} < \Delta_i \le \theta_{c_i} \le 0, \quad (3)$$

the choice between (2) and(3) depending on $\operatorname{sgn}\Delta_i$.

Hence, if Δ_i is positive (i.e. overuse in period i) and lies between thresholds θ_c and θ_{c+1}, then c cumulus points are assigned. If Δ_i is negative and between thresholds θ_{c-1} and θ_c, then c cumulus points are assigned, where c now is a negative number, hence the cumulus points are referred to as “green” ones, whereas for positive c the cumulus points are “red”.

Now the cumulus points c_i are accumulated over time according to

$$\Gamma_n = \sum_{i=1}^{n} c_i, \quad (4)$$

hence, Γ_n describes the total sum of cumulus points assigned since the start of the contract.

The reaction to CP accumulation is again basically up to the ISP and is the content of a second rule, the so-called “Reaction Rule”, typically looking like this:

Reaction Rule: *Define* Θ *to be the reaction threshold. Then the contract between customer and ISP is in the state of imbalance and needs to be renegotiated after period n if*

$$|\Gamma_n| \ge \Theta. \quad (5)$$

Depending on $\operatorname{sgn}\Gamma_n$, there may as well be two different thresholds Θ^+ and Θ^- for red and green CPs, respectively.

Note the renegotiation of the contract may have different forms, as shown in Section 3.4 with respect to the example.

3.3 ISP Policies

The central question to be answered is how to deal with these CPs. The concept of Cumulus Points has a long tradition, especially in Switzerland and Germany. Usually, the green version of them is used to stimulate customers towards buying more often e.g. in supermarkets of one specific chain see, *e.g.*, the Swiss Migros “Cumulus Card”. The red form of CPs may be found, *e.g.*, in the German system of dealing with unlucky behavior in car traffic. Here, crossing red lights or driving cars in an not too sober state yields (in case the policeman has been around the corner) a certain small number of points to be noted in a central file located in Flensburg.[4] As soon as in the course of two years 12 or more points have been accumulated there, you are supposed to loose automatically your driving license. On the other hand, points may also be deleted, *e.g.*, in the case of good long-term behavior, of visiting additional courses about correct traffic behavior etc.

4. Guess why this city located in the north of Germany has got a reputation that is somehow very special to car drivers.

These examples show that the rules mentioned in Section 3.2.6 of how to use the scheme with respect to communication services and bandwidth in fact is completely up to the ISP. Different ISP policies include especially the following ones:

- **Measurements:** It has already been argued that it is almost impossible to find a standard way of network monitoring and accounting that is compulsory for all ISPs. With our proposal it is up to the ISP, on which data measurements the distribution of CP is based.
- **CP assignment:** Being assigned CPs depends on violating certain thresholds in terms of utilization or bandwidth. Fixing these thresholds is up to the ISP. As the CP assignment depends crucially on the accounting mechanism, the thresholds should be set such that different measurement techniques applied to the same consumption pattern eventually do not differ by more than one CP. Moreover, setting $\theta_{\pm 1} \ne 0$ prevents smaller oscillations around *x* to result in superfluous CP assignments.
- **Accumulation:** Usually, CPs are supposed to be accumulated over subsequent billing periods. However, as the traffic sinner example shows, CPs may be allowed to expire, or red CPs may be charged up against green ones.
- **Contract renegotiation:** Another threshold to be set freely by the ISP concerns the point at which the contract with the customer is supposed to be renegotiated. The way of renegotiating is also open. Either the customer delivers a new statement about expected QoS requirements, and the provider offers a new flat rate, or the old contract remains valid, and the accumulated discrepancy is resolved, *e.g.*, by an extra payment. Note that offered flat rates will depend on the size of resource consumption, i.e. the higher the capacity requirement, the lower the price per MB due to usual discounting. Moreover, the size of the extra payment in principle corresponds to an extra contract “ex post” over the duration time of the original contract and the capacity discrepancy. Hence, the flat rate per MB of the additional contract is higher, and the customer’s behavior is driven to be in accordance with the contract (*e.g.*, as shown in the example below, cf. Figure 4 during April).

3.4 Example

Figure 4 describes a typical example of how CPs are used. Customer *C* has stated her expected bandwidth requirements to be *x* MB/s, but the actual bandwidth consumption exceeds the agreed upon one slightly in January and heavily in February. Accordingly the consumer receives one red CP at the end of January and two additional red CPs at the end of February. Afterwards, her consumption falls below the expected value (one green CP in March), before it behaves exactly according to the contract in April (which is apparently the ideal situation). Later on, in May and June this value is exceeded again. The accumulation of the CPs as of end of June sums up to five red CPs and eventually requires a renegotiation of the original contract.

Continuing the example it can be safely assumed that it relies to the following ISP policy:

For each month the customer is awarded up to two CPs in each direction (one for slight, two for heavy discrepancy between expected and actually used bandwidth). Five red CPs form the threshold for renegotiation of the contract.

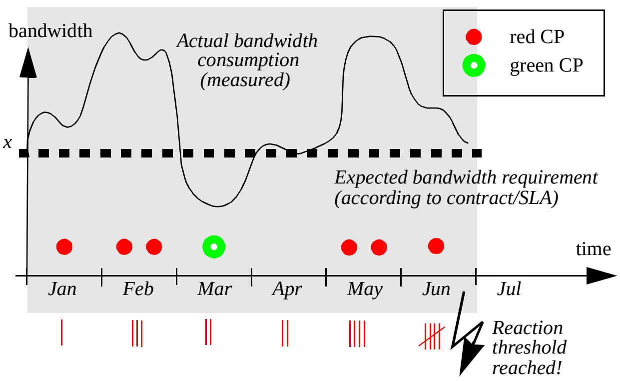

Figure 4: Red and Green Cumulus Points and Their Accumulation over Time

Note that these five red CPs actually indicate that over a longer time period there has been an extensive mismatch between expected and consumed bandwidth. To correct this mismatch, one could think of several possibilities.

- **Case 1:** The customer believes that the mismatch has been only temporarily and that her initial statement is still valid. Then the ISP may set an additional charge for each red CP, upon paying the charge the respective CP expire, and the entire system is basically reset to the initial condition (cf. Figure 5).

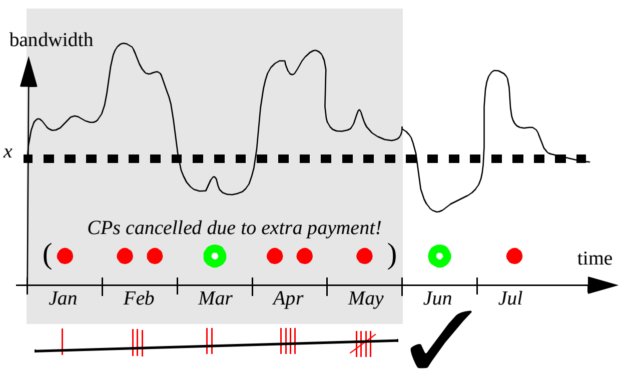

Figure 5: CPs Cancelled by an Extra Payment

- **Case 2:** Assume again that like in case 1 the customer does not want to change her commitment, but unlike in case 1 she does not want to make an extra payment in order to cancel the accumulated CPs. In this case, she could make a special additional contract (maybe also fixed in time, *e.g.*, over five months) over the average amount *x'* the bandwidth has been overconsumed over the last few months. If this additional contract is not used at all, it will solely produce green CPs month after month that may be charged up against the existing accumulation of red CPs until all of them have lost their validity. In this example, there are still three red CPs valid as of end of July (cf. Figure 6).

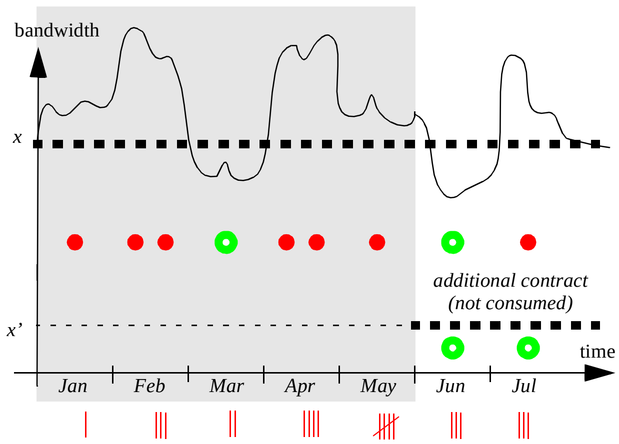

Figure 6: CPs Cancelled by an Additional Contract

- **Case 3:** Assume the customer now is convinced that the initial statement has not been correct. In this case a change of the contract would be necessary, now stating a higher expected bandwidth consumption *x"*. The accumulated sum of five red CPs could either be removed by an extra payment according to case 1, or by accidental under-utilization of the new statement *x"*, as it is the case in the following figure, leading to a status of 2 red CPs altogether as of end of July (cf. Figure 7).

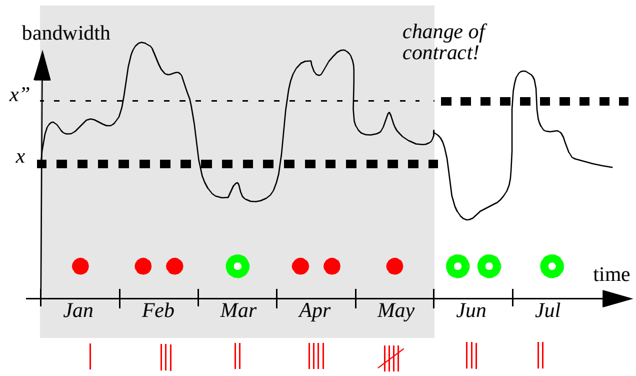

Figure 7: CPs Cancelled by a Contract Adaption

3.5 Evaluation of CPS

Summarizing shortly, the most important properties of the Cumulus Pricing Scheme include the following properties:

- **Discrete:** Instead of reporting every small fluctuation of the user behavior, CPS allows the user behavior to fluctuate within certain boundaries. The feedback is given in a quantized form.
- **Cumulative:** Not single bursts, but only continuous change of promised behavior triggers reactions, *e.g.*, in terms of changing prices.
- **Early warning:** If some user starts to change her behavior, this will have no immediate consequences, but nevertheless she will know at an early stage that her pattern has changed, and that continuing the new usage habits will require sooner or later a renegotiation of the current contract.

- **Predictable and transparent:** As the entire scheme presents itself to the user as a variation of the fully accepted flat rate schemes, there will not be any trouble in terms of user acceptance. Charges for the user remain stable over long time-scales, and necessary changes are transparent to the user because of the early warning character of the feedback mechanism.
- **Market managed:** The contract between customer and ISP is based on information delivered by the customer about her expected usage pattern. This information may be used by the ISP to determine the flat rate charge and moreover to optimize network utilization or revenue. As soon as the customer does not fulfill her commitments, CPS introduces a slow penalizing mechanism in order to trigger well-behavior of the customer or change of the contract.
- **ISP policy dependent:** There is a strong influence of ISP policy on the realization of the scheme. The ISP may freely decide not only over the measurement procedure, but also over the threshold values for awarding Cumulus Points, the thresholds concerning their accumulation etc. On the other hand, the semantics of CP are sufficiently clear to allow the mapping of such thresholds and values between different ISPs.
- **Technically feasible:** Awarding Cumulus Points most probably should be justified by respective measurement results, but it is largely up to the ISP of how to gain these data and how to interpret them. In the minimal case, a small number of very crude measurements or perhaps even estimations of user behavior may already be taken as sufficient by the ISP.

4 Design of a Generic and Modular Charging System

In existing charging systems of todays providers, setting prices, the function of charge calculation, and billing itself is integrated, even additionally combining the maintenance of service classes, user profiles, customer data, identities, and banking account data are included. Although these tasks can be distinguished clearly (cf. terminology as defined in Section 1.2), they are almost completely centralized within a single system. However, future charging systems need to be able to integrate a variety of different charging and accounting records, Call Detail Records (CDR) and Internet Protocol Detail Records (IPDR) [13], even from different communication providers or content providers, since customer's demand is determined by the so-called "one-stop billing" approach [54]. This suggests strongly to divide existing monolithic charging systems into several components with clearly defined and open interfaces. By doing this it will become possible to exchange individual components and to integrate different components supporting different technologies without having to adapt the entire system.

Therefore, the approach taken here, provides a generic and modular charging system in support of various pricing schemes applicable to different communication technologies. The initial goal is to identify relevant components and their relations to each other and to create an open system architecture, which allows charging of transport and services with different networking technologies available today, ranging from the data-orientated level tasks (such as metering) up to the money-related level tasks (such as a billing interface). Special focus will be put on the detailed design of an Internet Charge Calculation and Accounting System (ICCAS).

4.1 Tasks of an Internet Charging System

Task descriptions for a generic and modular charging system provide the basis for design guidelines of its internal components and external interfaces. Charging systems are supposed to support the following tasks:

- **Perform transport, service, and content charging**: The optimal design for an Internet Charge Calculation and Accounting System includes a combined approach for the three different levels of charging. Transport charging, sometimes termed network charging or network access charging as well, forms the basis for providing a system to deal with the transfer of data, mainly based on a general network infrastructure, such as the Internet. The service charging located on top of this level allows for the clear distinction of different services including different QoS requirements and resource consumptions. Certainly, the transport charging will be integrated into this concept and may even be completely hidden. Finally, the content charging includes the accounting tasks for information, the content, which is specifically monetary-sensitive and needs to be paid for by reading, using, or copying it. Based on the level of business interactions, it may be useful to apply content charges for certain information services only, integrating invisibly by customers all underlying transport and services charging. Transport and services include the ones provided by a variety of service providers (cf. [33] for their different definitions and distinctions), which are offered in an open market situation. This charging task needs to be as far as possible service-independent, to ensure future extensions and adaptations to yet unknown services. Content charging will be an emerging future need.
- **Perform accounting tasks according to transport and multi-service definitions**: Data gathered from the physical infrastructure and mediated due to policies, need to be accounted for. This requires as far as possible the knowledge of "sessions", "durations", or "flows". Mainly, these information are derived from data metered, such as "begin-of-session" or "end-of-flow". If such starting and end points can not be determined explicitly, heuristics need to be applied for session or flow detection purposes. In any case, the "length" of a communication relation will be recorded, if any usage-based charging approaches are to be supported. The accounting task for a single, well-known service is performed by an algorithm, which utilizes a clear service specification. In case of multi-service provisioning, these service specifications must exist concurrently and need to be maintained over some period of time. Therefore, the separation of incoming, metered data and their mapping onto the particular service in operation is essential.

- **Support different levels of security for charging information**: For every data and information, which can be mapped to monetary equivalents, these data show a certain degree of sensitivity. However, based on the dedicated level of interest, a single accounting record, a single metered routing data, or a charging record may not be a security problem, since their lifetime and validity, and therefore asset, are short. But other combinations of aggregated data, *e.g.*, flow-related information in terms of usage information, duration, billing records, and customer identification, form critical information. Therefore, appropriate security mechanisms are to be integrated.
- **Support auditing**: Communication services offered in a market environment need mechanisms which support the proof of service delivery under well-defined circumstances, mainly driven by legal requirements. Therefore, an auditing functionality need to be based on metered and accounted for data, which may be specifically restricted, structured, or stored depending on legal aspects, such as telecommunications acts.

4.2 Charging System Design

The design of a charging system requires a clear separation of concerns between different tasks. As depicted in Figure 8 and explained below, these components are based on IP routers, which are part of the ISP's network. Therefore, it has to be noted that the networking architecture itself with or without charging support remains similar, only metering needs to be integrated additionally. For advanced feedback control loops between service usage and the user, a QoS component inside the IP router is essential to perform QoS management tasks, highly depending on the IP technology in place. Therefore, the components located within the market management area of Figure 8 are inherently part of the network management architecture for economic-driven network operation and support systems.

4.2.1 Components and Operation Basics

Every provider maintains a network consisting of routers and network links between them, metering functions, mediation systems, accounting systems, and a billing system. Metering functions can be implemented in independent devices (passive metering) or integrated into routers (active metering). In either case, they generate usage data. Due to the fact that the amount of usage data is usually too large to be sensibly processed, firstly, it must go through mediation where it is transformed into a form suitable for further processing. The resulting accounting information (base accounting records) are gathered and accumulated in accounting systems. These accounting systems, in turn, forward all accumulated and perhaps abstracted accounting information through a charge calculation function towards the billing system interface. These raw usage data are not yet associated with a specific user or customer. Therefore, it is necessary to identify a customer responsible for every sequence of data units sent. This is performed within mediation, since the amount of data leaving the mediation components can be much smaller due to the aggregation of usage data. The charge calculation, which receives pricing information from the pricing component, translates the accounting information into charging records, hence, it maps the resource-oriented information from the accounting systems into monetary values. Within the charge calculation, discounting strategies, marketing-driven pricing schemes, or simply fixed prices can be applied. All these pricing and services-related information are specified in a tariff. Finally, the billing system uses these values to prepare the bills to be sent to customers. These tasks are conceptually coordinated within an overall high-level architecture and need to inter-operate as depicted in Figure 8.

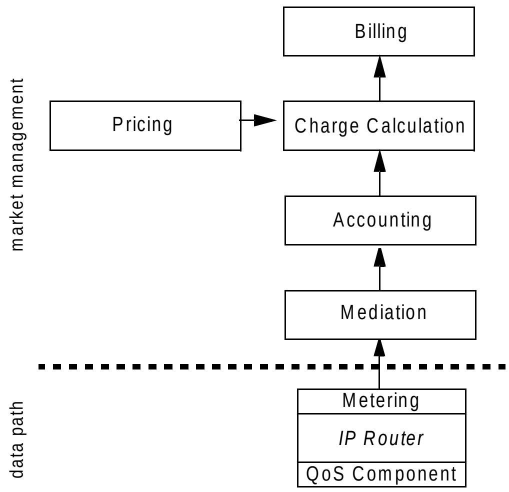

Figure 8: Concept of Charging

4.2.2 Design Dimensions

While the components for the charging system have been identified, still it must be determined how these components are implemented and deployed in possible scenarios with potentially several different ISPs. There exist three dimensions in which a charging system can vary according to the scenario (cf. Figure 9).

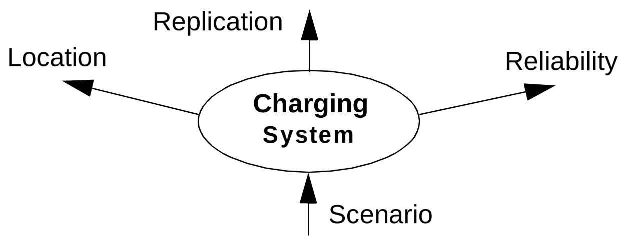

Figure 9: Charging System Dimensions

The dimension of ***location*** defines, where components are located. In particular, an "inhouse" location refers to the fact that the ISP itself hosts this component and provides the according functionality internally. The "out-sourced" location defines that this component and functionality are being performed outside the scope and administrative domain of the ISP. Mainly business case assumptions and the size of the ISP considered will determine the final location of components in a given ISP infrastructure. In addition, security-relevant questions may arise, once the out-sourcing of financial activities is intended.

The dimension ***replication*** defines how many of these components considered exist in a given environment. Mainly the number of clients served by an ISP and the number of interconnection points with peering ISPs will determine the number of replicated components required. However, besides the pure replication an important issue is the interaction between these replicated components. Appropriate protocols (open, ISP-specific, or vendor-specific) need to be selected for a suitable and correct design as well as the implementation. Open interface specifications are required to provide a chance for replication.

The dimension ***reliability*** defines how reliable these components have to be. The needed degree of reliability depends only indirectly on the ISP type. It rather depends on the afore mentioned other dimensions of location and replication as well as heavily on the component type itself. Nevertheless, the needed reliability of components is a dimension in which a specific charging system can differ from others.

When deploying a charging system for a specific ISP, the ISP type defines a set of different choices based on the distinction of roles for Access ISPs and Core ISPs [33]. Depending on the ISP type the location, the replication, and the reliability of components will determine suitable and less useful combinations of these components. However, currently there is no general set of criteria available depicting the optimal location and replication of components for a given ISP scenario. Nevertheless, work on cost modeling of ISPs may determine a suitable design process for this aspect. A sample for two neighboring ISP is given in Figure 10, where essentially every router is equipped with metering components, three accounting systems operate per ISP, two charge calculation instances are available, and a single billing system per ISP are drawn. This example assumes that these two ISPs are peered Access ISPs with a evenly distributed customer topology, since every router acts as an access router for users. However the number of users differs per access router, since the lowest router of ISP B requires to operate a single accounting and charge calculation system for a single router.

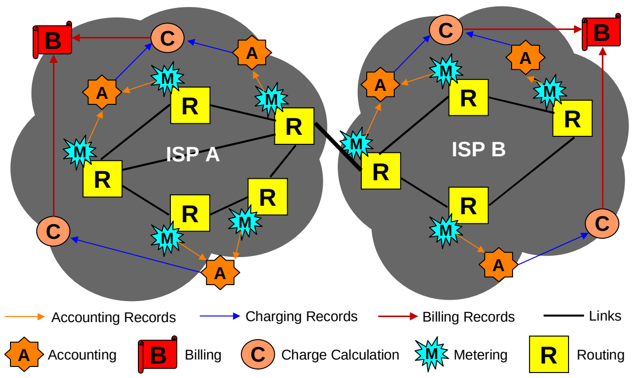

Figure 10: Possible Location and Replication of Charging Components in Provider Networks based on [27]

4.2.3 Time-scales

Besides these design dimensions discussed above, time-scales define the important dynamic criteria for feedback handling as discussed within the pricing model design (cf. Section 3.2). According to [24], overall management time-scales are distinguished as follows: Short-term in minutes, medium-term in hours, and long-term in weeks or months. These scales are extended by an atomic scale for ultra-short times in seconds and below to allow for the definition of potential feedback in a round trip-time or milliseconds. As summarized in Table 3, measurement intervals and units of measurement show the relevant timing data and information to be accounted for. The type of feedback is identified as well by describing its major content.

Table 3: Time-scales, Measurement, and Feedback Content

Time-scale Naming	Measurement Intervals	Measurement Units	Feedback Content
atomic	milliseconds, round trip-times	packets	communication-relevant data
short-term	minutes	flows/sessions	application data
medium-term	hours/days	billing periods	billing data
long-term	weeks/months	contract periods	contract data

Applying these time-scales onto pricing-controlled activities results in (1) the atomic monitoring and control level, (2) the short-term intervention level, (3) the medium-term service provisioning level, and (4) the long-term business/strategic level. Therefore as drawn in Figure 3, the proposed charging system operates as a management system being capable of supporting various pricing schemes. They include, *e.g.*, auctions and congestion pricing, as a means to intervene the current network utilization, and CPS, which is part of the strategic level, since different business models are developed for different user demand and user segments Any of these medium and long-term data are contracted in Service Level Agreements, including for todays Internet contracts only a minority of technical specifications, but large parts of business and legal data (cf Section 5.1 below).

4.3 Charging System and ICCAS Architecture

Understanding the charging system tasks, their clear separation, and the design dimensions as well as time-scales, an overall architecture for a charging systems and its important subsystem, an Internet Charge Calculation and Accounting Systems (ICCAS), is based on the mapping of the conceptual view onto distinguished components. These components are developed to specify the architecture of the charging system and equipped with interfaces and interaction paradigms.

4.3.1 Charging System Provider Interactions

Refining the concept into components and their interactions and interfaces results in the overall architecture as depicted in Figure 11 as it has been designed for the M3I project [41].[5] While important tasks of these components are discussed later in more detail in Section 4.3.3, interactions between two neighboring providers take place on two levels.

The first one is of course on the data path since providers must exchange data between their networks. Further inter-provider information exchange happens as part of the specific protocol processing as defined in the QoS model applied, *e.g.*, for resource reservation purposes such as using the Resource Reservation Protocol (RSVP) or inter-Band-

5. Special thanks are addressed to the Technische Universität Darmstadt, namely M. Karsten, J. Schmitt, O. Heckmann for the joint design of the full M3I pricing, charging, and accounting view, cf. [20] and [29] as well.

width-Broker communication, where messages are exchanged between border routers of neighboring providers. In these cases, a type of signaling or consolidation protocol has to take care of the distributed information scattered around in the network.

Since the transport of this data is not for free, ISPs will charge each other for the data transported. This leads directly to the second level of interaction. Each provider collects information of the amount of data transported and calculates a charge for it. He sends a bill through a billing system to the responsible neighboring ISP's entity. Either he bills another provider, if he provided a service for him or he sends the bill to one of his customers. So information exchange between providers occurs on the level of billing systems, where inter-provider invoices are exchanged. Instead of performing absolute billing between interconnected providers, they can also offset their claims against each other. A set of peering agreements and settlement schemes exist for today's ISPs, however, they are defined in a quite static manner and do not allow for immediate responses to bandwidth bottlenecks or further customer and user demands.

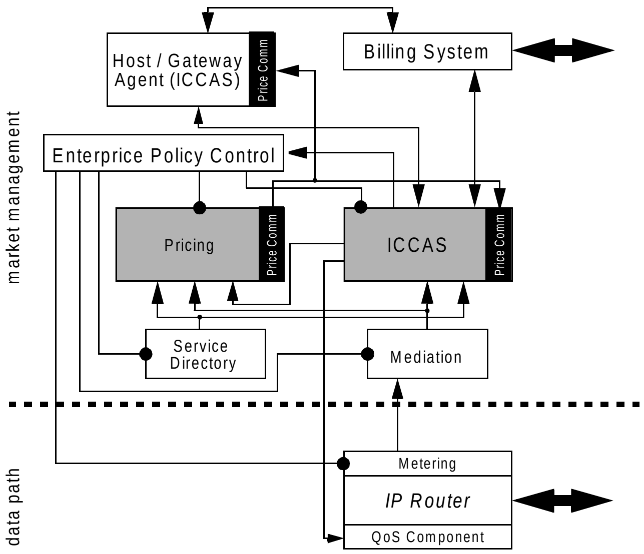

Figure 11: Architecture of the Overall Charging System

4.3.2 ICCAS External Components

For describing the charging system completely, an outside-first approach has been taken, which illustrates all components external for the ICCAS itself first. ICCAS details follow in Section 4.3.3 afterwards.

Metering

As shown in Figure 11, Metering is integrated in the IP router. Alternatively it could be placed directly on the wire. Indeed, such a solution introduces supplementary expenditures, *e.g.*, an entity needs its own IP address[6] or requires special protocols. Furthermore, it can only monitor the actual usage of the link and has no knowledge of usage of any critical resources relevant for congestion control within the router. Therefore, the interconnection of several metering units to reconstruct the current router status is not a feasible. Concluding from this, it would be necessary, in spite of having metering units on the wire, to know the state of the router, so an explicit interaction of the charging system and the router would be required anyway.

Mediation

In addition, the purpose of the mediation entity is to transform metered data (of each single meter), to merge data of different meters, and to reduce the amount of data metered. The mediation's effect of dramatically reducing the amount of data has a deep impact on the charging network design.

Pricing

For calculating charges of transmitted data, prices are important. These prices are applied in the charge calculation component (see Section 4.3.3) but they are not set at this point. There are many different ways to set prices and, therefore, a separate pricing component performs this task. It can make use of economic models or just use fixed prices set by hand. For dynamic pricing models most often an input from metering is needed, since the amount of data traffic influences prices.

Billing

Major assumptions on functions of a billing system are to be made. It affects the charging system and the ICCAS in particular in so far that interfaces need to be placed correctly and that plausible and complete scenarios can be created. It is used to describe inter-provider charging concerns as mentioned above.

Enterprice Policy Control

The enterprise policy control entity represents the ISP's interface for the management and supervision of all (except the billing system) ICCAS related entities. It covers and controls the ISP's business strategy with respect to its implementation and configuration of the given networking equipment.

Host/Gateway Agent

The host/gateway agent performs two different functions. The first one is to communicate charges to hosts (users) and gateways to provide an optional feedback channel for their service usage. In this case, the host agent acts on behalf of the user. This can include the negotiation of services with the ICCAS, an automatic reaction to communicated charges, or even payment information. A host agent can also restrict user's options, when the customer in control of these users wants to restrict the behavior of his users he pays for. In particular, this is the case for companies, where the company is in the roles of the customer of an IPS and the employees act as users of the services offered.

Service Directory and QoS Component

Finally, since in general a user tends to lack a complete understanding of Quality-of-Service (QoS) in technical terms, he will be unable to specify detailed requirements in a way that can be used as a direct input to the QoS component within the router. Instead the user has a higher, application level view of quality. This view must be translated into technical values, which can be used for setting parameters in QoS components and for charging according to according technical usage data. Therefore, this translation takes place in the Service Directory.

6. The IP address is required, since the metering entity needs to be configured according to an ISP's enterprise policy.

4.3.3 ICCAS Architecture and Internal Components

After having discussed external components of the ICCAS, its internal entities consists 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 these components can be physically distributed as discussed above. Embedding the ICCAS into the overall charging system is provided through eight distinct interfaces, which are described in Section 4.4 below.

Concerning the flow of data within the ICCAS, internally, the ICCAS has been divided in two logical paths as shown in Figure 12. On one hand, the Accounting Information Path (AIP), depicts the flow of pure charging-relevant data. On the other hand, the Control Policy Path (CPP) is used to manage and configure the ICCAS, especially all entities involved with the processing of the charging data. These two paths differ mainly in the order and direction they process data. The AIP starts from the bottom of the graph (taking raw technical data) by processing metered and mediated data as well as pricing information. It ends on the top of graph, where complete charging records are handed over to the billing system. In contrast, the CPP starts from the top of the graph (taking business-related data) by receiving enterprise policy control information and processes down to the bottom of the graph, resulting in QoS control data to be handed over to the underlying router and optionally in agent.

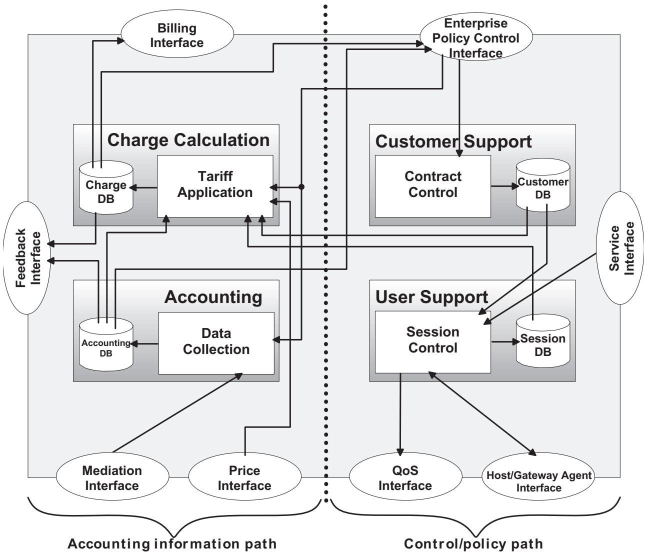

Figure 12: ICCAS Architecture in Detail

Accounting

The accounting component receives all metered and mediated usage data and is responsible for storing it. It must provide these stored data to other ICCAS components and interfaces for further processing, feedback, or statistic evaluation. Accounting is the central usage data storage component.

Charge Calculation

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.

Customer Support

An ISP may have many and a large number of different customers. Additionally, a customer is not the same as a user, *e.g.*, one customer might pay the bills of several users. Therefore, a customer defines the role, who negotiated a contract with the ISP. The content of this contract, *e.g.*, number of users covered by the contract and their names and accounts, are managed within the customer support.

User Support

While the customer support component is responsible for keeping all contract information the user support component is responsible for making sure that those contracts are kept. On one hand, this means that he blocks any user requests that are not covered by the contract the user (more exactly: the customer, who pays for the user) has. On the other hand, he must make sure that a service requested by a user is delivered to him, if the contract allows it.

4.4 ICCAS Interfaces

In general, all interfaces between the components described are designed to act (1) as protocols, allowing for the communication between two remote entities of components, or (2) as software interfaces, reflecting the clear architectural decision, that the interaction between those components happens within a common address space.

4.4.1 Interface Overview

The following list of interfaces summarizes all interfaces required, identifies its application path (Accounting Information Path, AIP or Control Policy Path, CPP), and outlines their main functionalities:

- **QoS Interface (CPP):** To provide services to the customer it is necessary to control the QoS component of routers. This interface can be used to set QoS parameters of routers, depending on the technology in place.
- **Mediation Interface (AIP):** The mediation interface is responsible for collecting data from several mediation entities, possibly even from mediation entities of different types.
- **Enterprise Policy Control Interface (CPP):** This is the interface for changing parameters of the ICCAS after the system has been deployed. By using this interface the enterprise policy control can install new services or request and receive charging or accounting data.
- **Service Interface (CPP):** This interface can be used to read service definitions out of the service directory.
- **Billing Interface (AIP):** This interface is responsible for sending calculated charging records to the billing system.
- **Pricing Interface (AIP):** Pricing is responsible for setting the prices used by the charge calculation component, therefore, this interface is used to send the calculated prices to the charge calculation component.
- **Feedback Interface (AIP):** To set suitable prices the pricing component uses price models with various input variables. Some price models need usage or charge information as input variables, hence, these can be communicated to the pricing component via this interface.

- **Host/Gateway Agent Interface (CPP):** This interface is responsible for the optional communication with the customer. Mainly, this includes the selection of services the customer can use or the transfer of a feedback signal form the service provider to a user. This interface is open for future enhancements.

4.4.2 Interfaces of the Accounting Information Path

To allow for a complete understanding of the Accounting Information Path, all relevant interfaces are discussed.

Interface: Mediation-Accounting (I-MA)

The usage data, which have been mediated after the data gathering took place, needs to be transferred to the ICCAS. Therefore, a protocol is designed, which defines rules and transmission units for transferring mediated data to the accounting component. Since the anticipated load for this interface will be high, the protocol must be highly efficient, yet extensible. The data exchanged across this interface (cf. Figure 13) will include one of the following alternatives, which depend on the particular scenario:

- A simple hand-over of data gathered by metering.
- A hand-over of data mediated based on the particular inputs from the enterprise policy control. This may result in the dedicated specification of specialized data to be required for the ICCAS, some special aggregation of these data, or even the neglecting of data resulting from the gathering process.

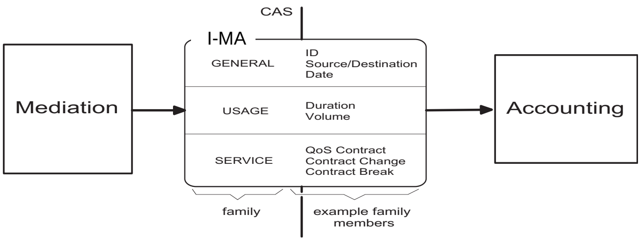

Figure 13: Mediation-Accounting Interface (I-MA)

Due to the fact that the type of data for I-MA and the Interface Mediation-Pricing are similar, only differing in its size, the number of data records, or its frequency of exchange, a similar type of protocol is defined.

Interface: Charge Calculation-Pricing (I-CP)

The ICCAS requires input from the pricing component, in particular details on prices and tariffs for services. A generic way to exchange this price information between different participants is defined, allowing for the widest scope of information, yet efficient transmission. Charge calculation-relevant details for this protocol encompass the following ones:

- Memory-efficient data structure for prices and tariffs.
- Processing-efficient data structures (context-free) for prices and tariffs.
- Inclusion of customer and/or user identification information for prices and tariffs.

Since the pricing does not reflect a full and self-existing component as such, rather a well-defined means for communicating prices and tariffs across networks and between components, the software-based interfaces between the "hosting" component and the pricing component are of local matter for these components only.

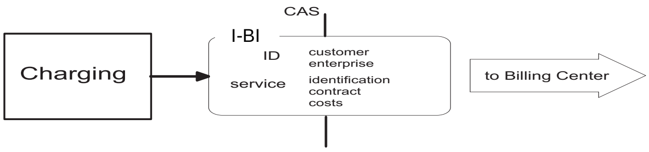

Figure 16: Billing-ICCAS Interface (I-BI)

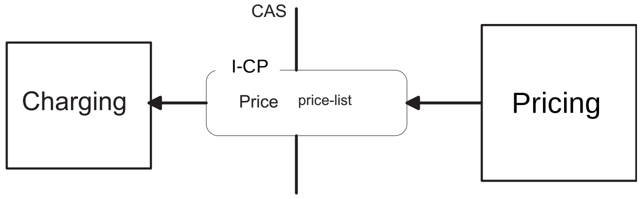

Figure 14: Charge Calculation Pricing Interface (I-CP)

Interface: Feedback-ICCAS (I-FI)

The feedback interface provides the necessary input for the pricing component to adapt and change prices according to the current network status. The feedback represents the status of the network in terms of resource usage, i.e. quantitative usage given by the amount and characteristic of the accounting and charging records, as well as in terms of the current charging and accounting policy.

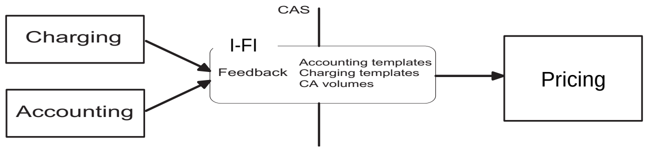

Figure 15: Feedback-ICCAS Interface (I-FI)

Interface: Billing-ICCAS (I-BI)

This interface between the ICCAS and a billing component provides the access of billing institutions. It supports facilities to clearly identify customers and enterprises in financial terms and is responsible for accessing economic accounting tasks. Furthermore, it identifies the offered service. This is done by giving it an ID (corresponding to the one in the charging data base) and by attaching all requested for characteristics of the service. The service contract and the costs are as well passed to the interface. The contract can be used by the billing component to legitimate and detail a list of costs.

Further details on interfaces of the Control Policy Path can be obtained from [20].

5 Results and Evaluation

The integration of various pricing schemes into the developed ICCAS can be performed. This section discusses the integration of the pricing scheme CPS into the ICCAS based on a DiffServ networking environment, since both concepts are well suited to charge differentiated data transmissions in a DiffServ network. They are combined in a way which does not lead to a large protocol overhead and which is feasible in small as well as in large networks, thus scaling well.

5.1 CPS in the Context of DiffServ SLAs

CPS is tailored to complement DiffServ in the provisioning of QoS and in the appeasement of customers. According to the DiffServ philosophy, CPS ties up at the edges of the network, i.e. at the border routers of any DiffServ domain. Consequently, CPS policies, which should be applied and enforced, need to be defined at and for edges of the network domain. Since DiffServ already uses SLAs at the edges to comply with service provisioning tasks, CPS policy definitions and agreements preferably ought to be specified in the same SLAs. Merging service provisioning tasks with service charging tasks is appropriate and economical, since both apply on the same service, a similar granularity, and at identical locations in the network. Therefore, an SLA includes besides legal commitments financial commitments, *e.g.*, the sum of financial forfeits on contract corruption, both of which not concerned with the charging task directly.

The process to set up an SLA is sketched in a generic fashion in Figure 17. The customer and the ISP initiate a negotiation phase in which technical requirements and possibilities are assessed, warranties on QoS and performance are given, the processes to verify them are specified, and in which charges are set, i.e. for CPS the price is fixed on the foundation of the customer's traffic volume estimation. The negotiation between the customer and the ISP is vulnerable to subtle inaccuracy, since it does not yet and not only handle technical topics, which are expressed in precise and reliable parameters or mathematical formulas. Thus it is indispensable for the customer as well as for the ISP to consult with an legal adviser. Concluding, it is up to the customer, the ISP, and a legal adviser to define the SLA.

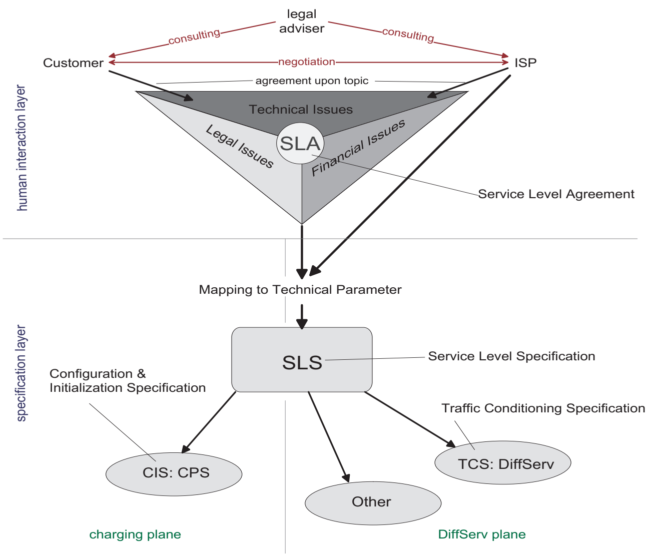

Figure 17: SLA of a DiffServ-CPS System.

Because SLAs do not already define technical parameters, ISPs have a lot of freedom to map the SLA contract to technical parameters. Furthermore, this freedom to comply with the contract allows to separate concerns of proper DiffServ technology from those of the pricing, CPS, although they are specified in the same SLA. Figure 17 illustrates this fact by the *specification layer*, which is separated in a *charging plane* and a *DiffServ plane*. Concerns of charging are located in the *charging plane*. Likewise, this generic representation allows to locate and specify parameters of charging systems and pricing schemes, other than CPS. Parameters necessary to be specified within an SLS (Service Level Specification) include configuration and initialization parameters. In the case of CPS these cover, *e.g.*, threshold values for collecting cumuli points, prices associated to cumuli points, or refresh periods.

5.2 CPS supported by the Charging System

Contracting between customers and a provider founds the basis for connectivity. In case of a DiffServ networking environment, SLAs and their handling by the charging system supports end-to-end services.

5.2.1 End-to-end Services

The main aspect of the DiffServ Architecture to concentrate network's and system's intelligence at the network edges supports an approach to meter at network edges only. Transmissions can than be charged for according to their associated Behavior Aggregate (BA). However, an important question remains, how can end-to-end services be accounted and charged for?

Each user has established a contract with his service provider. In a DiffServ world this contract is represented by the before mentioned SLA (cf. Section 2.3). This SLA determines, which services the customer may use and how he will be charged for them using the applied pricing scheme, here CPS. However, the customer has only established an SLA with his provider, not with other domain operators, who play an important role in his data transmissions as well. Therefore, to guarantee the provision of a service not only for one domain, but for the entire transmission path, a flow establishment phase is needed (cf. Figure 18). This setup can be performed by the User Support component of the ICCAS.

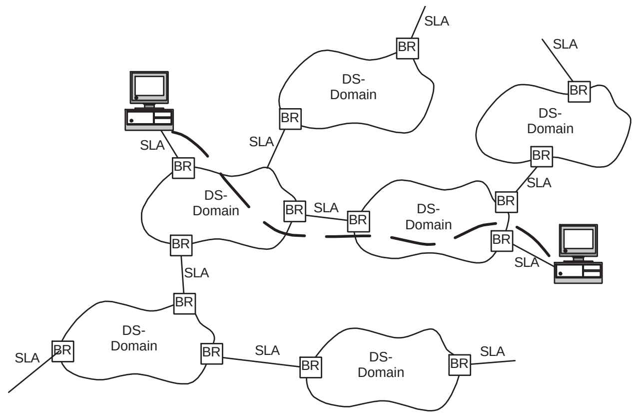

Figure 18: Different SLAs in a DiffServ Network and End-to-end Connections

This approach does not return to micro-flow orientation, it supports reservations for BAs, aggregated super-flows containing many micro-flows, by bandwidth brokers of the DiffServ network. A settlement between different domains also does not take place on a micro-flow basis, but only for BAs. Between domains SLAs exist also, which are negotiated on the basis of estimations domain operators have of their traffic, which originally is caused by the traffic of end users. The specific settlement between two neighboring domains is also performed applying the similar CPS scheme.

5.2.2 Interfaces of a DiffServ-CPS Charging System

The interfaces of the Accounting Information Path are considered at this stage to explain the use of CPS and DiffServ within the ICCAS. As illustrated in Figure 19, a single ICCAS for a DiffServ domain supervises and configures four Border Routers (BR). Including the billing system, the ICCAS has to support to protocols with remote entities, i.e., I-MA and I-BI as discussed below.

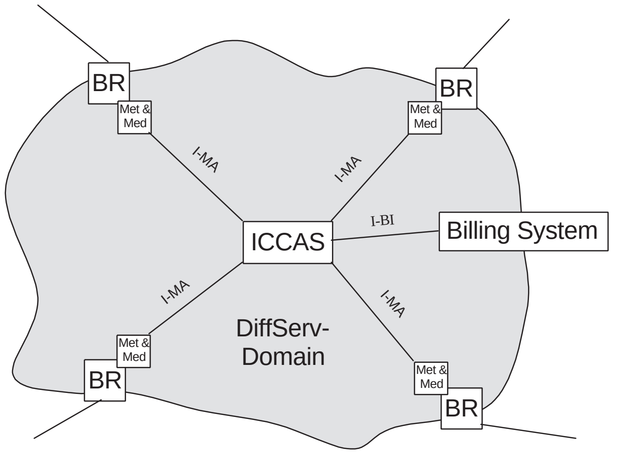

Figure 19: Transmission of Charging Data in a DiffServ Domain

- **Mediation-Accounting interface**: This interface is small. It is used to collect all metered usage data of each BA of each border router, where mediation components are located (cf. Figure 19). As shown in the message sequence chart in Figure 20 this data can either be transmitted periodically or upon request from the ICCAS. To set parameters in mediation components, *e.g.*, the interval for automatic transmission of metering data, the ICCAS must also be able to configure mediation components remotely.
- **Price Interface**: The price interface is not needed since dynamic prices do not exist by applying the CPS scheme. Prices changes on a longer time-scale are reflected by appropriate changes of SLAs themselves. Therefore, prices for data transmissions can easily be queried from the customer support inside the ICCAS (cf. Section 4.3.3) where the SLA information is stored in the contract control.

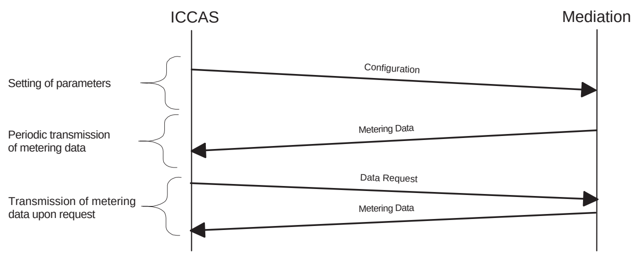

Figure 20: I-MA Message Sequence Chart

- **Feedback interface**: This interface is also not required in the case of DiffServ and CPS, since the only feedback a customer receives are cumulus points at the end of each billing period, which he receives via a bill. However, it would be possible to offer the user the option to query his current cumulus points' standing. Though, it appears to be more sensible for the user to measure his data traffic himself and, consequently, calculating his cumulus points using his SLA and comparing the result to the bill. This is one of the major advantages of the proposed integration of CPS into a DiffServ environment: The user can control his bills exactly and will immediately recognize any mistakes made by his provider. Measurement and calculation on the user's side will be performed automatically.
- **Billing interface**: The billing interface needs to receive not only charges for each customer to create a bill, but also all cumulus points for a current billing period. The numbers of cumulus points to be sent are calculated in the charge calculation component.

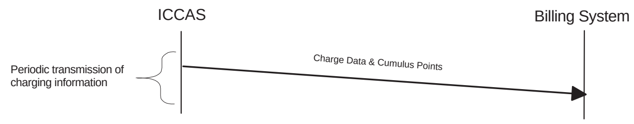

Figure 21: I-BI Message Sequence Chart

5.3 Exchange Data for Interfaces and Protocols

Based on the description of the interactions between all charging system components, overheads for charging information can be identified. Taking the Accounting Information Path (AIP) into closer considerations, interfaces are of local interest or data is exchanged between remote distances.

The Mediation-Accounting interface (I-MA) serves as an example for the need of an exchange protocol. As depicted in Table 4, three families of data are distinguished. General data contains informations such as identifications, originators, and date. Usage data includes duration and volume data. Finally, service-relevant data encompasses references on contracts and their handling. According to the message sequence chart in Figure 20, configuration data to indicate the level of aggregation are transferred once to the mediation component. Data requests follow the pop scheme, and will result in the sending of metering data. They may be sent periodically.

Table 4: Message Content for I-MA

Family	Member	Size	Description
General	ID	4 Byte	Identification of customer and session (or service).
	Source/Destination	32 Byte	Originator and Consumer identification.
	Date	4 Byte	Date of service/session allocation.
Usage	Duration	4 Byte	Duration of service/session use.
	Volume	4 Byte	Effective bandwidth consumption.
Service	QoS Contract	sizeof(Q)	Agreements on jitter, loss rate, delay.
	Contract Change	1 Byte	Customer initiated renegotiation.
	Contract Break	1 Byte	Notification of network failures.

Assuming a direct Internet Protocol (IP) encapsulation of the I-MA interface protocol data units, the configuration messages contain, *e.g.*, per usage-based service to be monitored 20 Byte IP header, its ID (4 Byte), the source and destination (64 Byte), and the date (4 Byte), in total 72 Byte. Periodic metering data include the IP header, the ID (4 Byte), the duration, if the flow has been finalized (4 Byte), and the volume (4 Byte), totalling to 32 Byte packets. Depending on the accounting period P_A and the number of services N_S to be mediated and accounted for, the total bandwidth $B_{I\text{-}MA}$ of I-MA for periodic data collected equals to:

$$B_{I\text{-}MA} \text{ [bit/s]} = 1/P_A * N_S * 32 * 8 \text{ [1/s * bit]}.$$

To continue with AIP interfaces, the Charge Calculation-Pricing interface (I-CP) contains a single price list only (cf. Table 5), which is distributed to the relevant ICCAS as often as the enterprice policy component decides on changes on the current prices per unit.

Table 5: Message Content for I-CP

Family	Member	Size	Description
Pricing	Price-list	sizeof(Pl)	List of price/unit, *e.g.*, per volume, per jitter, delay constraints. Prices may vary, *e.g.*, depending on time of day or congestion.

Considering the CPS over DiffServ scenario, the Feeback interface (I-FI) is not required, however, it may contain a number of different families of data, such as accounting and charging templates or volumes for quantitative feedback information (cf. Table 6).

Table 6: Message Content for I-FI

Family	Description
Accounting templates	Information about the characteristic of the Accounting Records.
Charging templates	Information about the tariffs and the resulting Charging Records.
Charging and accounting volumes	Quantitative feedback about accounting and charging.

Finally, to complete with the AIP interfaces, the Billing interface (I-BI) transports data from the ICCAS to the billing system in place. As illustrated in Figure 21, the billing instantiation requires for the send operation of service descriptions (4 Byte) and once its technical definition in terms of QoS parameters (sizeof(Q) Byte). Periodic billing records are transferred including the current service ID (4 Byte), the charge (3 Byte), and the monetary unit of the charge (1 Byte), *e.g.*, CHF or US$ Future enhancements include the integration of authentication means.

Table 7: Message Content for I-BI

Family	Member	Size	Description
Description	ID	4 Byte	General description of a service as seen by users.
Definition	QoS list	sizeof(Q)	The technical definition of a service in terms of QoS parameters.

Table 7: Message Content for I-BI

Family	Member	Size	Description
Billing Data	Charge	3 Byte	Charging value of the service utilized.
	Monetary	1 Byte	Monetary charging unit.

Depending on the particular billing period P_B, the total bandwidth $B_{I\text{-}BI}$ of I-BI for periodic billing record transmissions equals to:

$$B_{I\text{-}BI} \text{ [bit/s]} = 1/P_B * 8 * 8 \text{ [1/s * bit]}.$$

5.4 Example Discussion and Technology Choices

The presentation of a generic and modular charging system as well as a suitable and technically feasible pricing scheme have been studied so far, but why should a user use different service classes at all? The reason is that no user is like another one. They all differ in their personal judgement of QoS, in their personal satisfaction with a given QoS, and their various applications with different QoS requirements. Therefore, users will utilize different QoS classes, when offered and priced suitably. Probably, users will not use similar service classes all the time, but change service classes depending on the application they use, on the current QoS they get, and of course on their current mood.

For instance, a professional Internet user transmitting high quality real-time video streams over the Internet shows very specific requirements. He also has an estimation of how often he uses the web for such transmissions and can specify these requirements in an aggregate level, *e.g.*, within an SLA he has negotiated with his provider. However and most certain, these estimations will not be correct always. When the user is on vacation, he will produce less or no traffic at all and, therefore, he will receive ‘green’ cumulus points. If, however, he is working on a project, which involves more video transmissions than usual, he will collect ‘red’ cumulus points. If the amount of his traffic misses his estimations frequently, his contract will be adapted as shown in Section 3.

Of course, traffic of the above mentioned user in most cases does not sink in his access provider’s domain, but continuos through to other domains. Just as the single user has an estimation of his data traffic, the provider is able to do a similar estimation based on the sum of all his customers. These customers can be either users or other providers, who send data through his domain. There is a high probability that the error in users’ estimations normally will not be much noticed in the traffic of providers, since too low and too high estimations will even up. Furthermore, it is the ISP’s policy to meter heavy users and change their contract and SLA in time However, certain events might occur, when the estimation of a provider’s traffic will be wrong. For example, if olympic games take place, many users will utilize the Internet to access current results and background information. People will access the pages, journalists will send their articles to their newspapers or their video reports to their broadcasting company. The Internet might even be used for live reports. The olympic games are of course a very extreme example, where any network infrastructure will get into serious trouble, if no further capacities are planned or provisioned. However, it serves very well to demonstrate, what will also happen on a smaller scale with other events.

Finally, existing technology needs to be evaluated and a suitable solution for the ICCAS implementation to be designed. There are several products available, which can fulfill a subset of functions that are required for the designed charging system, while another important subset of functions has to be implemented.

For the metering component, a close networking hardware dependency is observed. *E.g.*, Cisco's NetFlow product is able to collect and transmit flow data, if they are metered at Cisco routers. An open software solution which performs metering functions for a desktop computer is NeTraMet [7]. To handle information gained by metering and process mediation tasks, the Smart Internet Usage (SIU) product by Hewlett Packard can be used. SIU is able to access NetFlow routers directly and can be adapted to access a NeTraMet meter via SNMP. One of the major restrictions imposed by these products is the non-existence of a real-time capability. Data can be measured only in the time-scale of seconds or even minutes. However, this is not a problem with the proposed integration of ICCAS, DiffServ, and CPS. While the CPS has been designed primarily with the user in mind, end-users and providers, it has also the major advantage to be implementable with technology available today. Products mentioned fulfill these requirements of metering and accounting of data traffic already. Therefore, time-scales are not a problem in case of CPS.

5.5 Flat Rate versus Usage-sensitive Pricing: Practical Considerations

When on December 1, 1996 AOL switched its pricing scheme from a usage-sensitive one (a US$9.95 monthly fee including 5 hours and US$2.95 for each additional hour) to a pure flat rate (i.e. US$19.95 per month for unlimited access), the result has been the well-known demand explosion that AOL initially could not cope with at all [40]. Hence, due to a couple of lost law-suits, AOL was forced to offer extensive refunding to unsatisfied users, but finally managed to cope with the situation and succeeded in becoming the most important player in the ISP world of the US. Breathe Freely, a UK-based ISP, has been less lucky. After introducing in May 2000 an unlimited access scheme for an on-off payment of 50£, the apparent lack of capacity forced the ISP recently to skim off the most internet-use intensive 1% of its users, because they apparently caused real problems for the rest [39].

These two examples are two of the most prominent ones concerning ISP troubles, when offering free unmetered access to the Internet. Without any doubt, flat rate pricing schemes are still the most popular ones, and hence a justified point to start while designing a new pricing scheme. The only thing to be added in order to prevent disasters like the mentioned ones basically consists of allowing at least a defined concept of feedback to customers on their current usage patterns and their compliance to the overall network situation. That is exactly what the proposed Cumulus Pricing Scheme CPS is offering. Its feedback mechanism is extremely flexible in terms of measuring data, very modest in terms of consequences for cases of abuse, and very easy to be understood as an early warning system.

Using CPS, the AOL customers would have been forced to declare their expected requirements, which would have had at least a psychological effect even during the transition phase. Moreover, to most customers their own over-utilization would have been the obvious reason for the service degradation they experienced during the first months of the new scheme. In the case of Breathe Free, one of the most emphatic complaints concerned the fact that the "heavy users" have not been made aware of their behavior and their possible consequences. Together with the fact that there must have been at least some metering in order to find out the most intensive 1% of the users, this points directly to a solution like the one proposed by CPS for the first time.

6 Summary and Conclusions

Todays network management functionalities have been extended within the work presented by a dedicated support of business-oriented network management aspects, in particular the supply of charging and pricing functions for differentiated services. Its necessity is driven by world of commercialization in networking and the demand for QoS-based services, but minimized technical effort for QoS provisioning in networks. The development of a discrete, predictable, transparent, and technically feasible pricing scheme termed Cumulus Pricing Scheme (CPS) and the generic and modular ingredients of an Internet Charging System, in particular the Internet Charge Calculation and Accounting System (ICCAS) illustrates a feasible technical solution to the support of pricing differentiated services.

The lessons learned from this work clearly indicate that the proposed pricing scheme would have avoided a recent failure of an all-you-can-eat offer as seen from a U.K.-based Internet Service Provider (ISP). In addition, a quite simple set of measurement policies, negotiable between the customer and the ISP, are sufficient and can be supported effectively by a charging system. Specifically, the ICCAS is suited for handling an extensible variety of different services, where particularly the CPS has been integrated without problems. In addition, all ICCAS interactions are functionally complete and sufficient for todays known pricing schemes, ranging from usage-based to flat-fee approaches.

The work in progress includes the ongoing implementation of the full ICCAS functionality. The optimization and tuning of interfaces and protocols will follow based on the first measurements taken, according to the measured processing overhead for the ICCAS. The simulation of the market-managed feature of CPS with multiple ISPs is under its way and will be extended to inter-ISP scenarios. This includes the simulation of various ISP policies and effects on prices, user behavior, and

Acknowledgements

This work has been performed partially in the framework of the EU IST project Market Managed Multi-service Internet (M3I, IST-1999-11429), where ETH Zürich has been funded by the Swiss Bundesministerium für Bildung and

Wissenschaft, Bern and partially in the ANAISOFT project funded by the Swiss National Science Foundation, Bern.

Special thanks go to M. Karsten, J. Schmitt, and O. Heckmann at the Technische Universität Darmstadt, Germany for the joint design of the overall architecture within M3I. Additionally, the authors like to acknowledge many discussions with all of their project partners.

References

[1] F. Baumgartner, T. Braun, P. Habegger: *Differentiated Services: A New Approach for Quality of Service in the Internet;* In: H. van As (edt.): *High Performance Networking,* Kluwer, Dordrecht, The Netherlands, 1998, pp. 255-274.

[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: On Overview;* Internet Engineering Task Force, RFC 1633, June 1994.

[4] R. Braden, L. Zhang, S. Berson, S. Herzog, S. Jamin: *Resource Reservation Protocol (RSVP);* Internet Engineering Task Force, RFC 2205, September 1997.

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

[6] N. Brownlee: *Internet Pricing in Practice;* In: L. McKnight, J. Bailey (edts.) *Internet Economics;* MIT Press, Cambridge, Massachusetts, U.S.A, 1997, pp 77-90.

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

[8] K. Chu: *User Reactions to Flat Rate Options under Time Charges with Differentiated Quality of Access: Preliminary Results from INDEX;* International Workshop on Internet Service Quality Economics, MIT, Cambridge, Massachusetts, U.S.A., December 2-3, 1999.

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

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

[11] D. Clark, W. Fang: *Explicit Allocation of Best-Effort Packet Delivery Service;* IEEE/ACM Transaction on Networking, Vol. 6, No. 4, August 1998.

[12] 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.

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

[14] 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.

[15] 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.

[16] R. Edell, P. P. Varaiya: *Providing Internet Access: What we learn from the INDEX trial;* Keynote Talk at Infocom'99 New York, INDEX Project Report #99-010W. URL: http://www.index.berkeley.edu/99-010W.

[17] 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.

[18] D. Ferrari, L. Delgrossi: *Charging for QoS;* Invited Talk, 9th International Workshop on Network and Operating System Support for Digital Audio and Video (NOSSDAV'99), London, 1999.

[19] P. Flury, P. Reichl, J. Gerke, B. Stiller: *Charging Considerations for Virtual Private DiffServ Networks;* TIK-Report No. 94, ETH Zürich, Switzerland, August 2000.

[20] J. Gerke, P. Flury, B. Stiller: *The Design of a Charging and Accounting System for the Internet;* TIK-Report No. 91, Computer Engineering and Networks Laboratory, ETH Zürich, Switzerland, August 2000.

[21] D. Grossman: *New Terminology for Diffserv*; Internet Draft, work in progress, August 2000.

[22] M. Günter, I. Khalil, T. Braun: *An Architecture for Managing QoS-enabled VPNs over the Internet;* 24 IEEE Annual Conference on Local Area Networks (LCN'99), Boston, Massachusetts, U.S.A., October 18-20, 1999, pp 122-131.

[23] 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.

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

[25] J. Heinanen, F. Baker, W. Weiss, J. Wroclawski: *Assured Forwarding PHB Group;* Internet Engineering Task Force, work in progress, June 1999.

[26] V. Jacobson, K. Nichols, K. Poduri: *An Expedited Forwarding PHB;* Internet Engineering Task Force, RFC 2598, June 1999.

[27] M. Karsten, J. Schmitt, B. Stiller, L. Wolf: *Charging for Packet-switched Network Communications - Motivation and Overview;* Computer Communications, Vol. 23, No. 3, February 1, 2000, pp. 290-302.

[28] 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.

[29] M. Karsten (Edt.): *Pricing Mechanisms Design;* M3I Project Deliverable 3, June 30, 2000.

[30] F. Kelly: *Charging and Accounting for Bursty Connections;* In: L. W. McKnigth, J. P. Bailey (edts.) *Internet Economics,* MIT Press, Massachusetts, U.S.A., 1996, pp 253-278.

[31] S. Keshav: *An Engineering Approach to Computer Networking;* Addison-Wesley, Reading Massachusetts, U.S.A., 1997.

[32] K. Kilkki et al.: *Internet Charging Reconsidered;* 5th Annual Networld and Interop Engineers Conference (Interop'98), Las Vegas, Nevada, U.S.A., May 1998.

[33] H. Kneer, U. Zurfluh, G. Dermler, B. Stiller: *A Business Model for Charging and Accounting of Internet Services;* First International Conference on Electronic Commerce and Web Technologies, Greenwich, U.K., September 4-6, 2000.

[34] C. Linnhoff-Popien, P. Reichl: *Netzmanagement;* Aachener Beiträge zur Informatik, Augustinus Verlag, Aachen, 1998.

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

[36] 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.

[37] J. MacKie-Mason, H. Varian: *Pricing the Internet;* In Public Access to the Internet, B. Kahn, J. Keller (edts.), Prentice Hall, Englewood Cliffs, New Jersey, U.S.A.,1995.

[38] M. MacLachlan: *AOL: Users Hate It, Wall Street Loves It;* CMP Techweb News, February 1998.

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

[40] A. J. Mund: *AOL's Breakdown: A Harbinger for the Internet's Future?* Cable TV and New Media - Law and Finance, Vol. 14, No. 12, February 1997.

[41] M3I: *Market Managed Multi-service Internet;* 5th Framework EU Project, IST Program, No. 11429, URL: http://www.m3i.org, January 2000.

[42] K. Nichols, V. Jacobson, L. Zhang: A *Two-bit Differentiated Services Architecture for the Internet.* RFC 2638, July 1999.

[43] A. Odlyzko: *Paris Metro Pricing for the Internet.* AT&T Labs, Research, New Jersey, U.S.A., May 1997.

[44] P. Reichl: *Kelly's Bound, RUAA and the Pricing of Multiclass Traffic in Loss Networks;* U.K. Performance Engineering Workshop, Bristol, U.K., July 1999.

[45] 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.

[46] 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

[47] 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.

[48] D. Songhurst (edt): *Charging Communication Networks - From Theory to Practice;* Elsevier, Amsterdam, The Netherlands, 1999.

[49] A. Sirbu: *Internet Billing Service Design and Prototype Implementation;* Carnegie Mellon University, Computer Science Department Technical Report, Pittsburgh, Pennsylvania, U.S.A., 1994.

[50] B. Stiller: *Pricing and Charging of Packet-based Networks;* Tutorial T12 presented at IEEE/IFIP Network Operations and Management Symposium (NOMS 2000), Honolulu, Hawaii, April 14, 2000.

[51] 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, Springer Verlag, Heidelberg, Vol. 1629, pp 281-296.

[52] 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 21-25, 1999.

[53] B. Stiller, P. Reichl, S. Leinen: *A Practical Review of Pricing and Cost Recovery for Internet Services;* to appear: NETNOMICS - Economic Research and Electronic Networking, Vol. 2, No. 4, 2000.

[54] B. Stiller, G. Fankhauser, N. Weiler, B. Plattner: *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, Session Commerce and Finance, Track 3.

[55] B. Teitelbaum, P. Chimento: *QBone Bandwidth Broker Work Group.* URL: http://qbone.ctit.utwente.nl/BBroker/, August 2000.