

CHARGING AND ACCOUNTING FOR QOS-ENHANCED IP MULTICAST

Georg Carle, Felix Hartanto, Michael Smirnov, Tanja Zseby
GMD FOKUS, Kaiserin-Augusta-Allee 31, D-10589 Berlin, Germany
{carle, hartanto, smirnov, zseby}@fokus.gmd.de

Abstract

In this paper we propose a layered framework for charging QoS-enhanced IP services with policy based configuration of its layers. As an example architecture we present our implementation - Value Added IP Charging and Accounting Service (VIPCAS). We further concentrate on VIPCAS elements for policies and data exchange. As flexible data structure for exchange of accounting information, we propose the Premium IP Network Accounting Record (PIP-NAR), which is suitable for a variety of charging schemes (e.g. charging of reserved and used resources), and which also makes provider-specific extensions possible. We specify the Tariff Formula Language (TFL) for representing tariff policies, and the Charging Information Protocol (CIP) for distributing tariff information expressed in TFL. For a fair charging of IP Multicast services, cost sharing functions should be supported at both the accounting and the charging layer. We evaluate different approaches in this area which support receiver heterogeneity with regard to the reserved QoS.

1. INTRODUCTION

In recent years there emerged an increasing number of applications such as video conferencing, distance learning, and video broadcast where data sent by a source is destined to multiple receivers. These applications also have more stringent requirements in terms of bandwidth, delay and jitter. The recent development in the Internet has been targeted to support these emerging applications. Integrated Services [33] and Differentiated Services [1] have been proposed to support service differentiation with quality of service guarantee. Multicast routing and multicast delivery have evolved from being a pure research topic and being experimentally deployed in the MBONE to being supported by major router manufacturers. Despite these developments, deployment of commercial IP multicast services is slow. One factor contributing to this is poor support for multicast charging [11].

Charging for IP multicast services introduces a number of additional problems as compared to unicast charging. One of the problems is to obtain

information on the cost of a multicast tree, and to distribute the information among multiple providers. Another problem is to share the cost of the communications among receivers which may have heterogeneous QoS requirements. In addressing these problems this paper presents a layered framework for charging of multicast IP services with QoS support and support for cost sharing between receivers.

The structure of this paper is as follows. Section 2 reviews related work, focusing on support for QoS and multicast. Section 3 presents our framework for charging and accounting for QoS enhanced IP services. Section 4 describes a charging and accounting architecture based on the framework and a flexible data structure for transferring accounting information among multiple providers. Section 5 discusses some examples in configuring the framework and introduces a tariff formula language and a charging information protocol which can be used to inform users about current tariffs. Section 6 addresses the problem of cost sharing among receivers with heterogeneous QoS and presents simulation results for different cost sharing schemes. Section 7 concludes the paper and discusses future work.

2. IP SERVICE CHARGING - STATE OF THE ART

QoS Support

Service differentiation with QoS support requires charging schemes to protect the network from abusive resource reservation. Usage-based charging also supports a funding mechanism to extend network capacity at the expense of those users that actually use these resources. Charging requirements are discussed in [22], focusing on Integrated Services, and in [2], focusing on Differentiated Services. Methods for designing charging schemes for IP services with QoS have been presented in [17].

It is well accepted that flat-rate charging for QoS-enhanced IP services has severe shortcomings [30]. Usage-based charging can be based on reserved and/or consumed resources, e.g. [17] proposed charging formula based on both reserved and consumed resources. Charging can be based on link bit rate share used by the service [22], or on resource rate and buffer space [23]. In [23], models are proposed for deriving resource costs from an IntServ flow specification. Different flow measurement methods for charging and accounting purposes are considered in [13], while different measurement granularity are discussed in [12]. Further related work can be found in [16,19,24,25,31,32].

Multicast Support

Charging for multicast services has raised several challenging tasks like the calculation of tree costs, multicast savings and cost sharing mechanisms between receivers. [8] introduces a charging and accounting architecture which supports multicast communications for IP Integrated Services over ATM. [10] suggests to use network resource cost for pricing multicast communications and derives an approximation for the cost of a multicast tree with known group size.

Group communication with payment by receivers introduces the problem of cost sharing. Additional difficulties arise for Integrated Services as each receiver can make an individual reservation. These individual reservations lead to the need of fair cost sharing techniques (see Section 6), where receivers are charged according to their reservation parameters. A scheme to share costs of a multicast tree among receivers is presented in [20,21]. The cost sharing scheme presented in [15] also allows to take distance into account.

3. QOS CHARGING: FRAMEWORK

We introduce a layered framework, the charging and accounting reference model, and an example architecture based on this framework. As part of the architecture, a data structure for exchange of accounting information has been designed. For representation of charging policies, a tariff formula language (TFL) is used. The charging information protocol (CIP) allows for distributing charging policies to charging instances at customer premises. This tariff information allows to provide online charging information to the user, and to assist users in tariff-based service selection.

The reference model classifies charging, accounting and closely related processes, and describes their interaction (see *Figure 1*). At the right, five layers are shown that encompass processing for charging and accounting. A configuration plane allows for providing configuration parameters for the processing layers.

Configuration parameters are derived from pricing policy, charging policy, accounting policy and metering policy. Since the metering layer has to provide the data needed for the charging formula, it is useful to derive the basic elements of a lower layer policy (e.g. metering) from higher layer policies (e.g. charging). These policies can be provided by interaction of dedicated policy servers with the corresponding entities of the configuration plane.

The *metering layer* provides the functions to obtain reservation information and to meter actual usage of network resources. In case of multicast, meters can be placed at the edge routers only or at multiple splitting points.

The *collecting layer* encompasses functional entities that access data provided by metering entities and forwards it for further processing to the ac-

counting layer. For supporting multicast charging, this layer is also responsible for selecting appropriate meters (meter placement). Transfer of metering data to collectors can be initiated explicitly (the collector initiates transfer of metering data) or implicitly (after a triggering event such as detection of a new flow, the meter initiates transfer of metering data to the collector(s)). Additionally, this layer may distribute collected usage data to other domains in a multi-provider environment.

Entities of the *accounting layer* process collected usage data (from collectors within the same domain and other domains) and reservation data, consolidate them based on service parameters, and create accounting data sets (i.e., accounting records) which are passed to the charging layer for pricing assignment. To support multicast charging, this layer is also responsible for reconstructing the multicast topology including splitting points where required by the cost sharing scheme.

The *charging layer* derives costs for accounting data sets based on service specific tariff parameters. Different cost metrics may be applied to the same usage of resources, and may be evaluated in parallel. A detailed evaluation of the resource usage can be used for generating bills to the customer, or for internal analysis (auditing) by the service provider. A simple evaluation of current costs can be used for displaying an estimation of accumulated costs for the service user, or for control purposes by the customer organization or by the provider. For charging multicast services, cost allocation assigns costs to specific endpoints, such as sender(s) and receivers of a multicast group.

The *billing layer* translates costs calculated by the charging layer into monetary units and generates a bill for a customer. This process may combine technical considerations with economic considerations, such as volume of resources used by the customers, and marketing methods (e.g. offered discounts).

The left part of the figure represents the policy plane, i.e. technical and

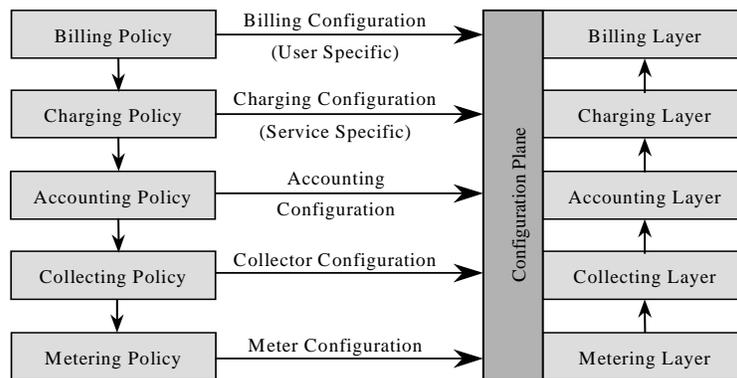


Figure 1. Charging and Accounting Framework

commercial rules for setting parameters based, e.g., on the network configuration (important for metering policies) or the market situation (important for charging and billing policies). The parameters are injected as temporary data to the layers in the right part of the figure through the configuration plane.

4. QOS CHARGING: ARCHITECTURE

Based on the framework introduced in the previous section, we designed an architecture for implementing a charging and accounting service. *Figure 2* shows a sample architecture referred to as VIPCAS (Value-Added IP Charging and Accounting Service). The architecture is being implemented in the ACTS project SUSIE. The network accounting part of the architecture is implemented using IP technology. The charging and billing part of the architecture supports TINA reference points [14] and is implemented using CORBA technology. A policy gateway and an accounting data gateway provide interfaces between these two parts.

The architecture shows that in order to support multicast charging, metering may take place in the edge routers and at the multicast routers. In order to charge, for example the sender, the collected information needs to be fed back to the sender domain (e.g. provider domain A), so that the provider can reconstruct the multicast tree structure in calculating the overall multicast costs.

The architecture uses metering conformant to the IETF Real Time Flow Measurement Architecture (RTFM) [3,5]. For configuration of a meter and for collecting of accounting data, SNMP and a special meter MIB [4,7] are used.

Data from the meters are collected by collecting entities at each domain.

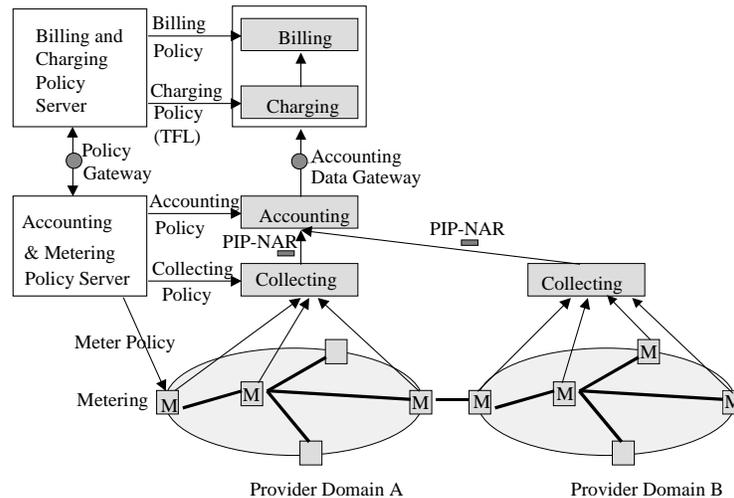
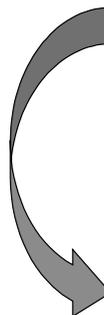


Figure 2. Architecture for Value-Added IP Multicast Charging and Accounting Service

We have specified a data structure, referred to as PIP-NAR (Premium IP Network Accounting Record), that is shown in *Figure 3*, and with more details in the appendix. This data structure is filled by the collecting layer and allows to transport usage information within one provider domain and also to exchange usage information in a multi-provider scenario.

The data structure contains reserved and used resources for an IP flow. It contains a measurement point identifier and a record description to support different styles of PIP-NARs (uni-directional/bi-directional, DiffServ/IntServ style, and others). For unique identification of the flow, a flow description is used. A detailed specification of the PIP-NAR can be found in [29].



Parameter	Type	Length [Bytes]
Record Description		
Version	Char	1
Length of Record	Char	1
Type of Record	Short	2
Measurement start time	Long	4
Measurement stop time	Long	4
Measurement point identification (meter IP address)		r_id
Flow Description (src,dst IP, ports)		f_id
Reserved Resources (e.g. flowspec)		rr
Used Resources (#packets, #bytes)		umd
Data Extension (e.g. distance, burstiness)		

Flagname	Set	Not set
Uni-directional	unidirectional record	Bi-directional record
IPv6	IPv6 Flow	IPv4 Flow
Reserved Resources	PIP-NAR contains Reserved Resources	PIP-NAR contains no Reserved Resources
Used Resources	PIP-NAR contains Used Resources	PIP-NAR contains no Used Resources
DiffServ	Differentiated Services	Integrated Services
Extension	Extension Present	No Extension present

Figure 3. Selected Elements of the PIP-NAR

The content of the PIP-NAR *reserved resources* section differ for Integrated Services and for Differentiated Services. *Used resources* contain the data volume and the number of transmitted packets. The PIP-NAR record format is extensible, where future metrics can be included in the *extension* section using the TLV (type, length, value) syntax. Currently defined extensions are distance and burstiness.

Filled PIP-NARs are forwarded to a charging server where a charging scheme is applied to the data. Subsequently, calculated charges are passed to a billing server, where information about network related charges is used to-

gether with additional tariff information (user information, discounts, etc.) to generate bills and to prepare current billing information for users in real-time (hot billing).

5. QOS CHARGING: CONFIGURATION

This section discusses possible setting of the configuration plane parameters within the charging and accounting architecture.

Metering is needed at the ingress routers in order to obtain the data volume sent. The flows that should be observed are specified by attributes in a rule set that is loaded to the meter MIB by the meter reader [3-7].

For monitoring IntServ reservation information, reservation style and the flow specification parameters have to be added to the flow attributes in the rule set. RSVP Reservation messages are sent from receivers to the sender. For unicast, information about the reservation of the receiver is available at the ingress router. In this case there is no need to explicitly obtain this information from the egress router.

For multicast communication in an IntServ environment it has to be taken into account that each receiver can make an individual reservation. These individual reservations lead to the need of fair cost sharing techniques (see Section 6), where receivers are charged according to their reservation parameters. Due to the RSVP merging capability [33] by merging reservation requests on the way to the sender, the reservation message that arrives at the ingress router may contain a merged flow specification that represents only the highest QoS level reserved. This RSVP message does not allow to extract individual reservations of the receivers. Therefore, reservation information needs to be captured at egress routers. This can be done by monitoring RSVP messages with an RSVP-extended meter. Information from the egress routers can be collected by using SNMP. A further possibility arises due to the capability of RSVP to carry additional objects. Charging and Accounting (CA) objects that carry the individual reservation information can be added to RSVP messages as proposed in [8]. This allows to limit collection of metering information from the ingress router, while still supporting receivers with heterogeneous QoS. Capturing of the additional objects in the RSVP messages requires some minor changes to the meter and the flow table at the ingress router. Scalability of this solution can be improved by using the CA object aggregation technique introduced in [8].

The IntServ flow specification parameters can be captured in the ingress router by using the existing extension to the RTFM meter described in [26]. For enabling a RSVP message parser, a new variable to the control part of the meter MIB (RSVP_acc) has been introduced. Since RSVP allows the renegotiation of QoS parameters, it has to be ensured that all changes in the reservation are noticed. This can be achieved by a collection interval for ac-

counting records that is smaller than or equal to the RSVP update interval. The accounting record collection interval can be specified with an option when the meter reader is started.

The DiffServ model uses per-packet identifiers to specify which service class a packet belongs to. These classes are specified by the provider in few parameter sets. As a result of this, the packet matching engine has to consider the DiffServ field (codepoint) in the packet header as an additional attribute [18]. Flow table and configuration files (rule sets) have to be modified so that these fields are considered in the matching process. The meaning of the codepoint in terms of QoS parameters can vary for different administrative domains (AD). A translation of codepoints to the assigned parameters can be performed above the collecting layer.

Tariff Formula Language

A tariff contains the information on how the price of a service is calculated. A tariff can consist of multiple charging formulas and some rules that define under which conditions each formula is valid. An example would be a tariff with one formula for business hours and one for the night time.

For the representation of tariffs we have developed the Tariff Formula Language (TFL). This context-free language contains essential mathematical operations (e.g. addition, multiplication) and basic mathematical and logical functions (e.g. exponential function, square root, AND, OR etc.). Furthermore, conditional expressions (if/then/else) allow to express e.g. the variation of prices for different time phases. A set of commonly used charging variables (including volume, duration, time of day) are pre-defined. The TFL provides a simple structured way to represent even complex tariffs. Tariffs are expressed in plain ASCII characters and therefore easy to parse and human-readable.

The following example presents the tariff formula from [23] for guaranteed services expressed in TFL. The approach from [23] uses virtual resource parameters: token rate and residual rate, which can be determined using the Int-Serv parameters token bucket rate and service rate. Two coefficients a and b represent the prices per resource unit. In our example we make the prices (parameter a and b) dependent on the time of the day (variable td). The globally defined variables tr and sr denote the reserved token bucket rate r and the service rate R , respectively.

```

# parameter a
a = IF(AND(td>=TIME("00:00:00"), td<TIME("05:00:00")), 0.5,
IF(AND(td>=TIME("05:00:00"), td<TIME("21:00:00"), 0.8, 0.5))
# parameter b
b = IF(AND(td>=TIME("00:00:00"), td<TIME("05:00:00")), 0.2,
IF(AND(td>=TIME("05:00:00"), td<TIME("21:00:00"), 0.4, 0.2))
# tariff formula
p = a*tr + b * (sr-tr)

```

Within our architecture, a TFL parser has been implemented using the GNU Bison tool.

Charging Information Protocol for Tariff Information Exchange

In many situations there is a desire by providers to rapidly adapt tariffs to the market situation (see current telephone market situation where newspapers hardly can follow up the tariff changes). For this purpose we developed the charging information protocol (CIP). The protocol follows the client-server approach. A charging information server maintains the information and distributes it to clients on demand.

The protocol is flexible with regard to the transport protocol. Furthermore, it can be adapted to the special needs of a small or large group of clients by using either unicast or multicast for the announcement of tariff information.

Distribution of charging information can be done by unicast or multicast transmission. Tariffs for the offered service classes are sent in a sequence of information messages (INFO). In order to allow clients to recognize a loss of a packet, the INFO messages contain sequence numbers. With these numbers it is possible to request a retransmission of lost packets.

If a unicast connection is used for the announcement of tariffs, all clients that want to receive information about current tariffs have to register with the CIP sever first. In the registration request clients can choose between two modes to get information from the server. In the push mode (default setting) information messages are sent periodically to the client. In order to prevent sending to non-existent or non-operational stations, messages are acknowledged by the client.

In the pull mode information is only sent on demand. Clients need to send a request (GET_INFO) in order to get the information messages. In the unicast case CIP uses timeout and retransmissions to provide a reliable transport. Besides the reliability and the possibility to use TCP for transport, unicast distribution allows selective individually adapted advertisements. This means that the information can be reduced to tariffs that are new to a particular client.

Furthermore, special offers for certain customers can be conveyed individually.

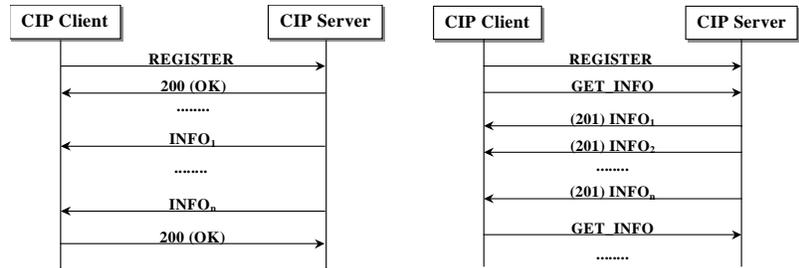


Figure 4. CIP protocol messages for push mode(a) and pull mode(b)

An alternative that provides a better scalability is the distribution of tariff information via multicast. Clients that want to obtain charging information just need to join a particular group. With multicast only the push mode is used. Since no acknowledgements and retransmissions are provided in this case, reliability is lower than for the unicast case. If an INFO message gets lost, the client can recognize this because of the sequence numbers, but he has to wait until a new INFO message is transmitted within the next regular announcement.

The information messages (INFO) contain the following fields: Identification (service name, provider), Validity, Tariff, QoS guarantees, Information about the reservation, and Transaction ID.

The identification field contains the name of the service class and the provider identification. The validity field defines the time interval for which this tariff is valid. This can be used for example to offer special tariffs for specific dates. The tariff field contains the tariff in TFL. The QoS guarantees give a service description regarding the offered QoS. Furthermore, the message gives information about the mechanism used for the reservation. Each INFO message contains a transaction ID to allow the detection of lost messages.

Selection of a service

The selection of an IntServ service can be done by sending an RSVP reservation message to the sender. In order to allow informing the service provider about the chosen tariff, a service selection mechanism is supported by CIP. To choose a service, the client sends a SELECT message to the server. The message is acknowledge by the server. SELECT messages are sent until the server responds or a timer exceeds.

This service selection mechanism is useful if multiple tariffs are offered for the same service class (for different usage profiles or different user groups, or by different providers). Furthermore, the mechanism allows the control of

policy rules if the CIP server is combined with a policy server that enforces the rules within the network (accept or reject reservation requests).

6. MULTICAST CHARGING WITH COST SHARING

In general two approaches can be followed for sharing the costs for multicast communication among the participating nodes:

- Sharing the costs of individual links among all connected nodes (link cost sharing)
- Sharing the total costs of the multicast tree among all members of the multicast group (total cost sharing)

Links Cost Sharing (LCS) schemes split savings from the joint usage of a link among the nodes that are directly connected.

Total Cost Sharing (TCS) can be treated as a simplified, abstract LCS scheme, where the whole network cloud is considered as one link, with all egress routers treated as directly connected nodes. While LCS schemes usually require recursive calculation (for all nodes/links of the multicast tree) to allocate costs to the receiver, for TCS schemes it can be sufficient (depending on the total cost determination strategy) to consider only the border nodes. Therefore, it is possible to implement a TCS with less overhead than a LCS.

Simple cost sharing schemes take into account the number of receivers when evaluating the charge for a receiver, without considering differences in QoS. More sophisticated cost sharing schemes also can take into account reserved resources. As the VIPCAS architecture allows to meter individual reservations at egress routers, sophisticated cost sharing schemes can be supported. In the following, we investigate the following two methods for QoS-sensitive multicast cost sharing: Incremental Sharing (IS), and Proportional Sharing (PS).

The Incremental Sharing (IS) method splits the costs of a link with regard to the individual reservations made by the directly connected nodes (downstream nodes). It implies that reservations can be ordered by QoS levels. The receivers share equally the incremental cost of all QoS levels which are less than or equal to their reserved QoS [21].

The Proportional Sharing (PS) is a novel cost sharing scheme. It is based on the assumption that the shared link is of a greater value to someone who reserved a higher QoS. For PS, the share of the link costs a receiver has to pay is proportional to its reserved QoS. In comparison to IS which shares link costs only based on local information (the QoS requested by a specific receiver/node), PS takes into account the end-to-end 'value' a flow can be associated with for a specific receiver. This is achieved by sharing link costs depending on costs between the splitting point and the end point. PS can be used

as a LCS, or as TCS. If used as TCS, it allows for a QoS-sensitive multicast cost sharing which can be implemented efficiently.

We consider a generic tariffing scheme where costs of a receiver v_n may depend on the reserved resources R_{v_n} , and the distance d_{v_n, μ_m} to the ingress router μ_m . This generic scheme covers more specific charging schemes such as those presented in [17] and [23]. For a unicast connection, the costs $C_{v_n, unicast}$ can be expressed as

$$C_{v_n, unicast} = f(R_{v_n}, d_{v_n, \mu_m}) \quad (3)$$

The total costs of a multicast tree $C_{MC-Tree}$ depend on the number and distribution of the receivers and on their QoS parameters. To express the savings from using multicast instead of unicast, we introduce the multicast gain factor G_{MC} with

$$G_{MC} = \frac{\sum_{n=1}^N C_{v_n, unicast}}{C_{MC-Tree}} \quad (4)$$

By using multicast, the cost share of a receiver v_n gets reduced to

$$C_{v_n} = \frac{C_{v_n, unicast}}{G_{MC}} \quad (5)$$

This method is also useful for a simplified tariff model with a reduced implementation effort. Often, a network provider is much more interested in a comprehensive accounting scheme for the customer than in computing the actual costs of each connection. By estimating the multicast gain factor instead of ascertaining the actual costs of a multicast tree, the computation expenditure for the tariff can be significantly reduced. For N receivers with homogeneous QoS, [10] has estimated the multicast gain factor to be $N/N^{0.8}$ the 0.8 power of the multicast group size based on comparing the summation of edge costs between unicast and multicast routing.

For N receivers with homogeneous QoS, [10] has estimated the multicast tree cost to be $N^{0.8}$ the corresponding unicast costs for sparse multicast trees, resulting in a multicast gain factor of $N/N^{0.8}$.

For load-based tariffing, the reservation-based charging of Equation (1) can be extended to cover actual usage of resources by taking measurements of the effective bandwidth BW_{eff} [12] into account. By including the normalized effective bandwidth BW_{eff} into the cost function C_{v_n} of Equation (1), this leads to

$$C_{v_n} = f(R_{v_n}, d_{v_n \mu_n}) \cdot (\alpha + \beta \cdot BW_{eff_{\mu_n}}) \cdot \frac{1}{G_{MC}} \quad (6)$$

With the weighting factors α and β , the dependence of the total costs from reserved and actually used resources can be adjusted.

Cost Sharing: Simulation Results

For investigating the impact of the two cost sharing schemes (IS, PS) we use a basic multicast scenario with a single fan-out node as topology for our simulation. We simulated the dynamic behavior of one multicast group with receivers dynamically joining and leaving. While joining, each receiver may request an individual QoS. We set the maximum group size to 16. At each time unit receivers change their member status (join or leave the group) with a probability of 0.5. The QoS level for a joining receiver is chosen from 8 different levels and remains unchanged until a receiver leaves the group. Low QoS levels are assumed to be more likely than high QoS levels, using a $1/x$ distribution. Costs were allocated in proportion to the QoS levels. We assume that the costs of the shared link depend not only on the QoS but also on other factors like distance. For a distance of 1, costs for QoS level 1 is 1 monetary unit (MU), costs for QoS level 2 is 2 MU, etc. The QoS level of the shared link is the highest requested QoS at the observed time interval. In our multicast scenario, we set the distance for the shared link to 8, resulting in a cost of 8 MU for QoS level 1, and cost of 64 MU for the maximum requested QoS level 8.

Figure 5 shows the multicast gain that can be achieved depending on the number of group members. (The unicast case is included as a reference.) The true multicast gain is compared with a theoretically derived gain using an estimation of the multicast gain according to the results of Chuang [10]. The figure shows that due to heterogeneous QoS, the results of Chuang only give a rough estimation of the true multicast gain.

Figure 6(a) shows the maximum, minimum and average costs a receiver has to pay when joining the group with a selected QoS level, if one of the two cost sharing schemes IS and PS are applied. Maximum costs occur for cases where the joined receiver is the only receiver of the group. As no cost sharing is possible in this case, the link cost share for both schemes represents the true link costs. For smaller QoS levels the mean link cost share differ only little. For high QoS levels (above level 5) the mean link cost share increases more for the IS scheme than for the DPS scheme. This shows that with the IS scheme, receivers with low QoS pay a smaller share of the common link than with the DPS scheme, while the DPS scheme is more attractive for receivers with high QoS.

For the investigation of the effects that a cost sharing schemes has to individual group members, we performed a second type of simulation. This time we looked at one particular receiver that joins the group with the maximum QoS level and remains in the group until the end of the simulation. The join/leave behavior and the QoS levels of other 15 receivers are assigned randomly like in the first experiment. The share that the receiver has to pay of the costs for the jointly used link varies depending on the number of group members and their chosen QoS level at the observed time unit. *Figure 6(b)* shows the probability of the link costs share that the observed receiver has to pay if the different cost sharing schemes are applied. In all cases the link cost share the selected receiver has to pay varies depending on the reservations of the other receivers in the observed time interval. The differences in the mean charge for the two cost sharing schemes (DPS is more attractive for receivers with high QoS) become again visible.

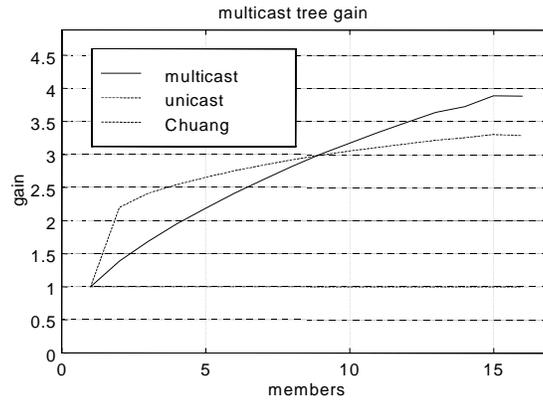


Figure 5. Simulation: true multicast gain vs. estimated multicast gain

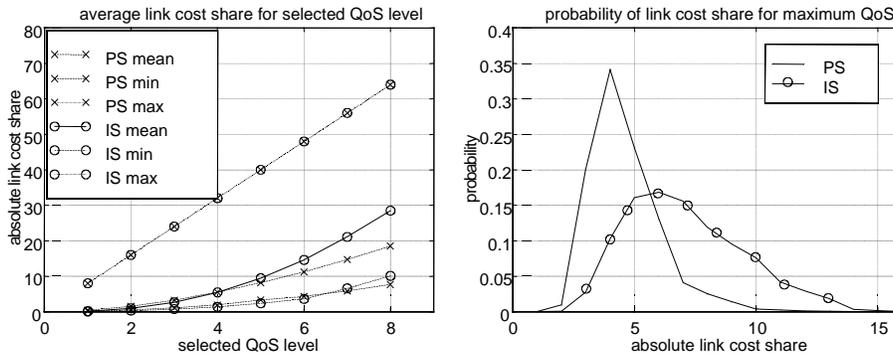


Figure 6. Simulation of end-to-end costs for cost sharing schemes PS and IS:

(a) mean/min/max costs for selected QoS, (b) costs distribution for selection of maximum QoS

7. CONCLUSION AND FUTURE WORK

We have presented a layered framework and described an initial implementation of a charging and accounting architecture based on the framework. The architecture allows its components to be configured for meeting charging and accounting requirements of IntServ or DiffServ in unicast or multicast scenarios. A data structure, the PIP-NAR, is used for transporting usage information from an accounting processing entity to a charging entity, and allows exchange of accounting information in multi-provider scenarios. We presented the Tariff Formula Language (TFL) that allows providers to express a wide variety of charging schemes. A TFL parser has been implemented based on the GNU Bison tool. We have developed a charging information protocol (CIP) that allows to inform clients about tariffs. A prototype of this protocol has been implemented using C++.

For the support of multicast charging, two cost sharing schemes which may be implemented by the architecture have been considered: a link cost sharing scheme, and a novel total cost sharing scheme. The schemes allow a fair allocation of common link costs among receivers, supporting fair charging for QoS-enhanced multicast services.

Future work is planned for deploying and testing the architecture within a multi-provider environment. For this purpose, the PIP-NAR data structure will be used by the collecting layer to exchange usage information among different providers. Future work will also cover optimal selection of IP service class by receivers based on charging information and on user preferences expressed by price limits and utility curves. For this purpose we have implemented a tool that applies numerical optimization techniques (Direction Set Method and Downhill Simplex Method) to tariffs expressed in TFL, and to user preferences expressed in a simplified version of TFL.

8. ACKNOWLEDGEMENT

The authors would like to thank Sebastian Zander and Torsten Ackemann for valuable contributions.

References

- [1] S. Blake, et al. An Architecture for Differentiated Services. IETF RFC2475, December 1998.
- [2] M. S. Borella, V. Upadhyay and I. Sidhu. Pricing Framework for a Differential Services Internet. European Transactions on Telecommunications, Vol. 10(2), March/April 1999.
- [3] N. Brownlee, C. Mills, and G. Ruth. Traffic Flow Measurement: Architecture. IETF RFC2063, January 1997.
- [4] N. Brownlee. Traffic Flow Measurement: Experiences with NeTraMet. IETF RFC2123, March 1997.

- [5] N. Brownlee, C. Mills and G. Ruth. Traffic Flow Measurement: Architecture. IETF Internet draft, Work in progress, June 1999.
- [6] N. Brownlee. SRL: A Language for Describing Traffic Flows and Specifying Actions for Flow Groups. IETF Internet draft, Work in progress, June 1999.
- [7] N. Brownlee. Traffic Flow Measurement: Meter MIB. IETF Internet Draft, Work in progress, June 1999.
- [8] G. Carle, M. Smirnow, and T. Zseby. Charging and Accounting Architecture for IP Multicast Integrated Services. Proc. of Interworking '98, Ottawa, Canada, July 1998.
- [9] M. Canosa, M. De Marco, and A. Maiocchi. Traffic Accounting Mechanisms for Internet Integrated Services. Proc. of SPIE VV'98, Vol. 3529 Internet Routing and Quality of Service, Boston, MA, November 1998, pp.254-264.
- [10] J. Chuang, and M. Sirbu. Pricing Multicast Communication: A Cost-Based Approach. Proc. of INET'98, Geneva, July 1998.
- [11] R. Comerford. State of the Internet: Roundtable 4.0. IEEE Spectrum, Vol. 35(10), October 1998, pp. 69-79. <http://www.spectrum.ieee.org/select/1098/int.html>.
- [12] C. Courcoubetis, F. Kelly, and R. Weber. Measurement-based Charging in Communication Networks. Statistical Laboratory Research Report 1997-19, University of Cambridge, 1997
- [13] D. Decasper, M. Waldvogel, Z. Dittia, H. Adishesu, G. Parulkar, and B. Plattner. Crossbow - A Toolkit for Integrated Services over Cell-Switched IPv6. Proc. of IEEE ATM'97 Workshop, Lisboa, Portugal, June 1997.
- [14] F. Dupuy, G. Nilsson, and Y. Inoue. The TINA Consortium: Toward Networking Telecommunications Information Services. IEEE Communication Magazine, Vol. 33(11), November 1995, pp. 78-83.
- [15] H. Einsiedler, and P. Hurley. Link Weighting: An Important Basis for Charging in the Internet. Proc. of Global Internet Mini-Conference, Sydney, Australia, November 1998.
- [16] G. Fankhauser, B. Stiller, and B. Plattner. Arrow: A Flexible Architecture for an Accounting and Charging Infrastructure in the Next Generation Internet. NETNOMICS: Economic Research and Electronic Networking, Vol. 1(2), 1999.
- [17] D. Ferrari and L. Delgrossi. Charging for QoS. Proceedings of the Sixth IEEE/IFIP International Workshop on Quality of Service - IWQoS'98, Napa, CA, May 1998.
- [18] S. Handelman, N. Brownlee, G. Ruth, and S. Stibler. RTFM Working Group - New Attributes for Traffic Flow Measurement. IETF Internet draft, Work in progress, June 1999.
- [19] F. Hartanto and G. Carle. Policy-based Billing Architecture for Internet Differentiated Services. Proc. of BC'99, Hong Kong, November 1999.
- [20] S. Herzog. Accounting and Access Control for Multicast Distributions: Models and Mechanisms. PhD dissertation, University of Southern California, U.S.A., 1996.
- [21] S. Herzog, S. Shenker, and D. Estrin. Sharing the Cost of Multicast Trees: An Axiomatic Analysis. Proc. of SIGCOMM'95, Computer Communications Review, vol. 25, Oct. 1995, pp. 315-327.
- [22] M. Karsten, J. Schmitt, L. Wolf, and R. Steinmetz. An Embedded Charging Approach for RSVP. Proc. of IWQoS'98, Napa, CA, pp. 91-100, May 1998.
- [23] M. Karsten, J. Schmitt, L. Wolf, and R. Steinmetz. Cost and Price Calculation for Internet Integrated Services. Proc. of KiVS'99, March 1999, Darmstadt, Germany, pp. 46-57.
- [24] K. Kilkki, and J. Ruutu. Simple Integrated Media Access (SIMA) with TCP. Proc. of 4th INFORMS Telecommunications Conference, Boca Raton, FL, March 1998.
- [25] J. MacKie-Mason. A Smart Market for Resource Reservation in a Multiple Quality of Service Information Network. Technical Report, University of Michigan, September 1997.
- [26] A. Maiocchi. NeTraMet & NeMaC for IIS Accounting: User's Guide. CEFRIEL, Politecnico di Milan, May 1998.
- [27] L.W. McKnight, and J.P. Bailey. An Introduction to Internet Economics. The Journal of Electronic Publishing 2(1), May 1996.

- [28] OMG, The Common Object Request Broker: Architecture and Specification, Revision 2.1, August 1997.
- [29] H. Orlamünder (Editor). Parameters and Mechanisms for Charging in IP based Networks [Network Aspects]. TR/NA-080301 V1.0.7 (1999-06), ETSI Working Group NA8 Technical Document, 1999
- [30] B. Rupp, R. Edell, H. Chand, P. Varaiya. INDEX: A Platform for Determining how People Value the Quality of their Internet Access. Proceedings of the Sixth IEEE/IFIP International Workshop on Quality of Service - IWQoS'98, Napa, CA, May 1998, pp. 85-90.
- [31] S. Shenker, D. Clark, D. Estrin and S. Herzog. Pricing in Computer Networks: Reshaping the Research Agenda. Communications Policy. Vol. 20(3), 1996, pp. 183-201.
ftp://parcftp.xerox.com/pub/net-research/picn.ps.
- [32] B. Stiller, G. Fankhauser, B. Plattner, and N. Weiler. Charging and Accounting for Integrated Internet Services – State of the Art, Problems, and Trends. Proc. of INET'98, Geneva, Switzerland, July 1998.
- [33] L. Zhang, R. Braden, S. Berson, S. Herzog, and S. Jamin. Resource ReSerVation Protocol (RSVP) - Version 1 Functional Specification. IETF RFC2205, September 1997.

9. APPENDIX

Elements of the Premium IP Network Accounting Record (PIP NAR)

The types used for the PIP-NAR elements correspond to the OMG IDL syntax defined in [28].

Table 1: Measurement Point Identification and Record Type

Parameter	Type	Length [Bytes]
Record Description		
Version	Char	1
Length of Record	Char	1
Type of Record	Short	2
Measurement Start Time	Long	4
Measurement Stop Time	Long	4
Measurement point identification		
IP Address	String [4/16]	4/16

Table 2: Flags for the Type of Record field

Flagname	Set	Not set
IPv6	IPv6 Flow	IPv4 Flow
Reserved Resources	PIP-NAR contains reserved resources	PIP-NAR contains no reserved resources
Used Resources	PIP-NAR contains used resources	PIP-NAR contains no used resources
DiffServ	Differentiated Services	Integrated Services
Extension	Extension Present	No Extension present

Table 3: Flow Description (IPv4 Flows)

Parameter	Type	Range/unit
Flow Description		
Destination Address	String [4]	
Destination Netmask	String [4]	
Source Address	String [4]	
Source Netmask	String [4]	
Destination Port	String [2]	
Source Port	String [2]	
Flow Start Time	Long	
Protocol	Short	(UDP=17,TCP=6)
Flow Owner	Short	QoS Setup Protocol, Flow, management or others

Table 4: Reserved Resources for Integrated Services

Parameter	Type	Range/unit
Reserved Resources		
Service	Short	Guaranteed Delay=1, Controlled Load=2
Rate	Long	[bits/s]
Depth	Long	[bytes]
Peak Rate	Long	[bytes]
Service Rate	Long	[bytes]
Min. Transmission Unit	Long	[bytes]
Max. Transmission Unit	Long	[bytes]
Backlog	Integer	[Bytes]
Delay	Integer	[μ s]
Slack Term	Integer	[ms]

Table 5: Used Resources

Parameter	Type	Range/unit
Used Resources		
Volume sent	Long	[bytes]
# of IP packets sent	Long	

Table 6: Extension Element

Parameter	Type	Range/unit
Number of TLVs	Short	
Data Extension		
Type	Short	
Length	Short	
Value	Dependent on type	