CHARGING AND ACCOUNTING ARCHITECTURE FOR IP MULTICAST INTEGRATED SERVICES OVER ATM

Georg Carle
Michael Smirnov
Tanja Zseby

GMD FOKUS
Kaiserin-Augusta-Allee 31
D-10589 Berlin
Germany

ABSTRACT

We introduce a charging and accounting (CA) architecture for IP multicast services with QoS guarantees over ATM for QoS-sensitive fair charging. The architecture is based on an extension of the Multicast Integration Server (MIS, [SaCS97]) that efficiently provides IP multicast over ATM, supporting Internet Integrated Services (IS) with receiver heterogeneity and shortcut management. We define requirements for a Charging and Accounting Protocol (CAP) for transporting CA objects, involving CAP clients and a CAP server. We demonstrate that RSVP extensions can be used for transporting CA objects, and propose interworking mechanisms between CAP server, MIS and a Billing Server.

KEYWORDS:

IP MULTICAST, ATM, INTEGRATED SERVICES, RSVP, CHARGING, ACCOUNTING

1 INTRODUCTION

1.1 Charging and Accounting in Computer Networks - State of the Art.

Charging and accounting (CA) for IP Integrated Services is crucially needed for commercial service providers, and allows to shape user behavior for achieving common benefits. CA can be based on reserved or on actually used resource or on a combination of both. Usage measurements for charging can be of different granularity (e.g. virtual connections [Kell96, CoKW97], or TCP flows). Cocchi et al. [CoES93] introduced a simple priority model with an exemplary price discrimination for different levels of service giving preference to bandwidth and/or delay sensitive applications. Their simulation results proof that flat pricing is inferior to a priority model. However, the assumption of fixed base prices for service classes and the difficulty of describing user utility functions leave the applicability questionable. Support for multiple service classes in combination with admission control and usage-based pricing was shown in [Shen95, ShCE96]. The Expected Capacity Framework approach [ClFa97] adds soft QoS guarantees.

Other approaches still lack adequate implementation support. IPng [RFC1671] introduces policy-based routing and accounting by an accounting tag (a source, destination, transaction triplet) which could be used as a voucher and changed at various points along a packet's path thus reflecting responsible parties [RFC1672]. Pricing of QoS levels based on auctions [MaVa94, 95] allows to achieve certain optimality criteria, but causes significant implementation overhead in routers. In [MacK97] the auction-based 'Smart Market' approach is re-examined to be used with RSVP [RFC2205].

1.2 Process Perspective on Charging and Accounting

A number of processes related to charging and accounting can be distinguished. This section introduces required terminology. [StFP98] uses a slightly different taxonomy.

This paper focuses on an IP service as the network service provided by the *service provider*. The *customer* of this service may be the person that is also using the service (the *user*), as typical for IP services in residential areas. In a commercial environment, the customer may be a large organization, with many user having access to the service via a customer premises network (CPN).

Usage metering describes the process of measuring resource usage. Two types of using network resources can be distinguished: reservation of network resources, and consumption (i.e. actual usage) of network resources. This distinction is useful as resources that are reserved by a user and not consumed by this user may be offered to a different user, but usually to different conditions. Charging schemes may reflect this difference, e.g. by charging separately for reservation, and for consumption.

Accounting. The process of accounting involves the following functions: collection of usage data by usage meters, creation of accounting records (data structures, or protocol data units of an accounting protocol), transport of accounting records, collection of usage data by an accounting server.

Charging is the process of evaluating costs for usage of resources. Different cost metrics may be applied to the same usage of resources, and may be allocated in parallel. An example would be a detailed evaluation of resource consumption for further processing by the service provider, and a simple evaluation of resource usage for online display of current costs. A detailed evaluation of the resource consumption can be used for generating bills to the customer, or for internal analysis 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. *Cost allocation* assigns costs to specific endpoints, such as sender and receivers of a multicast group.

Pricing is the process of assigning a price (expressed in monetary units) to a specified service. This process may combine technical considerations, such as resources used for a service, and economical considerations, such as applying tariffing theory and marketing methods. Charging can be performed centrally, processing data collected by an accounting server, or may be performed decentralized. In the latter case, the process of accounting is not limited to collecting of usage information, but comprises also of collection of charging information.

Billing is the process of generating a bill for a customer, based on prices applicable for this customer. Billing is performed by a billing server that collects and processes charging information relevant for the customer.

2 ARCHITECTURAL CONSIDERATIONS ON CHARGING AND ACCOUNTING.

2.1 Charging and Accounting Architecture: Related Work.

Defining a CA architecture that can be customised to a variety of requirements has to address:

Measurement placement. Proposed solutions include moving the measurement to the edge of the network [ShCE96] or integrating specialized modules into switches or routers [EdMV95, DWDA97]. Alternative approaches define dedicated measurement devices [EdMV95] and statistical sampling methods [Vie90]. Contractual traffic estimation based on user profiles [ClFa97] or explicit reservations allow to avoid measurements.

Payment by Senders or Receivers. Sender payment, receiver payment, or a combination of both [ShCE96] are to be distinguished. It is possible to use sender payment at the network layer, and to apply a higher layer protocol for enforcing full or partial payment by receivers. However, the IP multicast service model, where the sender is not aware of individual receivers, and the RSVP service model (sender is not aware of the QoS of individual receivers) suggest that it is very attractive to support a receiver-based payment or a sharing of payment at the network layer.

Cost Sharing between Receivers. Group communication introduces the problem of *cost sharing* between different receivers of a multicast tree. Cost sharing schemes are presented in [HeES95, Herz96].

Charging Schemes. A charging and accounting architecture can be designed for support of specific charging schemes. Metering of resource consumption can be avoided by limiting usage metering to metering of resource reservation [FaSP98], or by contractual traffic estimation based on user profiles [ClFa97]. The architecture presented in [FaSP98] is designed for reservation-based charging. It has been validated by implementation of two pricing schemes: a dynamic volume-based price model, and a delta auction based on customer bids.

Charging in an IP/ATM environment. Different services classes may require different pricing. ATM tariffs may differ largely for the different ATM service classes that can be used for providing IP services. The role of distance and duration of a connection may decease. One example is the proposed set of ATM tariffs for VBR, CBR, ABR, and UBR by Walker et. al. (see [WaKS97]), where the price per volume for VBR traffic is two orders of magnitude higher than the price per volume for UBR service. These prices are set independently of a distance and duration to focus on the error-rate obtained, only. This price difference gives a high incentive to select the UBR service class whenever possible. However, a pure volume-based approach for all service classes as in [WaKS97] does not reflect the complete picture. Since constant bit rate service classes are in principle very similar to the traditional telephone system, a distance-related component may be necessary to incorporate the provider perspective. As cell-based charging in ATM networks is computationally demanding, dedicated hardware-support may be required for metering of ATM resource consumption.

3 PROBLEM STATEMENT.

We identified need for an architecture that allows for QoS-sensitive and fair charging of multicast services for a variety of network and service provision scenarios. The generic Architecture for Charging and Accounting (GENACA) presented in chapter 5 has been designed to provide a generic framework that addresses this need. GENACA has to meet the following requirements: (1) supporting IP multicast services, (2) supporting the IP integrating services model, and (3) the requirements of the CAP protocol presented in chapter 4.

A second goal of our work is to develop an specific architecture for charging and accounting that supports the following additional functionalities: IP multicast over ATM with (4) shortcuts and (5) sender-initiated ATM connection setup. An additional functional requirement (6) is the support of charging schemes that are calculated on a combination of reserved and used resources. To allow for (7) QoS-sensitive and fair charging of multicast services, measurements at senders and receivers have to be supported. The architecture that fulfils requirements (1) to (7) is called Multicast Integration Charging and Accounting Service (MICAS).

4 REQUIREMENTS OF A CHARGING AND ACCOUNTING PROTOCOL (CAP).

Reserved resources (accounted for before actual usage) and consumed resources (accounted for after measuring) are key input parameters for charging. CAP follows the client server paradigm: Usage metering is up to CAP clients while the CAP server[s] provide a charging service for them. The capability to provide *fairness* is an essential requirement for the CA architecture.

Measuring of reserved and consumed resources at the ingress point allows for a direct implementation of source-based charging. According to the IP Integrated Services architecture, reservations originating from different receivers of a multicast group may be merged. Accounting of reserved resources at egress points of the provider network supports this fairness requirement, as it allows to charge receivers according to their reservations. As different receivers of a multicast group may receive delivered data with different QoS, measuring of the received end-to-end QoS allows for QoS-sensitive and fair cost sharing between receivers of a multicast group.

These requirements lead to the need of collecting data from both, the ingress and the egress points within the provider network. Multicast charging does not require the availability of all potential usage metering information for reservation, resource consumption and end-to-end QoS from the ingress point and at all egress points at the CAP server. Data aggregation from CAP clients (representing the measurement points) allows for improving scalability.

CAP follows a *soft state* approach. It tolerates rapid state changes. A need for an interaction between RSVP and CAP originates from possible merging, rejection and re-negotiation of receiver-initiated reservation requests. Additional CAP requirements identified are: support for both, *unicast and multicast* delivery, logging of the resource consumption at *ingress and egress* points of provider network, and desired aggregation of charging information at nodes of the multicast distribution tree.

Monitoring of reserved *and* used resources is needed to check whether a receiver has really obtained the requested end-to-end QoS. If requested and actual received QoS differ, the costs for this particular receiver should be decreased.

5 SOLUTION.

5.1 A generic solution based on the Internet Integrated Services (IS) architecture

IS considers the RSVP protocol [RFC2205] to be one of its QoS setup protocols together with the ST-II, Q.2931, etc. [RFC2216]. In this paper we suggest to encapsulate the Charging and Accounting Protocol (CAP) in RSVP.

5.1.1 Motivations for *CAP in RSVP* encapsulation

RSVP can be used as a *general purpose IP signalling* protocol because of its ability to encapsulate and carry additional opaque data objects; for the CAP these objects are already partly carried by RESV messages. RSVP is a *soft-state* protocol and thus meets one of the important requirements for CAP. RSVP deals directly with the process of receiver driven QoS setup, therefore in all scenarios where receivers pay we rely on RSVP mechanisms. RSVP provides the needed *robustness* with respect to charging for reservations and also gives a good basis for providing fairness of charging and accounting. RSVP directly controls main QoS provisioning building blocks of a packet switched network, packet classifier and packet scheduler, or access to link layer QoS control, and is itself at the same time under policy and admission control restrictions. We see in this an additional benefit: CAP encapsulated in RSVP could make use of these existing RSVP interfaces. The architectural scope of RSVP is, like the scope of CAP, only layer 3. However CAP could also benefit from a large number of existing mappings to specific link layer technologies [ISSLL98]. Last, but not least: RSVP is essentially supportive for IP multicast, which is of particular importance for this work.

What do we actually mean by CAP in RSVP encapsulation? One of the most generic definitions of a packet switched network protocol [Pou78] says: a protocol is a mean for a networked communication and is comprised out of a set of *messages*, processing *rules* at involved parties and a set of *virtual paths*. Therefore, when encapsulation is concerned we could distinguish between a brand new protocol when all the components are defined ad hoc and some flavours of encapsulations of components of a new protocol within another one which could be dabbed here a *carrier* protocol. One flavour of an encapsulation could be to use processing rules and messages of a carrier protocol but with a new [sub]set of virtual paths; another flavour could be with modified messages only and so on. As it is shown below in a section for our specific solution for charging and accounting for IP multicast integrated services over ATM we make use of RSVP virtual paths, we make use of some of the RSVP processing rules and of some of its messages, however the virtual paths of RSVP can be changed as suggested in [SaCS97].

5.1.2 Usage Metering Data Aggregation

To meet our CAP requirements formulated above we address the case of a single provider and we assume here that the provider's network (subnet of internetwork) is comprised of a single link layer technology. However, as it will become clear later, the generalization of our approach for the case of multiple link layer domains is obvious. We need a single (per accounted flow) Usage Metering Data Aggregation Point (UMDAP) within the domain. We assume that this is a requirement for an interaction with Charging and Accounting Server (CA-S), or, if the CA-S is co-located with the UMDAP, a requirement of an interaction with the Billing Server.

The resource consumption logging at both ingress and egress points of the domain requires that we introduce Usage Meters (UM) at these points, while the data aggregation requirement forces us to make a decision about the *placement* of the UMDAP. We identified four placement alternatives for the UMDAP: (a) the ingress router, (b) the egress router, (c) the splitting points of the MCT and (d) separated from both egress and ingress router. With regard to the impact of the UMDAP placement on the CAP in RSVP encapsulation, we found out that little overhead can be achieved by placing the UMDAP at the ingress router. With this solution, the PATH message remains unchanged while the Usage Metering Data (UMD) from receivers is conveyed by the modified RESV message.

An additional requirement comes from the fact that we are looking for a solution being conformant to the IntServ architecture (section 5.1.3) and its QoS setup protocol - RSVP in particular (section 5.1.4).

5.1.3 Conformance to the IS Architecture

The IP datagram flow data path can be considered as a sequence of transmission hops followed by the processing of a flow's datagrams within the "IP module". The IS architecture defines the QoS aspects within this scenario: the quality of service which has been received already by the packet in previous hops is to be taken into account when deciding on a *characterisation* - the required QoS which has to be obtained by the flow's datagrams within the current module in order to meet application level end-to-end QoS requirements [RFC2215, RFC2216].

A characterization is a computed approximation of the actual end-to-end behaviour which would be seen by a flow requesting specific QoS services from the network. By providing additional information to the end-nodes before a flow is established, characterizations assist the end-nodes in choosing the services to be requested from the network [RFC2216]. Within the lifetime of the flow the real QoS setup is influenced not only by the characterization but by the policy and the admission controls. The policy control is out of scope of this paper, however we assume that policy modules within routers are checking permission of the source application to request network services which are charged more than best effort services.

Characterizations are computed from a set of *characterization parameters* provided by each network element on the flow's path, with the use of a *composition function* which computes the end-to-end characterization from those parameters. The composition function may in practice be executed in a distributed fashion by the setup or routing protocol, or the characterization parameters may be gathered to a single point and the characterization computed at that point. [RFC2215].

The IS architecture is designed for layer 3 (IP packets processing and forwarding), however lots of work have been done recently to provide mapping of IS to specific link layers [ISSLL98]. Based on the IS framework, Figure 1 shows our proposal for combination of generic charging (CAP) with the IS. The leftmost part of Figure 1 is a sequence of protocols (including Adaptation protocol for a link *layer specific IS hop*

(LL_IS_HOP)) dealing with the flow on its physical path (the next column of Figure 1) from a source to destination. IS functionalities (the middle column of Figure 1) correspond to IS protocols to the left. Finally, the rightmost column of the Figure 1 positions CAP functionalities: modules for reservation logging and usage metering are co-located with the characterization or composition functionalities of the IS at ingress and egress points of the provider's domain. The issue of UMD aggregation is further addressed in the next section with regard to RSVP.

5.1.4 UMD Aggregation Styles: Following the RSVP Merging

CAP messages (UMD and maybe other data objects) are encapsulated in RESV messages of RSVP as additional opaque for RSVP data objects, in line with the generic RSVP message specification [RFC2205]. If the UMD should be conveyed from each individual usage metering point to the UMDAP unchanged we call this *Passive Aggregation* (PA) style of the UMD. If the amount of UMD could be safely made less then the raw data before it reaches the UMDAP we call it *Active Aggregation* (AA) style of the UMD. CAP's UMD object is specified by the authors of this paper in [CLZ98] as a data structure containing the following substructures: `UMD = <UMID, RecD, FlowD, RR, UR>`, where `UMID` - `Usage` Measurement point identification; `RecD` - Record Description; `FlowD` - Flow Description; `RR` - Reserved Resources; `UR` - Used Resources. The `FlowD` contains elements of the Integrated Services Management Information Base described in [RFC 2213] and [RFC 2214]. The substructures `FlowD` and `RR` could be derived from the original RESV message, therefore we should consider the possibility to aggregate the remaining elements. The `UMID` (receiver IP address type and value) as the necessary identification of the termination point of the multicast tree should stay within the modified RESV until it reaches the UMDAP. We concentrate below on UR aggregation.

The handling of regular RESV messages is hop based, i.e. the content of the message is modified by each IS router. For example merging of the receiver reservations performed with the use of reservations styles and merger rules defined by the IS makes it necessary to convey the receiver's reservation only to the nearest upstream router where a particular reservation state will be established. Thus RSVP (with shared reservations) scales well in terms of a number of receiver initiated data objects and message sizes: regardless the number of receivers and the size of the MCT only a single RESV message of nearly the same size is sent upstream on each hop every 30 s [ZhDE93]. Our aim is to make use of this nice feature.

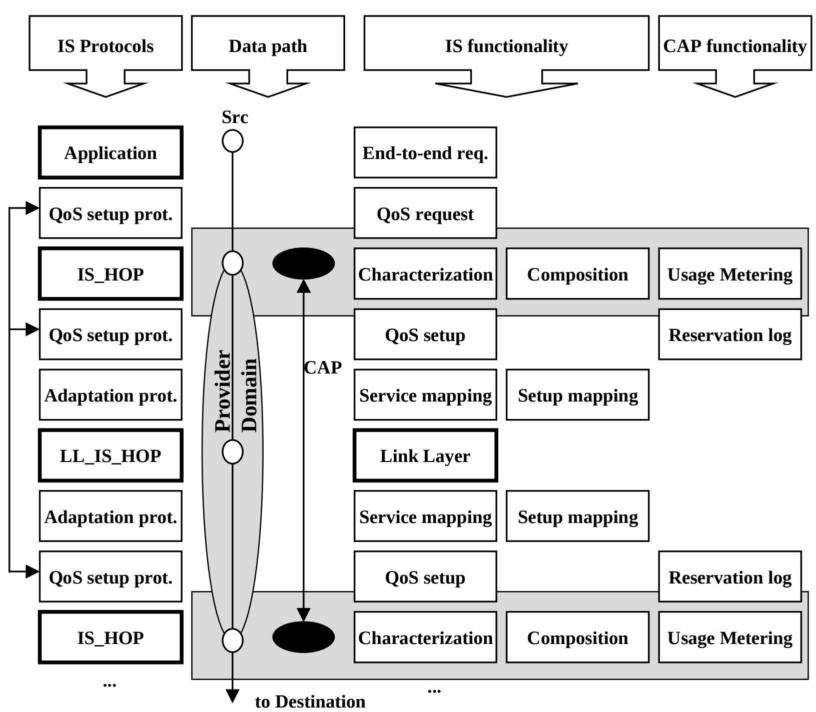

Figure 1: Sequence of protocols at the path for charging and accounting within IS

For passive aggregation each *individual* UMD should be forwarded to the ingress point from each of the receivers. While the processing point of the UMD is located at the ingress router the UMD object carried by each RESV message should be forwarded unprocessed to the next hop RESV message. Thus, the size of the RESV message with CAP UMD objects could grow significantly. This size will depend linearly on the number of receivers within the domain.

Table 1 provides our proposal for the *active* UMD aggregation along the MCT which follows the taxonomy of RSVP merging styles. The basic idea is very simple: if the UMD in a splitting point of the MCT (an IS router) could be *characterized*[1] as "the same" for several receivers then CAP should not transport duplicated raw data to the UMDAP but rather use some sort of aggregation encoding for this subset of users (grey fields in table 1). This could be a data structure like, e.g.: `< extreme UR value, {receiver, diff}>`.

Table 1: UMD and Reservations Merging

		Reservations	
		Distinct	Shared
Sender Selection	Explicit	FF(Si{Qi}[,...])	SE(S1,..., Sn{Q})
Sender Selection	Explicit	**PA(S,G,Q)**	**SAA(S,G,*)**
Sender Selection	Wild-card	not defined	WF(*{Q})
Sender Selection	Wild-card	**not defined**	**AA(*,G,*)**

Table 1 uses the following notation for the MCT: AA is applied to (*,G,*) - shared MCT and shared QoS requirements for the group; Source specific AA (SAA) is applied to (S,G,*) - source specific subtrees[2] for shared QoS requirements; PA is applied to (S,G,Q) - source specific subtrees and distinct QoS requirements for the group.

Like for merging in the original RSVP we follow the same set of basic principles for the applicability of UMD aggregation in CAP: (i) aggregation of UMD with different styles is disallowed, (ii) the aggregation of UMD for the distinct QoS request and wildcard source selection is not defined and (iii) UMD aggregation follows the same principles as for RSVP merging. Therefore, the size of the UMD being forwarded by the modified RSVP upstream will grow anyway, but only because of the difference in aggregation styles.

1 Note that in fact we try to *reconstruct* the original IS characterization out of the UMD.

2 We use the term source-specific to identify only CAP's viewpoint on the MCT: flows from different sources are charged separately as if they are using separate MCTs.

5.2 Applying GENACA to ATM

In order to show that the GENACA architecture can be applied to a wide variety of network technologies, we have chosen as example the technology that varies most from legacy networks: ATM. The different nature of IP and ATM requires additional mechanisms to provide IP and in particular IP multicast services over ATM. In this section we show not only how GENACA can work over ATM but also how the GENACA architecture can profit from ATM features and reuse mechanisms already applied for the provisioning of IP multicast communication over ATM. The most sophisticated solution for the provisioning of IP multicast service over ATM can be reached by using a Multicast Integration Server (MIS) [SaCS97] because the MIS supports - beside multicast address resolution - QoS levels and shortcut connections.

5.2.1 The MIS Architecture

The MIS architecture integrates two setup protocols: EARTH [Smir97a, b], for IP multicast address mapping to ATM addresses and layer 2 establishment of QoS connections, and RSVP for layer 3 resource reservation. Analysis of possible design alternatives [CaCS98] shows that for efficient mapping for IP Multicast Integrated Services to ATM, a new network entity - the RSVP server (RSVP-S) - is beneficial. Since the MIS architecture permits the usage of shortcut connections within the ATM network, the RSVP-S acts as an additional hop for the control (RSVP) messages. The data transfer via the shortcut connection remains unchanged.

The MIS is a particular node within the ATM network that combines RSVP-S and EARTH-S (EARTH-Server). The EARTH-S keeps the multicast address resolution table and answers requests from EARTH clients, whereas the RSVP-S co-ordinates the correct distribution and modification of RSVP messages. The EARTH and the RSVP protocol share the control connection to the clients. Within the MIS the two entities communicate via two interfaces, the extended Routing Support for Resource Reservation

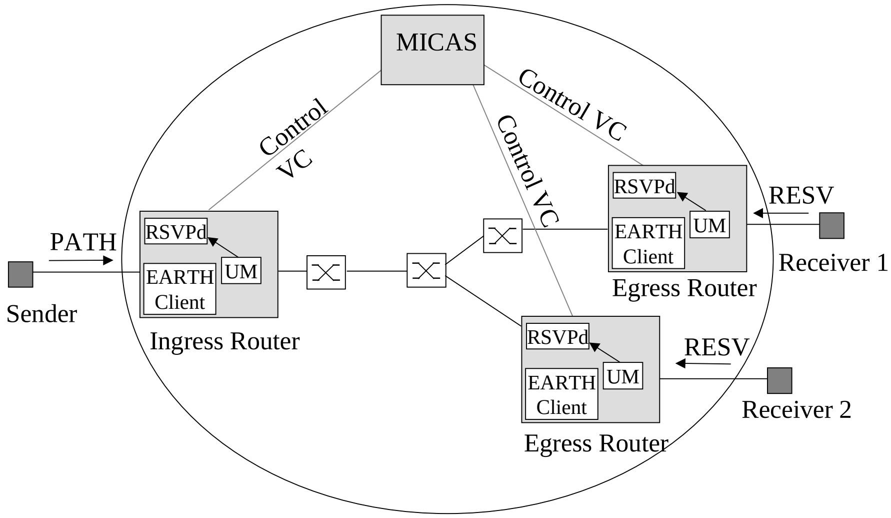

Figure 2: MICAS Architecture

Interface (eRSRR) and the Quality of Service Support Interface (QSSI).
All RSVP messages are communicated to the RSVP server. If a PATH message arrives, the RSVP-S asks the EARTH-S via the eRSRR interface about the location of the members of the multicast group. For this the two primitives EARTH_QUERY and EARTH_RESPONSE are defined. The membership information is used to ensure the distribution of PATH messages to all group members. If a RESV message arrives at the RSVP-S, the EARTH-S is informed about the reservation request by sending an EARTH_RESV message, which contains the desired QoS settings, via the QSSI interface. The mapping of RSVP QoS notation to ATM QoS notation is done in the RSVP-S as specified in [GB97].

The QoS parameters are communicated to the EARTH client that resides at the sender within an EARTH_MULTI message. The EARTH_MULTI is part of the EARTH protocol and contains the ATM addresses of the group members. This information is needed at the sender in order to establish the point-to-multipoint connection. In the MIS architecture the admission control is performed by the layer 2 signaling. Since the connection establishment takes place at the senders side instead of being co-located with the RSVP-S, the RSVP-S must be considered slightly different from standard RSVP entities. This remote admission control mechanism is introduced in [SaCS97].

After the connection establishment has been performed, the EARTH client informs the EARTH-S about the new QoS by sending an EARTH_QoS_NOTIFY message. This information is forwarded to the RSVP-S by sending an EARTH_RESV_ACK message via the QSSI interface.

5.3 Multicast Integration Charging and Accounting Service (MICAS)

The MICAS provides a Charging and Accounting Service for IP multicast delivery in ATM networks. It is based on the GENACA concept but uses additional features of the MIS architecture. Since all RSVP messages have to traverse the RSVP-S within the MIS, the MIS turns out as the optimal location for the UMDAP. In the MIS architecture both collaborating protocols EARTH and RSVP are already using the same set of control VCs.

In order to process the UM data immediately, the CA server is introduced as a third entity within the MIS. A new interface between RSVP-S and CA-S is defined. Figure 3 shows the interaction of the three entities within the MICAS. Figure 4 shows the MICAS components.

If a PATH message which carries a UM data object arrives at the RSVP-S, the UM data is forwarded to the CA-S. After the members of the multicast group are discovered via message exchange with the EARTH-S, the PATH message is forwarded to all receivers. If a RESV message with UM data is received, the reservation information carried in the FLOWSPEC and the UM data carried in the CA objects are extracted from the message and passed to the CA-S.

The CA-S stores the information in a CA table. This table can be accessed by a billing server in order to get the technical basis for the computation of charges for each user. In addition to technical costs the price calculation within the billing server can take into account economic decisions and marketing rules introduced by the provider.

6 CONCLUSION

The presented charging and accounting architecture efficiently supports IP multicast services with QoS guarantees over ATM, receiver heterogeneity and shortcut management. The MIS turns out as a suitable aggregation point for the collection of charging and accounting data. RSVP - already a part of the MIS architecture - was identified as well-suited carrier of CA objects for both reserved and used resources. Therefore, QoS sensitive fair charging for IP multicast services over ATM can be provided in a scalable way for a variety of network and multicast service provision scenarios.

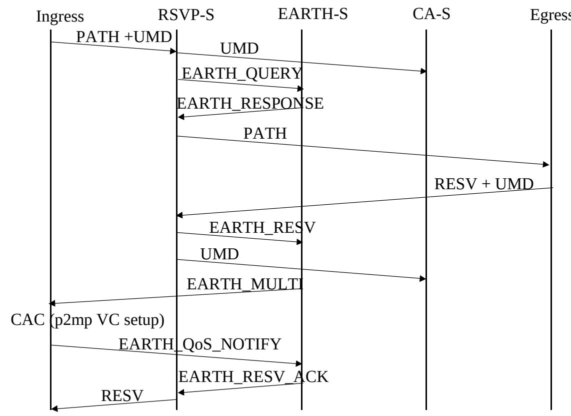

Figure 3: Interaction of MICAS components

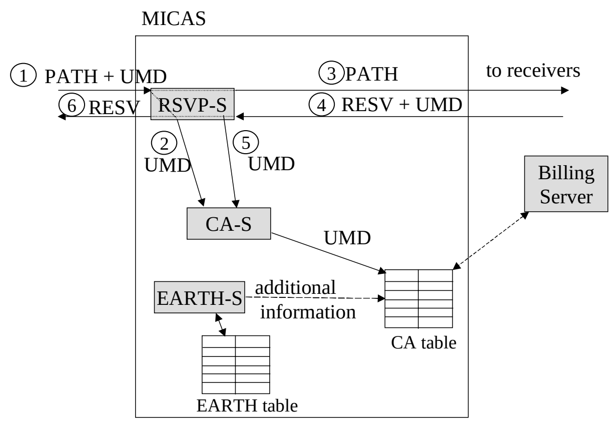

Figure 4: MICAS components

ACKNOWLEDGEMENTS

Foundation of this work: MIS is partially funded by the CEC ACTS Program under the project MULTICUBE. The research agenda presented in the paper is intended to be partially covered by the ACTS project SUSIE.

REFERENCES

[CaSS96] G. Carle; T. Schmidt, J. Seitz: Fair ATM Charging with Consideration of Traffic Characteristics and QoS Parameters, IEE Colloquium on ‘Charging for ATM‘, London, U.K., 12. November 1996

[CaCS98] M.Canosa, A. Corghi, H. Sanneck, M. Smirnov, L. Vismara, D. Witaszek, T. Zseby: Development of early-stage prototype IP/ATM integration, MULTICUBE Deliverable, April 1998

[Clar97] D. Clark: Internet cost allocation and pricing. In L. W. McKnight & J. P. Bailey, (Eds.), Internet Economics (pp. 215-252). Cambridge, MA: The MIT Press, 1997

[ClFa97] D. Clark, W. Fang, Explicit Allocation of Best Effort Packet Delivery Service, Technical Report, Laboratory for Computer Sciences, MIT, 1997

[CLZ98] G. Carle (Ed.), C. Lübbecke, T. Zseby:. Definition of the Premium IP Network Accounting Record (PIP-NAR), Work in Progress, GMD FOKUS, SUSIE project draft specification, March 98.

[CoES93] 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.

[CoKW97] C. Courcoubetis; F. Kelly; R. Weber: Measurement-based charging in communication networks, Statistical Laboratory Research Report 1997-19, University of Cambridge, 1997

[DWDA97] D. Decasper, M. Waldvogel, Z. Dittia, H. Adiseshu, G. Parulkar, B. Plattner: Crossbow - A Toolkit for Integrated Services over Cell-Switched IPv6; Proceedings of the IEEE ATM'97 Workshop, Lisboa, Portugal, June 1997.

[EdMV95] R. Edell, N. McKeown, P. Varaiya: Billing Users and Pricing for TCP, IEEE Journal on Selected Areas in Communications, Volume 13, Number 7, September 1995, pp. 1162-1162

[FaSP97] G. Fankhauser, B. Stiller, B. Plattner: Arrow: A Flexible Architecture for an Accounting and Charing Infrastructure in the Next Generation Internet, First Berlin Internet Economics Workshop, Berlin, Germany, October 24 - 25, 1997

[FaSP98] G. Fankhauser, B. Stiller, B. Plattner, Arrow: A Flexible Architecture for an Accounting and Charging Infrastructure in the Next Generation Internet, Accepted for publication in “Netnomics”, 1998.

[FSVP98] G. Fankhauser, B. Stiller, C. Vögtli, B. Plattner: Reservation-based Charging in an Integrated Services Network, 4th INFORMS Telecommunications Conference, Boca Raton, Florida, U.S.A., March 1998.

[GB97] M. W. Garrett, M. Borden, "Interoperation of Controlled-Load Service and Guaranteed Service with ATM", Work-in-progress, IETF Internet-Draft, <draft-ietf-issll-atm-mapping-04.txt>, November 1997.

[Herz96] S. Herzog: Accounting and Access Control for Multicast Distributions: Models and Mechanisms, PhD dissertation, University of Southern California, U.S.A., 1996

[HeSE95] S. Herzog, S. Shenker, D. Estrin: Sharing the Cost of Multicast Trees: An Axiomatic Analysis; In Proceedings of SIGCOMM 95, August 1995.

[ISSLL98] Integrated Services over Specific Link Layers (issll), IETF WG Charter description, URL http:// www.ietf.org/html.charters/issll-charter.html.

[Kell96] F. P. Kelly. Charging and Accounting for Bursty Connections, in "Internet Economics", J. P. Bailey and L. McKnight (editors). MIT Press, Massachusetts, 1996.

[MacK97] J. MacKie-Mason, A Smart Market for Resource Reservation in a Multiple Quality of Service Information Network, Report, University of Michigan, September 1997.

[MaVa94] J. MacKie-Mason; H. Varian: Pricing the Internet; In Public Access to the Internet (Kahin B. and Keller J., Eds.), Prentice Hall 1994

[MaVa95] J. MacKie-Mason; H. Varian: Pricing Congestibel Network Resources, IEEE Journal on Selected Areas in Communications, Volume 13, Number 7, September 1995, pp. 1141-1149

[Pou78] L. Pouzin, H. Zimmermann: Tutorial on Protocols, in IEEE Transaction on Communications, special issue Packet Switched Networks, 11(66), November, 1978.

[RFC1671] B.Carpenter, IPng White Paper on Transition and Other Considerations, RFC 1671, CERN, August 1994.

[RFC1672] N. Brownlee, Accounting Requirements for IPng, University of Auckland, August 1994.

[RFC2205] L. Zhang, R. Braden, S. Berson, S. Herzog, and S. Jamin. Resource ReSerVation Protocol (RSVP) - Vesion 1 Functional Specification. RFC2205, IETF, September 1997.

[RFC2213] F. Baker, J. Krawczyk, A. Sastry, Integrated Services Management Information Base using SMIv2, RFC 2213, Standards Track, September 1997.

[RFC2214] F. Baker, J. Krawczyk, A. Sastry, Integrated Services Management Information Base, Guaranteed Service Extension using SMIv2, RFC 2214, Standards Track, September 1997.

[RFC2215] S. Shenker, J. Wroclawski: General Characterization Parameters for Integrated Service Network Elements, RFC2215, Standards Track, September 1997.

[RFC2216] S. Shenker, J. Wroclawski: Network Element Service Specification Template, RFC2216, Informational, September 1997.

[SaCS97] L. Salgarelli, A. Corghi, M. Smirnow, H. Sanneck, D. Witaszek: Supporting IP Multicast Integrated Services in ATM Networks, Work in Progress, Internet Draft draft-salgarelli-issll-mis-00-11.ps; Nov. 97

[ShCE96] S. Shenker, D. Clark, D. Estrin and S. Herzog, Pricing in Computer Networks: Reshaping the Research Agenda; Communications Policy, Vol. 20, No. 1, 1996; ftp://parcftp.xerox.com/pub/net-research/picn.ps

[Shen95] S. Shenker: Some Fundamental Design Decisions for the Future Internet; IEEE Journal on Selected Areas in Communications, Vol. 13, No. 7, 1995, pp 1176 - 1188.

[Smir97a] M. Smirnov EARTH - EAsy IP multicast Routing THrough ATM clouds, Work in Progress, Internet Draft draft-smirnov-ion-earth-02.txt, March 22, 1997

[Smir97b] M. Smirnov. Scalable and Efficient Multiprotocol IP Multicast over ATM. In Proceedings of SPIE VV'97 - Broadband Networking Technologies, Dallas - Texas, Nov 1997.

[StFP98] B. Stiller, G. Fankhauser, B. Plattner, N. Weiler: Charging and Accounting for Integrated Internet Services - State of the Art, Problems and Trends, Proceedings of INET'98, Geneva, Switzerland, July 1998

[Vie90] B. Viero: Traffic Measurement on Variable Bit Rate (VBR) Sources with Application to Charging Principles, Computer Networks and ISDN Systems 20, Elsevier Science Publishers B.V. (North Holland), 1990, pp. 435-445

[WaKS97] D. Walker, F. Kelly, J. Solomon: Tariffing in the New IP/ATM Environment, Telecommunications Policy, Vol. 21, pp. 283-295, May 1997.

[ZhDE93] L. Zhang, S. Deering, B. Braden, D. Estrin, S. Shenker, and D. Zappala: RSVP: A new Resource ReSerVation Protocol, IEEE Network, 1993