Charging of Multimedia Flows in an Integrated Services Network

Burkhard Stiller, George Fankhauser, Bernhard Plattner
Computer Engineering and Networks Laboratory, Swiss Federal Institute of Technology, ETH Zürich
Gloriastrasse 35, CH–8092 Zürich, Switzerland
Phone: +41 1 632 7016, Fax: +41 1 632 1035
E-Mail: [stiller | gfa | plattner] @ tik.ee.ethz.ch

Abstract

Admission to communication services for multimedia applications does not imply purely technical control, it has to integrate economic admission policies as well. This is mainly due to the fact that multimedia communication services within the Internet cannot be assigned in a cooperative fashion as they represent bandwidth-hungry applications. Internet economics make various assumptions on pricing policies and services, encompassing features of the economic model, the pricing mechanism, the network model, and their practicability in the Internet today. Based on general architectural aspects for integrating and supporting charging and accounting properties of integrated services networks, the implementation of a flow-based approach to charge multimedia data in the Internet is presented. Ongoing work in the design, implementation, and evaluation of a control protocol for charging and accounting resource reservations in the integrated services Internet applies usage-based pricing models employing an auction-based pricing model (delta auction) and an adaptive, load-sensitive, volume pricing model. In contrast to packet-based charging schemes, the developed approach of flow-based charging features low processing overhead and low bandwidth requirements of less than 1% for fine-grained and frequent reservation setups are achieved.

1 Introduction to Flow-based Charging

Concerning networking traffic, todays Internet is able to support a best-effort type of traffic only. For various reasons, however, some applications, especially multimedia applications require networking support for at least a little better service assurances, may be even some guarantees. These service guarantees have to be provided by end-system resources, *e.g.*, supported by a QoS-aware operating system, and by the network itself, *e.g.*, supported by resource reservation protocols. Therefore, a first approach for charging deals with the flow-based approach, since it deals with a significant amount of resources that has to be reserved exclusively for a certain application. As these networking resources cannot be shared any more between other applications, an adequate policy has to ensure that this resource consumption is payed for. This policy includes as the main component economic incentives besides traditional protocol support as mentioned. Since the amount of required resources per application – or in general per flow – may vary, a flat-fee approach does not provide the right economic incentive. Therefore, usage-based pricing schemes need to be applied which take into account the actual amount of currently consumed resources.

Once this need for suitable methods for a usage-based service principle has been accepted, prerequisites to implement accounting, charging, and pricing have to be studied in detail. Generally required components for these tasks and their interdependency are depicted in Figure 1 which does not distinguish between end-systems and routers.

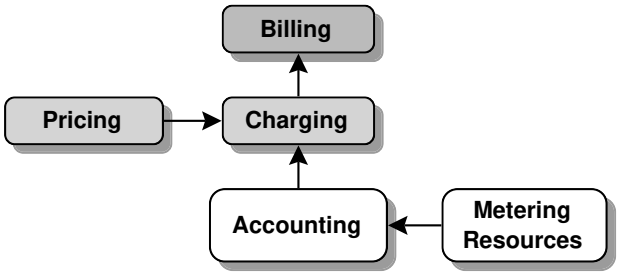

Figure 1: Functional Components of an Architecture for Flow-based Charging.

At a first step, the particular utilization of resources is metered or monitored and has to be accounted for. Accounting data can be available for end-system resources, *e.g.*, in terms of utilized central processing unit time, memory, or special devices. However, accounting data for network resources and network access are more important for the networking perspective of evaluating communication costs. These resource data may encompass, *e.g.*, the acquired network access bandwidth, duration of a communication, expected delays, loss rates, or time-of-day information. To support multimedia flows, accounting data must include Quality-of-Service (QoS) specifications for resources that are utilized to provide a service.

In a second step, this accounting information has to be collected for a particular user or application. Based on different networking infrastructures, communication scenarios, or network provider concepts a single or multiple collecting centers exist within an administrative networking domain. For that reason, a general accounting protocol has to be integrated on top of or within the networking

architecture, such as the Internet protocol suite and service architecture, to exchange well-defined accounting records. Accounting information even may be piggy-backed on protocols already in use, like a resource reservation protocol.

In the third step, a charging scheme calculates the currently valid price for these utilized resources which are precisely determined by the collected accounting records. A pricing scheme specifies in a generic and flexible manner prices for certain units of goods, *e.g.*, per one megabyte of data at noon over a distance of ten hops. This pricing information and the real use of resources determines the amount of money to be payed for the usage of these resources. The local pricing settles base prices per unit which are dependent on many factors, such as time-of-day, thresholds, or discounts. In turn, these calculated charges are determined in charging records and may be transferred and charged directly to the user, directly to the application, indirectly via host institutions, or indirectly via the utilized workstation. It has to be mentioned that charging may take place in the same system where the accounting took place. In this case the system itself is the accounting and charging center which may be designed for every involved system in the network, resulting in a completely distributed and decentralized approach.

In the final step, these charging records are handled by a billing system. The billing system decides on a provider-based or on a user-requested policy, when and how the final communications bill is delivered to a user or the application. Normally, billing is done in a more centralized manner, based on transmitted charging records. In addition, interoperator agreements may be determined to allow for a collective billing between single operators.

Obviously, these three steps involve a well-structured interaction between these components in a distributed manner, where the information on their location is of secondary interest, but the accounting, charging, and pricing information and their handling in terms of new or enhanced control protocols is of primary interest. In addition, of overall interest in this case is the secure exchange of accounting and charging records by these control protocols and a secure handling of corresponding payments.

2 Designing a Charging Protocol for Flows

Retaining the old flat-fee model for charging and accounting of multimedia flows does not reflect the new service classes being offered and it would lead to a fatal admission policy problem. As it can be assumed that the problem of reserving resources will be technically solved in the near future, a consistent and stable policy to grant admission to these resources is still missing. At least for expensive high-end communication services, such as multimedia data flows, this policy should be implemented using financial incentives. Focusing on reservations has the advantage of treating large bandwidth quantities that last over a longer period of time than packet-based charging schemes can. Such reservations for flows build the basis for an efficient charging mechanism with low overhead which is required to make usage-based pricing mechanisms economically feasible. Concerning QoS requirements and data rates generated by multimedia applications, it is very likely that service classes needed are based on reservations. Since they will occupy a significant part of the network bandwidth, they should be the premier target for usage-based pricing.

Using multimedia flows based on resource reservations as basic units for charging and accounting, the reservation protocol transports signaling information. A simplified version of RSVP [4] including basic functionality only, is used to carry charging relevant information. As shown in Figure 2, the basic receiver-oriented control messages are extended. PATH messages convey additional payment information and quote requests. This information is echoed by the receiver and returned to the sender in RESV messages containing receiver payment information and quotes. This round trip messaging can be used to have either sender or receiver pay for the reservation. To allow for locally different pricing models, the willingness to pay is reduced at each hop and price quotes are added up by the local market price. Payments are requested for bandwidth used on each hop's outgoing link over a short period of time.

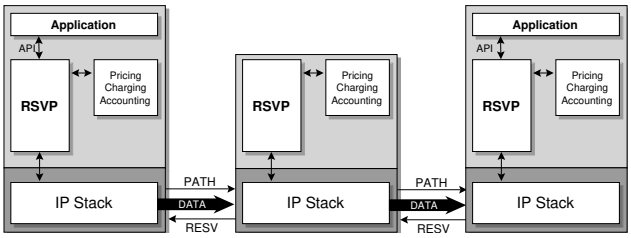

Figure 2: Applying RSVP to Charging Tasks

The developed approach supports local pricing decisions to enable easier interconnection of multiple providers which is based on peer agreements at aggregated links. It minimizes charging protocol overheads and allows for splitting payments between senders and receivers.

2.1 Pricing Models

Two pricing models have been used to evaluate the charging protocol. The major difference between these models is the way how they calculate prices. The delta auction approach calculates market prices based on consumers' bids and has been shown to be incentive-compatible [3]. Consumers announce their willingness to pay by bidding for the bandwidth they currently request. All requests are split into small sized units in order to perform an auction on bandwidth only. In situations, where no congestion occurs and all the bids fit into the link's offered bandwidth, no second price is available and the service can be given for free or optionally for a low volume-fee in case of non-zero marginal costs.

The other approach, a dynamic volume price model, is provider oriented with base-prices that are modified according to demand. Although it calculates prices dynamically according to the load on the network, the base-prices are set by the operator of each node [2]. The model features different service classes, best-effort and reserved bandwidth have been considered. The basic unit sold is a bandwidth reservation over a fixed period of time. This model provides two options which control the price when congestion is about to occur. One is a function that describes the increase in price when demand exceeds supply provided by the link bandwidth. As a parameter to this function, the actual demand/supply ratio is used. The other option is a parameter which specifies at what load the before mentioned function is applied to increase the price.

3 Implementing the Charging Protocol

Simplified RSVP messages used for reservation setup, such as PATH and RESV messages, are enhanced by adding payment information and price queries. Depending on the economic models used at each router the necessary information is extracted from a uniform message and processed. For example, the willingness-to-pay field of a RESV message is used to transmit bids in case of an auction. For more provider-oriented models with predefined prices, the reservation is only granted when the required amount is available from this field. PATH messages are used to pin a path and setup a state to make sure that RESV messages follow the same route back to the sender. PATH messages may contain a request to quote market prices or a sender provided payment. All the messages share the same protocol data unit (PDU) structure with type fields and flags indicating the operation. Figure 3 shows the basic PDU of the reservation protocol. The highlighted parts have been added for the charging and accounting protocol.

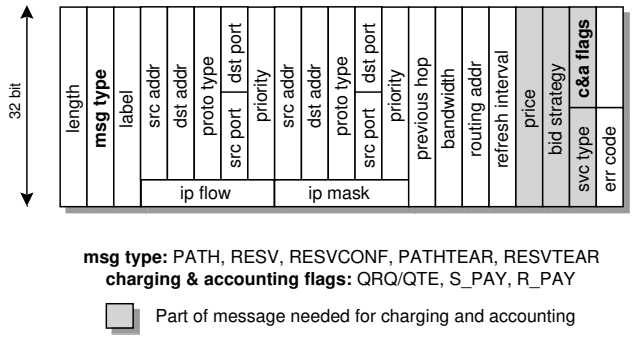

Figure 3: Enhanced Reservation Message Format for Charging Resources.

Normal operation includes quote requests and reservation requests that are initiated by applications via an extended reservation protocol API. Applications are expected to react in an adaptive fashion to the prices quoted and reservations performed. Based on minimal user input, such as profiles and budget constraints, middleware modules are responsible for negotiation of details with the extended API. Failures, such as unavailable resources or insufficient payments, are signaled via the standard PATHTEAR and RESVTEAR messages of the reservation protocol.

4 Results

Concerning first evaluations of our prototypical implementation, the general overhead of the charging and accounting protocol as described above was measured in relation to the requested resources. Primarily, bandwidth is used as a measure, but also the requited memory for accounting and extended flow and reservation tables has been studied. Furthermore, the delay introduced by the charging protocol and the scalability for large numbers of flows have been investigated. Interworking of different pricing models and payment modes still needs to be assessed. Measurements were performed on a flexible testbed of Crossbow routers and hosts [1]. Machines used included PCs with 233 MHz AMD K6 processors, running NetBSD 1.21, having 512 kByte of second level cache and sufficient amounts of memory to hold large flow tables.

As an overview and comparison of memory needed to store flow states extended by charging and pricing information at each router, some figures specify the requirements of total 194 Byte in the RSVP daemon in user space: The flow state requires 120 Byte, pricing 14 Byte, and the account 12 Byte. Basic authentication data (28 Byte) is also included in the form of an integrity object with an MD5 checksum, sequence number, previous hop, and time stamp. The charging relevant data uses 28% of the space for the total flow state information.

The protocol overhead needed to transport charging relevant information is most significant (cf. Table 1) since it is using up resources that could be otherwise sold. As an example we are using an IP telephony stream with very fine grained reservation periods of 5 seconds and a constant bit rate of 64 kbit/s (40,000 Byte for the 5 second period). Charging information for this stream uses only 0.75% of the flow's total bandwidth. During one reservation period 500 data packets are processed and a message round trip for querying the price and refreshing the reservation is needed.

Table 1: Protocol Overhead for IP Telephony Stream.

	[Byte/ packet]	Packets []	Sum [Byte]	[%]
Payload	80	500	40000	73.2
IP/UDP Headers	28	500	14000	25.6
Simplified RSVP Msg	96	2	192	0.4
Charging Information	12	2	24	0.1
Price Queries	108	2	216	0.4
Authentication Data	28	2	56	0.1
Optional RESVCONF	136	1	136	0.2
Total	-	-	54624	100.0

Processing at routers to perform price queries, the delta auction, and authentication requires for a long living flow on the average 322 µs per packet on a router which in total encounters for 2.3% for the sample scenario of IP Tele-

phony flows. In order to put a realistic load on the system, flow tables had to maintain 100,000 flows concurrently in a single router while performing this measurement and the Crossbow traffic control functions to setup or remove flows were used.

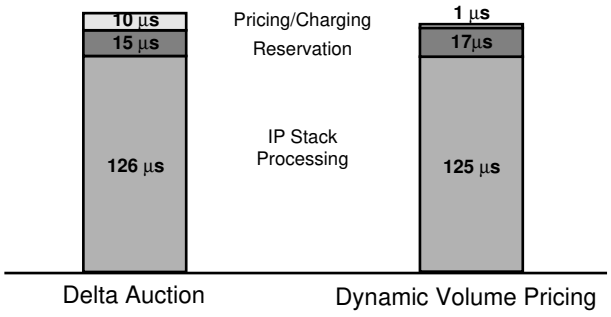

Figure 4: Processing Times for Pricing Schemes

Figure 4 shows the processing times for a single reservation request including processing of the charging information for different pricing schemes. In contrast to the above stated average numbers, it does only include processing figures for the reservation and charging protocol and not for traffic control. These numbers have measured interface-to-interface delays on a router connected to two 10 Mbit/s Ethernet subnetworks. Total delays for a reservation message encountered 151 µs and 143 µs for the delta auction and the dynamic volume pricing model, respectively. While the latter does not add significantly to the processing delay (<0.1%), the more complex delta auction uses up less <7% of the overall delay. However, it must be noted that the IP protocol stack processing delay including the protocol overhead, the required socket processing, and the use of timers, remains similar at 126 µs and 125 µs for the delta auction and the dynamic volume pricing scheme, respectively. In addition, the reservation fraction includes RSVP processing and the management of time-outs for reservation periods.

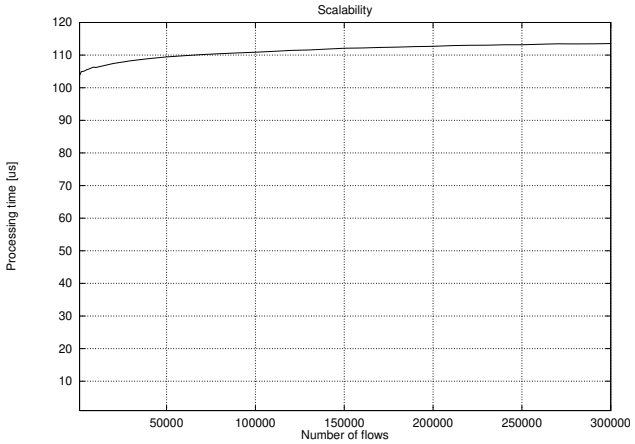

Figure 5: Reservation Request Processing Times for the Delta Auction Pricing Scheme.

While simple pricing models, such as volume pricing (even if the price is dynamic), can be implemented with simple lookup functions and do not involve complex sorting or searching procedures, the presented delta auction has to maintain proper data structures to process incoming bids and to decide whether they can be admitted or have to be rejected. Searching can affect the scalability of such a solution. For the prototype implementation, balanced trees were used to keep processing complexity at the *O(log n)* where *n* is the number of flows. Using load generators for *n* up to 300,000 flows shows the desirable behavior for reservation processing with the delta auction (cf. Figure 5). Balancing trees with information about the ongoing auctions is a strict requirement, since bidding strategies of participating end-users (or applications and daemons that act on behalf of them) will lead to degenerated trees.

5 Conclusions and Future Work

Adding charging and accounting functionality to the integrated services network architecture is feasible. By applying flows instead of packets, it is suitable to provide charging and accounting functionality without prohibitive overhead. For a sample scenario of 100,000 IP Telephony flows at one router the processing overhead for the auction-based, dynamic pricing and the associated charging functions was only at 2.3%. The most important result determines that the introduced protocol overhead was determined by 0.75% only for a reservation interval of 5 s, handling sender as well as receiver payments and supporting local pricing methods concurrently.

Future work includes new routing schemes considering QoS and pricing information to enhance the currently static routing approach. In addition, an interesting economic problem is how to integrate best-effort flows into the architecture. The presented current approach focuses deliberately on reservations only, under the assumptions that the revenue from reserved flows will exceed that from best-effort traffic. Future empirical studies will also be needed to address the reliability of this assumption.

References

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

[2] 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 – Economic Research and Electronic Networking, Baltzer, The Netherlands, 1998.

[3] W. Vickrey: *Counterspeculation, Auctions and Competitive Sealed Tenders*; Journal of Finance, 1961.

[4] L. Zhang, S. Deering, D. Estrin, S. Shenker, D. Zappala: *RSVP: A New Resource Reservation Protocol;* IEEE Networks Magazine, Vol. 31, No. 9, pp 8-18, September 1993.