Charging Multicast Communications Based on a Tree Metric

Hans Joachim Einsiedler*, Paul Hurley†,
Burkhard Stiller‡, Torsten Braun*

* Institute of Computer Science and Applied Mathematics (IAM)
University of Berne
CH-3012 Berne
einsiedl@iam.unibe.ch
braun@iam.unibe.ch

† Institute for Computer Communications and Applications (ICA)
Swiss Federal Institute of Technology - Lausanne (EPFL)
CH-1015 Lausanne
paul.hurley@epfl.ch

‡ Computer Engineering and Networks Laboratory (TIK)
Swiss Federal Institute of Technology - Zürich (ETHZ)
CH-8092 Zürich
stiller@tik.ee.ethz.ch

April 22, 1999

Abstract

Current charging mechanisms in the Internet are restricted mainly to volume and time of the day. Links or network clouds can have weights which may be dynamic, e.g. based on congestion level. The weight information must travel from the sender to the receiver. We provide a mechanism, which sends a message down the multicast tree for collecting the necessary information. This information can be used in the edges and the core of the network for charging. We also discuss how charging could be implemented by using these weights along links or network clouds from edge to edge.

Keywords

Charging, Differentiated Services, Fair Distribution of Costs, Internet Protocol, Multicast Border Router.

1 Introduction

Multicasting offers an important Internet service to a variety of applications, such as conferencing, collaborative group work, or software update distributions. As the growth of the Internet in terms of traffic volume, users, hosts, links, and routers does not appear to be slowing down, it is essential to try to keep the amount of identical traffic small which is being distributed to a number of different receivers. Since mechanisms, principles, and protocols applicable to multicast traffic have been developed in approximately the last 10 years, a number of solutions have emerged to the questions of scalable and technical multicast support.

However, the payment for this type of traffic has been neglected to date. Appropriate charging methods for multicast traffic are in their initial design steps or are even missing completely. This is partly due to the fact that basic charging for unicast traffic proposals have only been recently launched. In addition, the number of influencing factors for multicast traffic in support of appropriate charging models is much larger when compared to unicast traffic. These factors encompass at least the following: (1) heterogeneous receivers in terms of transmission rates, bandwidth, and reliability models, (2) payments of senders, receivers, or split payments, (3) handling the frequent joining and leaving of multicast groups and (4) establishment of a multicast tree and its exploitation for routing and accounting decisions.

Charging for multicast traffic is described within this paper at the Internet Protocol (IP) level. There is no application-level or content charging included, although future work is needed to address this question as well. Therefore, this paper is organised as follows. Section 2 describes the new Quality of Service (QoS) based services for the Internet, the idea of weighted links for cost computation and in which way the link weights can be used for Differentiated Service multicast. The cost computation is the content of Section 3. Section 4 gives an outlook on how the charging can be solved in the next generation Internet Protocol (Internet Protocol version 6). The charging mechanism is introduced in Section 5. Section 6 contains conclusions and future work.

2 Internet Services

2.1 Quality of Service in the Internet

In the current Internet, we have only Best-Effort service, a service which provides no guaranteed Quality of Service (QoS). The Internet Engineering Task Force (IETF, [1]) realised this problem and developed a protocol for resource reservation called called Resource Reservation Protocol (RSVP, [2], [3]). It was foreseeable that it would not scale with the enormous growing Internet.

The IETF extended the discussion and built Differentiated Service working group (Diff-Serv, [4]) to focus on the scalability problem. The idea is to offer different services in the Internet. Different proposals were made such as *Assured service* [5], *Premium service* [6], *Differential Link Sharing* [7], *Scalable Reservation Protocol* [8], *Simple Integrated Media Access* [9], *Weighted Proportional Fair Sharing* [10] and *Olympic service* [11]. Of these, currently only two services are discussed, an extended version of the *Assured service* [12] with four different classes and a modified version of the *Premium service* [13].

In contrast to RSVP, where the policy is fully dynamic and hop-by-hop, in Differentiated Services the Service Level Agreement (SLA) between hosts/clients and providers or between service providers is long-term and requires minimum signaling.

2.2 Link Weights

Based on the idea of different services in the Internet, the authors in [14] developed a mechanism to charge for the use of the resources which the packet consumes during its journey through the network. This information is stored in the Time to Life (TTL) field of both, the Internet Protocol (IP) headers (IP version 4, [15] and IP version 6, [16]. There is no change to the original functionality of this field. The TTL field is initialised to a certain value when the packet is sent. Each router decrements this field by one. If this field becomes zero before it reaches its destination, the packet will be discarded by the router because it stayed too long in the network. In [14], links and *Internet Service Provider* (ISP) clouds have weights, which represent a value for the network resources. Edge routers of the ISP clouds or even interior routers will use these weights to decrement the TTL field by a value greater than or equal to one. The decrementation of this field is retained which preserves the looping prevention functionality of the TTL field. The start value has to be set to a maximum for all packets in the beginning with the value 255 (8 bit).

The link weights or ISP cloud weights do not have to be constant over time, but have to be defined in agreement with neighbour ISPs.

The costs of the used resources on the way through the Internet can be computed with this information in the edge of the network and also in the network. The costs are then accumulated and directed in the opposite direction of the data flow.

2.3 Multicast and Differentiated Services

The IETF does not currently discuss the use of Differentiated Services for multicast and the combination of the various routing protocols such as *Distance Vector Multicast Routing Protocol* (DVMRP, [17]), *Multicast extension for OSPF* (open shortest path first, [18]), *Protocol-Independent Multicast Sparse Mode* (PIM-SM, [19]), and *Core Based Trees* (CBT, [20]).

Independent of the routing protocol, it can be decided whether or not a host/node is allowed to receive or send in a given service class and how much it is permitted to send. When a host sends multicast data and requires a specific service class for its multicast session, it must communicate with the first trustable multicast border router. This router has to decide if the host is in or out of profile. The exact semantics have yet to be defined and they are out of the scope of this proposal. One possibility is to extend the numbers of *Internet Group Management Protocol* (IGMP, [21]) messages for this communication/signaling.

For illustration purposes, however, consider the following. If the host is allowed to send multicast traffic in the requested service class, then the multicast router announces the session by broadcasting including the payment direction which we describe later. If not, the router rejects the request and the host must either send a new request (i.e. reduce the resource request or change traffic class) or give up.

In order to proceed, we make the following assumptions for service class multicast sessions.

- At every branching point, there is a trustable multicast router which knows its number of branches (for definition please see Section 3). E.g., this information can be provided from the routing protocol. Normally, edge routers on shared mediums do not explicitly know how many different branches they have. However, given a scenario where these routers are also Differentiated Service enabled, they will have profiles for each of their branches. Each branch must explicitly register to obtain their profile and a router can ascertain its number of branches from this.
- In a shared medium, the data, which is forwarded to a member of the multicast group, is encrypted if the data is payment sensitive. This is a fundamental requirement that is independent of and complementary to our proposal.

3 Fair Distribution of Cost

Consider a multicast session where each receiver is responsible for paying. The cost of a given multicast session should be shared fairly amongst its receivers in relation to its cost to the network i.e. the distance traveled is shared.

We define the *number of branches* as the number of active receiving nodes (hosts or routers) for a given multicast session connected to a multicast router. This value includes the number of local receivers i.e. the number of receivers who partake in the session directly from the router.

A fair distribution of distance cost amongst subtrees is given upon application of the recursive formula

$$cost_i = \frac{cost_{i-1} + \sum_{n=1}^{m} weight_n}{br_i} \qquad (1)$$

where:

$cost_i$	Cost at branching point where $cost_0 = 0$
m	Number of links between the two branching points
$weight_n$	Weight of this links
br_i	Number of branches at branching point i

Thus, branching and cost information is needed at every multicast router in order to provide fair distribution of cost, and this information needs to be passed to the edge. For this we propose a message which will carry this information down along a multicast tree.

Ways to share the cost of a multicast tree are analysed in [22]. Our scheme can be considered to be "Local members and next-hops are allocated identical hops" (see [22]) in their notation which is a one-pass mechanism. We do not require the total number of receivers downstream from each multicast router to be known; just its branching information. We share the cost fairly amongst subtrees and not among total receivers. Cost division into areas facilitates collective charging and we believe it provides better control than direct division of each link's cost amongst all receivers. Our scheme, if desired, would require minute changes to work with the "All locals are considered as one next hop" (ENHS) scheme discussed in [22].

3.1 Weights in Multicast

The weight of the link and number of branches must be supplied from each router to the edge to enable fair cost distribution. We refer to this information in each router as its (branch, weight) pair.

This data is supplied by a message packet per service class per multicast session. The message can be the payload of a newly defined IGMP "cost" message or embedded in a *User Datagram Protocol* (UDP, [15]) packet which is sent out in the same service class with a different but related multicast address to the multicast data. It is then forwarded along to each branch.

The first trustable multicast border router next to the sender within the multicast tree generates the message periodically. Lost messages are not that important due to the periodicity. This message also announces the payment direction i.e. if the sender or receiver should pay. The message is built by adding this information into the first field and then the first pair (number of branches, 0). The weight field is set to zero because the multicast router does not yet know the weight of the outgoing link.

Every subsequent router is responsible for adding a (branch, weight) pair to every arriving message packet for a multicast session within the specified service class before forwarding. Each pair is built as follows.

- Each router increments the weight of the previously inserted pair by the weight of the previous link.
- If the number of branches from the router is more than one, it adds a new pair consisting of (number of branches, 0).

This ensures that routers with only one branch does not need to add another pair. For intermediate non-multicast routers which do not support our scheme, the tagging of weights is done at the next multicast border router.

3.2 Example

Figure 1 shows a part of a multicast tree. Each branch of the multicast is given a weight. The number of branches are known to the trustable multicast border routers. Also in this figure is the Branch/Weight message, which, for clarity, is shown without the IP/IGMP header.

It is periodically generated and goes down from the root to the leaf. This information is then passed to the charging application in each router. The last trustable multicast border router (leaf router) does not generate this message but adds the needed weight of the last link (w4) and the last Branch/Weight pair (2,0) and passes this information to the charging application.

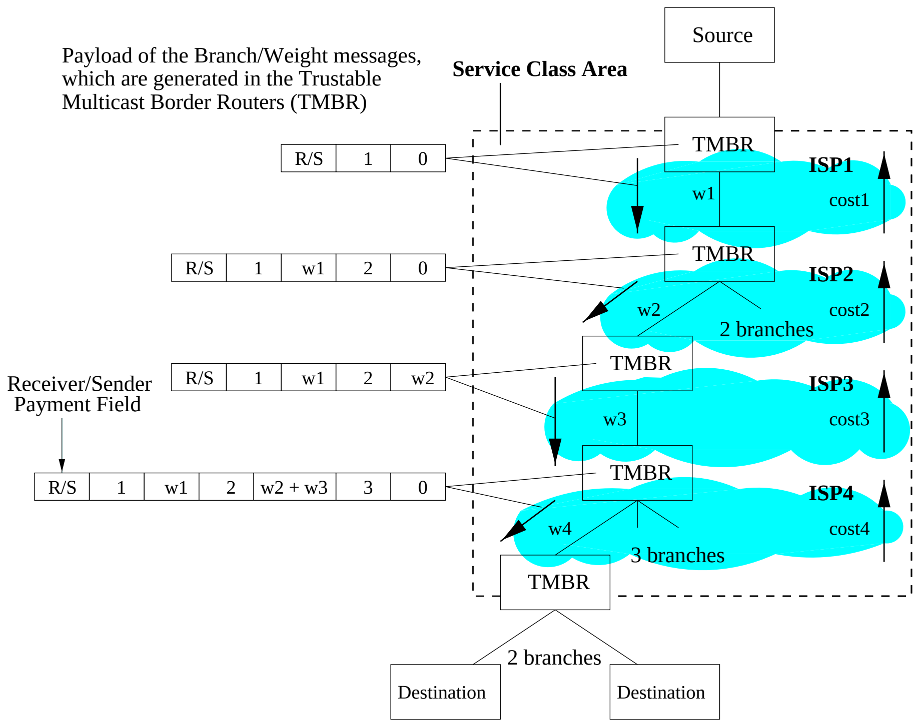

Figure 1: Branch of a Multicast Tree

4 Implementation in IPv6

Using IPv6, there are two ways one could implement the Branch/Weight information packet.

- *Separate UDP or IGMP message in IPv6 as in IPv4*
  The trustable multicast border router next to the sender periodically generates this message and the message is forwarded down the multicast tree.

- *Additional Header for IPv6*
  IPv6 allows additional information to be conveyed from the source to intermediate systems along the path, by the use of extension headers. Extension headers follow the basic header as shown in [16] and [23]. Below we define a multicast Branch/Weight header.

We propose an IPv6 extension header with the following structure (see also Figure 2).

- As defined, the first byte is the mark for the next header. The choice of header number is left open. (There are some headers already defined [16]). Then there is one byte for length. The length byte size allows us to add 84 Branch/Weight pairs which is more than acceptable.

- The third byte contains the payment information.

- Two bytes follow with the number of branches. We assume that two bytes are enough to store all possible branches in a future multicast network.
- The next byte contains the weight of the path.
- Then a new Branch/Weight pair, three bytes, is created and the process repeated.

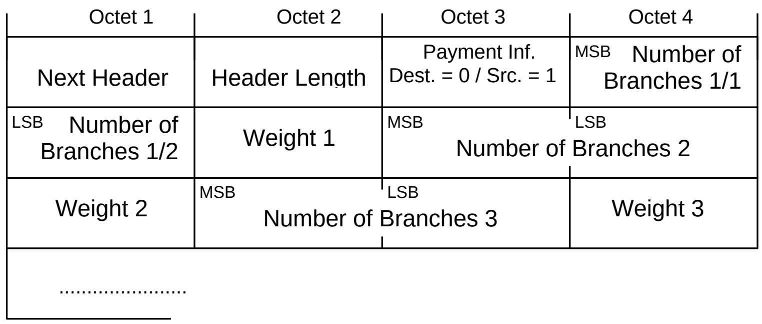

Figure 2: The Branch/Weight IPv6 header

There are two ways how this header can be used. Either the first trustable multicast router as seen by the sender can extend the new header to each multicast datagram or this can be done periodically with a multicast datagram. In the latter case, the subsequent multicast routers should update only this specific multicast datagram.

5 Charging Mechanism

5.1 Multicast Transmission Computation of Costs

Recall that a fair distribution of the costs between branching points was given by Formula 1 in Section 3.

Multicast Session Where Receiver Pays

The computation of the cost using the formula is done by the last trustable multicast border router to which the receiver is connected.

For example, the cost in example 3.2 is

$$cost_{Receiver} = \frac{\frac{\frac{w1}{2}+w2+w3}{3} + w4}{2}$$

Multicast Session Where Sender Pays

When the sender pays for the transmission, the final field for the number of branches is not needed because there is no distinction amongst service classes on the last hop (i.e. the last hop is “free”).

The cost to the sender is the sum of all branches,

$$cost = \sum_{br=1}^{m} cost_{br}$$

where:

$cost_{br}$	Cost of the multicast branch
m	Number of all multicast branches

For the same example 3.2, the cost for the shown branch is

$$cost_{br(anch)} = \frac{\frac{w1}{2} + w2 + w3}{3} + w4$$

5.2 Charging and Cost Direction in the Network

In Figure 1, one can see the path of the Branch/Weight message from the root to the leafs of a multicast tree. Let us assume that each link is owned by a different ISP. The costs along the path is shown for such a case in Table 1.

The charging direction is the **opposite** direction to the information/packet flow direction from the leafs to the root. The charges are accumulated and each ISP will decrement the charge by the cost of its resources and will forward the rest to the connected ISP in the opposite direction to the packet flow.

ISP link and its weight	Costs
ISP 4, w4	$cost_4 = \frac{\frac{w1}{2}+w2+w3}{3} + w4$
ISP 3, w3	$cost_3 = \frac{w1}{2} + w2 + w3$
ISP 2, w2	$cost_2 = \frac{w1}{2} + w2$
ISP 1, w1	$cost_1 = w1$

Table 1: Cost Distribution of a Multicast Tree

5.3 Signaling for Attending a Session

In Section 2.3, we assumed a trustable border router is placed at each branching point. As such, there is no need for a special signaling protocol to request the participation for a session since the forwarding of the data is dependent from the fact if the next hop can pay for the session or not.

In principle, the leafs are not trustable. The trustable multicast border router has to decide if a leaf is allowed to participate in a session, which is transmitted using Differentiated Services. For the case when receiver has to pay for its participation in the session, the trustable border router must check if the receiver can pay for it. The decision is done after a receiver has requested the participation. This problem will not occur when the sender pays for the transmission because the charging aspects are already covered when the sender/root requests to send the data using Differentiated Services.

6 Conclusions and Future Work

We propose a special cost packet to ensure fair distribution of cost. We showed also that the branching factor is needed in the cost computation. Weights and branching factor as well as the payment direction are the information contained in this cost packet.

It is possible that the cost packet could also be implemented by putting the aggregate fraction of cost calculation in the router thus keeping the size of the packet small and independent of the number of routers traversed. This has the disadvantage of requiring the router to perform these calculations. We are investigating the feasibility and advantages of this as part of our ongoing work.

We also have shown how the costs can be transfered between ISPs. The direction of the accumulated cost is from the leafs to the root of the multicast tree, which is the opposite direction to the data flow.

The importance of multicast in the Internet will increase and it will not only be multimedia applications which will use multicast and with a requested QoS at the same time. Other data transfer application such as updating of distributed databases, distribution of security keys, etc. will need to request QoS to satisfy a certain amount of reliability. QoS traffic should be charged differently to Best-Effort.

One possibility is to include the use of the resource in this charge. We have shown an easy way to do so without any change to the network infrastructure or existing protocols, facilitated by the simple addition of a new message type.

Acknowledgments

We would like to thanks Ljubica Blazevic and Eric Gauthier for the interesting discussions.

References

[1] IETF Web Page, "Internet Engineering Task Force (IETF)." http://www.ietf.org/.

[2] R. Braden, L. Zhang, S. Berson, S. Herzog, and S. Jamin, "Resource ReSerVation Protocol (RSVP) – Version 1 Functional Specification," September, 1997. Internet RFC 2205 (draft-ietf-rsvp-spec-15.txt).

[3] IETF-IntServ-Working Group, "Integrated Services for the Internet." http://www.ietf.org/html.charters/intserv-charter.html/.

[4] IETF-DiffServ-Working Group, "Differential Services for the Internet." http://www.ietf.org/html.charters/diffserv-charter.html/.

[5] D. Clark and J. Wroclawski, "An Approach to Service Allocation in the Internet," July 1997. Internet Draft: draft-clark-diff-svc-alloc-00.txt.

[6] K. Nichols, V. Jacobson, and L.Zhang, "A Two-bit Differentiated Services Architecture for the Internet," Nov 1997. Internet Draft: draft-nichols-diff-svr-arch-00.txt.

[7] Z. Wang, "User-Share Differentiation (USD) - Scalable bandwith allocation for differentiated services," Nov. 1997. Internet Draft: draft-wang-diff-serv-usd-00.txt.

[8] W. Almesberger, T. Ferrari, and J.-Y. L. Boudec, "Scalable Resource Reservation for the Internet," Nov. 1997. Internet Draft: draft-almesberger-srp-00.txt.

[9] K. Kilkki, "Simple Integrated Media Access (SIMA)," June 1997. Internet Draft: draft-kalevi-simple-media-access-01.txt.

[10] J. Crowcroft and P. Oechslin, "Differentiated End to End Services using a Weight Proportional Fair Sharing TCP," Dec. 18, 1997. University College of London.

[11] K. Nichols and S. Blake, "Differentiated Services Operational Model and Definitions," Feb. 1998. Internet Draft: draft-nichols-dsopef-00.txt.

[12] J. Heinanen, F. Baker, W. Weiss, and J. Wroclawski, "Assured Forwarding PHB Group," Feb. 1999. Internet Draft: draft-ietf-diffserv-af-06.txt.

[13] V. Jacobson, K. Nichols, and K. Poduri, "An Expedited Forwarding PHB," Feb. 1999. Internet Draft: draft-ietf-diffserv-phb-ef-02.txt.

[14] H. J. Einsiedler and P. Hurley, "Link Weighting: An Important Basis for Charging in the Internet," in *Proceedings of the Third Global Internet Mini-Conference in conjunction with Globcom '98*, (Sydney/Australia), November 1998.

[15] W. R. Stevens, *TCP/IP Illustrated Volume 1 - The Protocols.* Addison Wesley Publishing Company, March, 1996.

[16] S. Deering and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification," Dec. 1998. Internet RFC 2460.

[17] D. Waitzman, C. Partridge, and S. Deering, "Distance Vector Multicast Routing Protocol," Nov. 1988. Internet RFC 1075.

[18] J. Moy, "Multicast Extensions to OSPF," March 1994. Internet RFC 1584.

[19] D. Estrin, D. Farinacci, A. Helmy, D. Thaler, S. Deering, M. Handley, V. Jacobson, C. Liu, P. Sharma, and L. Wei, "Protocol Idependent Multicast-Sparse Mode (PIM-SM): Protocol Specification," June 1997. Internet RFC 2117.

[20] A. Ballardie, "Core Based Tree (CBT) Multicast Routing Architecture," Sept. 1997. Internet RFC 2201.

[21] W. Fenner, "Internet Group Management Protocol, Version 2," Nov. 1997. Internet RFC 2236.

[22] S. Herzog, S. Shenker, and D. Estrin, "Sharing the "Cost" of Multicast Trees: An Axiomatic Analysis," in *Proceedings of ACM/SIGCOMM'95*, Aug. 1995.

[23] S. A. Thomas, *IPng and TCP/IP Protocols - Implementing the Next Generation Internet.* John Wiley and Son, Inc, 1996.