

IP Multicast over Cable TV Networks

A. Selcuk Uluagac and Jon M. Peha

Carnegie Mellon University¹

Dept. of Electrical and Computer Engineering

Pittsburgh, PA15213, USA

selcuk@alumni.cmu.edu, peha@cmu.edu, www.ece.cmu.edu/~peha

Abstract. When a cable TV network that provides Internet access is connected to multiple ISPs, there are instances where multicast does not work or works inefficiently. This paper identifies causes of these problems, and proposes solutions, demonstrating that it is possible to provide efficient multicast with any of the architectures under consideration. In addition, the de facto industry standard for data transmission over cable networks, DOCSISTM, guarantees that a cable company will have the ability to block certain multicast traffic (such as traffic generated by Internet television broadcasters which compete with the cable company's core business.) This paper describes how an ISP can circumvent this. Under the assumption that there is a significant amount of multicast traffic, we show that cable companies and ISPs would be motivated to provide multicast services in all cases, but there are cases where they are not motivated to choose an efficient approach. Finally, we consider the impact of possible regulation that prohibits cable companies from blocking. In cases where this regulation has any impact, it can increase or decrease the cost of multicast services.

1 Introduction

Many companies multicast television and radio programming over the Internet [1]-[4], even though dial-up access cannot support real-time video and is marginal for quality audio. As more subscribers get broadband Internet access [5], these and other multicast applications may become common, making multicast efficiency important.

Efficient multicast mechanisms have been developed for a cable network connected to one Internet Service Provider (ISP), as was typical when Internet over cable began. However, this is changing. The US Government required Time Warner to connect its cable networks to multiple ISPs as a condition of the merger with America On Line [6]. There have also been trials with multiple ISPs in the U.S. [7]-[11]. Moreover, multiple ISPs can connect to cable data networks in Canada [12].

In this paper, we address multicast over cable platforms that support multiple ISPs. With current methods of enabling connections between cable networks and multiple ISPs, there are realistic cases in which multicast does not work or works inefficiently. We describe the causes of these problems and possible solutions. Additionally, the dominant standard for IP over cable [13] allows a cable company to block any multi-

¹ The authors would like to thank Professor Marvin Sirbu of Carnegie Mellon University and the reviewers for their valuable comments.

cast stream it wishes. For example, a cable company might use this capability to block services such as Internet TV that would compete with the cable company's core business [14]. We will describe the mechanism to block competing multicast traffic, and how ISPs can provide multicast services even when cable companies try to block.

Under the assumption that multicast becomes common, we will examine the conditions in which cable companies and ISPs are motivated to provide efficient multicast services, and the conditions in which a cable company would have incentive to use its blocking capability. There has been controversy regarding whether cable companies should be regulated so they cannot limit competition, or interfere with the free flow of information [15]-[17]. Under the same assumption that multicast is common, we also examine the possible consequences of regulating multicast mechanisms.

In section 2, we review architectures that permit multiple ISPs to operate over a cable network. Section 3 presents associated multicast problems on these architectures and proposes solutions. Section 4 explains how cable companies can block multicast traffic, and what an ISP can do to prevent this. Section 5 evaluates which technical design decisions would best support an effective business strategy from the perspective of a cable company and the perspective of an ISP. This section also considers the effect of a regulatory prohibition on blocking. Section 6 concludes this paper.

2 Forwarding Mechanisms that Support Multiple ISPs

Figure 1 shows a typical cable architecture [18]-[19]. Subscribers access the network through a *cable modem (CM)*. Traffic from CMs flows to a *Cable Modem Termination System (CMTS)* [20] at the cable *Head-End*. The CMTS manages traffic from or to the CMs, and allocates cable resources (i.e. bandwidth) both upstream and downstream. The CMTS is generally connected to an upstream interconnection device located at the *Regional Head-End* [20]. ISPs connect to the network at this *Regional Head-End*. The protocol suite that defines the interface and operation requirements for these devices is Data Over Cable Service Interface Specification (DOCSIS) [13].

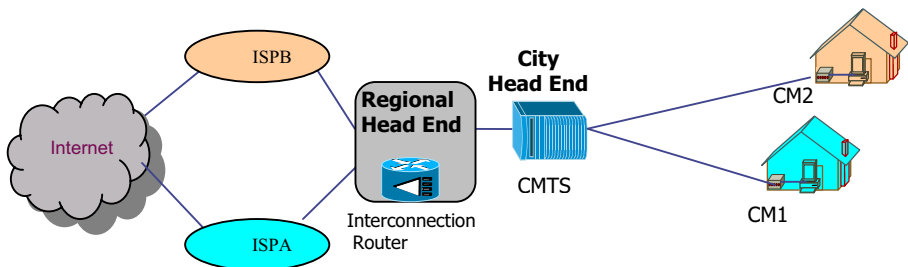


Fig. 1. Sample Cable Network Architecture with Multiple ISPs

When there are multiple ISPs connected to a cable network that uses typical routing techniques based on destination IP address at the regional head end, a subscriber's

packets are not always sent upstream to their preferred ISP. Many techniques have been proposed to address this problem [21]-[26]. Among those, *Source-address based routing* (SABR) and *encapsulation* have attracted the most attention.

With source-address based routing, each cable modem's source address is assigned from the IP number pool associated with the customer's ISP [22], [23]. The router at the cable head end forwards upstream packets to the appropriate ISP based on source IP address instead of destination address. In effect, the router at the regional head end acts as though it were N *virtual routers*, one for each ISP. Each virtual router handles the packets that have source addresses in their range.

With encapsulation, a cable user's data packets are put into another packet. Generally, the outer packet is addressed to the preferred ISP and the inner packet is addressed to the final destination. The router at the cable head-end sees only the outer packets and forwards accordingly. The ISP delivers the inner packet to its final destination over the Internet. Encapsulation enables tunneling, where outer packets support tunnels between a cable user and a tunnel-terminating device (e.g. Access Concentrator). Inner packets travel inside those tunnels; each cable user has a separate tunnel running transparently through the cable network to the ISP [22], with a unique session or tunnel IDs to identify it. The encapsulation mechanism can be managed by either the cable company or the ISP [23]. In the former case, the tunnel ends on the cable network; the cable network strips off the outer packet and forwards the inner packet to the correct ISP. In the latter case, the tunnel ends at the ISP; the ISP strips off the outer packet, and the cable company knows nothing about the inner packets.

PPP over Ethernet (PPPoE) [27] and Layer Two Tunneling Protocol (L2TP) [28] are the encapsulation protocols most discussed for cable networks. Both encapsulate Point-to-Point Protocol (PPP) frames [29]. Figure 2 depicts a sample PPPoE architecture, where CM#1 and CM#2 users subscribe to ISPA and ISPB respectively.

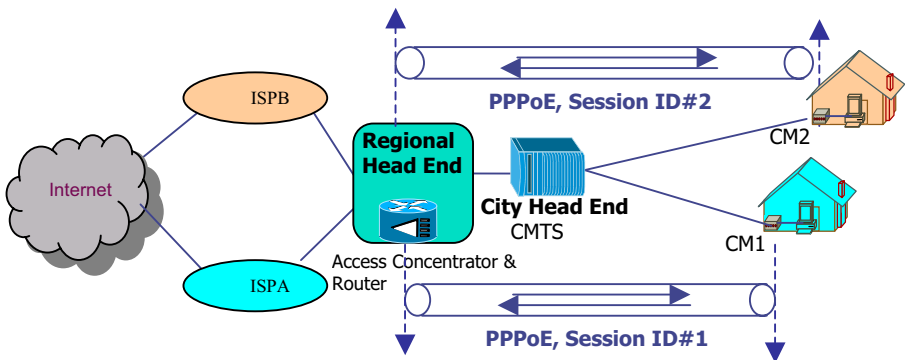


Fig. 2. Sample Cable Network Architecture with PPPoE

3 Multicast Problems and Solutions

3.1 Multicast Problems & Solutions with Source Address Based Routing

When source-address based routing (SABR) is used as the underlying technology in the cable network, the multicast transmissions pose different problems depending on the IGMP [30] *mode* of the DOCSIS devices (i.e. CM and CMTS), and whether an ISP can interpret multicast routing messages received from the cable network.

DOCSIS specifies two operational modes for the CMTS and CM devices: *active* and *passive* [31]. Active mode devices are able to terminate and initiate IGMP messages (e.g. send and receive a membership query message) like a multicast-enabled router [31]; passive mode devices cannot. Passive mode devices can only forward IGMP messages. If a DOCSIS-compliant device (i.e. CMTS or CM) acts as a router, it has to be in active mode; if it acts as bridge, it can be either active or passive.

If the CMTS is in *active mode*, multicast is not possible without protocol changes. The CMTS in active mode collects membership information from the cable modems via IGMP membership reports such as IGMP Join or Leave messages, and then summarizes these multicast activities by sending a message upstream to the regional head-end. Depending on the protocol deployed, this upstream message could be a Protocol Independent Multicast-Sparse Mode (PIM-SM) Join/Prune, a Distance Vector Multicast Routing Protocol (DVMRP) Prune/Graft, or a Multicast Open Shortest Path First (MOSPF) Link State Advertisement [32], [33]. The source address of a message sent upstream to the regional head-end is the address of the CMTS, so a source-addressed-based router at the regional head-end router cannot determine where to forward the packet. If it simply drops that packet, all multicast packets are lost.

There are two ways to make multicast possible. In one, which we call *Selective Forwarding*, the router at the regional head-end forwards upstream unicast packets based on source, and multicast packets sent by the CMTS (e.g. PIM-SM Join or DVMRP Prune) based on destination address. Whenever a new multicast connection is established, the router randomly selects an ISP, and forwards upstream multicast messages to that ISP for the duration of the multicast. Alternatively, this problem could be solved at the CMTS rather than at the regional head-end. Instead of using its own IP address as the source of membership reports sent upstream to the regional head-end (e.g. PIM-SM Join), the router at the CMTS could use the IP address of the first cable modem to subscribe to the multicast. The router at the regional head-end can then route this packet based on source address to an appropriate ISP.

Both schemes work if and only if the ISP understands these upstream multicast messages. If the cable network and the ISPs do not adopt the same protocols, this can be achieved by using an interdomain multicast protocol such as Border Gateway Multicast Protocol (BGMP) [34]. Alternatively, the router at the regional head-end could send an IGMP join message upstream to an arbitrarily selected ISP router instead of sending the multicast summary message.

With both of these approaches, it is preferable to send all packets through the ISP associated with the first cable modem to subscribe, even if that cable modem has since gone off line. Otherwise, the regional head-end may send upstream reports first to one ISP, and then to another, which would cause both ISPs to simultaneously forward the stream for extended periods.

All traffic from a given multicast stream would go through this same ISP. To be fair to all ISPs, the probability of selecting a given ISP should be proportional to the amount of multicast traffic received by that ISP's customers. If (and only if) the amount of traffic carried by an ISP is *consistently* out of proportion, then there is a fairness problem. This may occur when ISPs are selected randomly by the regional head-end, but the correct probability distribution is not known. Fairness should not be a problem if the matter is handled at the CMTS. The problem can also be prevented outside the cable network through a new ISP Signaling Protocol, where ISPs dynamically negotiate with each other to decide which ISP delivers multicast data to the customers of all of the ISPs. ISPs would exchange information about the ongoing multicast groups/applications and their members. However, the signaling protocol adds significant processing and communications overhead.

If the CMTS is operating in *passive mode*, multicast efficiency depends on the IGMP mode of the CMs. Consider the case where the two cable users in Figure 1 are associated with different ISPs and are interested in receiving the same multicast stream. If both users successfully send membership reports (MR) (i.e. IGMP Join) to their associated ISPs, then both ISPs will forward the multicast stream, causing each CM to receive two copies of each packet. If the CMs are in active mode, duplicate membership can be suppressed as follows. Every 10 seconds, customer devices send MRs upstream to the CM, where they can be forwarded upstream to the CMTS. When the CMTS receives an MR traveling upstream, the CMTS also forwards the MR on all downstream links, so all CMs see the MR. An active mode CM will suppress upstream transmission of an MR that merely duplicates an MR it has seen. Thus, multicast is efficient when CMs are active mode. In contrast, a passive mode CM will suppress an unnecessary MR only if it is holding the unnecessary MR in its buffer at the instant when it observes an equivalent MR from another CM. To solve the resulting redundancy problem that occurs when CMTS and CMs are passive mode, the Selective Forwarding Mechanism explained above could be deployed.

In summary, without protocol changes, multicast is efficient when CMs are active mode and the CMTS is passive mode, but multicast is inefficient when CMs and the CMTS are passive, and multicast fails entirely when the CMTS is active.

3.2 Multicast Problems & Solutions with Encapsulation Mechanisms

If the cable network uses a standard encapsulation scheme such as L2TP or PPPoE as described in Section 2, users will not be able to receive multicast packets without protocol modifications. With these protocols, the cable network cannot observe the multicast-related messages, such as membership reports. Consequently, when a router at the regional cable head-end receives a downstream multicast packet, the router will not recognize the multicast address, and will discard the packet.

There are two ways to make multicast work. The first is *multiple unicast*. With this technique, a router converts each downstream multicast packet that arrives from the Internet backbone to multiple unicast packets, and a packet is sent separately over the cable network to each of the interested users. If the tunnel-terminating device is within the ISP, then this conversion takes place at the ISP's gateway. If the tunnel-terminating device is within the cable network at the regional head-end, then this con-

version can take place either at the head-end or the ISP's gateway. This technique requires some changes in the networking algorithms of either the ISP network or the cable network. The disadvantage of this solution is the excess resources consumed by redundant transmission of duplicate packets.

The second solution is *discriminatory encapsulation*, where multicast traffic is no longer encapsulated in the upstream and the downstream directions, but unicast traffic is encapsulated. Whenever a user wants to join a multicast application, encapsulation software residing in the user's PC does not encapsulate multicast membership report packets (i.e. IGMP Join.) This allows the cable network to be aware of the interested multicast receivers. In the same way, the tunnel-terminating device does not encapsulate multicast packets in the downstream. Thus, the implementation of this scheme requires new encapsulation software that discriminates between multicast and unicast traffic. This technique is more efficient than multiple unicast. However, as was discussed in Section 3.1, the ISP must be able to accept upstream reports from the regional head-end, either because there is an inter-domain multicast protocol such as BGMP [34], both networks use the same multicast routing protocol, or the regional head-end sends IGMP messages upstream.

4 Blocking Multicast Traffic

The DOCSIS protocol suite guarantees that a cable company can choose to block certain multicast traffic as a matter of its network policy [31]. For instance, a cable company may use this ability to block video streaming, which potentially competes with the cable company's core business. A DOCSIS-compliant CMTS can encrypt downstream multicast traffic, so that a cable modem must request the decryption key from the CMTS. The CMTS blocks the stream by denying the request [35]. In such a case, some interested customers will not be able to receive multicast streams.

However, an ISP can provide its customers with a multicast service that the cable company cannot block by using the multiple unicast technique described in Section 3.2; an ISP unicasts each packet downstream to every cable customer that wishes to receive the multicast service. The redundant transmissions obviously create higher loads than would be seen with an efficient multicast mechanism.

For multiple unicast to work, CMs must send membership reports upstream such that they reach the ISP, without being terminated en route at the CMTS or the regional head-end. Thus, upstream membership reports should be encapsulated. In systems where all upstream traffic is encapsulated, multiple unicast is always an option for ISPs. The method cannot be used in cases where software at the customer premises only encapsulates non-multicast upstream traffic as described in Section 3.2, or where source-address based routing is used and no upstream traffic is encapsulated. Thus, an ISP can use the multiple unicast technique where all upstream traffic is encapsulated, or where the ISP (and not just the cable company) provides the customer with the relevant piece of software to encapsulate the upstream membership report.

5 Multicast Strategies that Advance Business Objectives

This section examines which multicast mechanisms are likely to be adopted, assuming that multicast traffic is abundant, so there is incentive for protocol enhancements that improve multicast efficiency. Nevertheless, cable companies and ISPs do not always maximize profit by cooperating in attempts to improve efficiency.

Section 5.1 describes the extent to which cable companies and ISPs would be motivated to deploy mechanisms that would influence the efficiency of multicast under the assumption that cable companies are not restricted by regulation, or influenced by fear of regulation. Sections 5.2 and 5.3 discuss the technical options available to cable companies and ISPs, and identify which options these companies would select based on these motivations, with SABR and encapsulation respectively. In Section 5.4, we show how results would differ if regulators prohibited cable companies from using the blocking mechanism described in Section 4. Results both with and without this prohibition on blocking are summarized in Table 1 at the end of this section.

5.1 Motivation for Technical Efficiency

The pricing of Internet services influences who pays the price for inefficiency. The prices charged by cable companies can be categorized as *usage-based* or *non-usage-based* [36]. With the former, price depends on the amount of data sent and received. With the latter, price is fixed, independent of the amount of data sent and received.

With usage-based pricing, ISPs have incentive to facilitate efficient multicast where possible, so multicast becomes less expensive. Cable companies share this motivation for efficiency, except for multicast applications that directly compete with a cable company's core revenue sources such as broadcast television and pay-per-view movies. A cable company may increase its profit by causing competing services (e.g. Internet video broadcasting) to be more expensive. Nevertheless, cable companies benefit from an efficient multicast capability, because DOCSIS guarantees them the ability to block, thereby preventing competing services from using this capability.

ISPs can still provide multicast for competing services through multiple unicast, but the resulting redundancy makes multicast more expensive. Alternatively, it is conceivable that some networks may someday have the ability to offer discriminatory prices for different multicast streams, perhaps based on source address; in this case, the cable company could charge competing services more instead of blocking.

With non-usage-based pricing, cable companies always have incentive to make multicast efficient because efficiency reduces resource consumption without affecting the cable company's revenue. In contrast, ISPs have much less incentive to make multicast efficient, because they do not pay for the inefficiency. One implication is that cable companies should choose not to block multicast streams, because ISPs have the ability to circumvent this feature using the multiple unicast technique, and the resulting inefficiency increases the cost to cable companies and not to ISPs.

5.2 Source Address Based Routing

As discussed in Section 3.1, multicast traffic poses different problems in different architectural configurations. In each case, the cable company can address the problems without aid from an ISP. When the CMTS is in active mode, regardless of the IGMP mode of the CM devices, upstream multicast reports sent by the CMTS are dropped at the regional head-end so that the subscribers cannot receive multicast data. The cable company can solve this problem at the cable head-end by inserting the IP address of a subscribing cable modem as the source of membership reports sent upstream instead of using its own IP address. Alternatively, it can solve the problem at the regional head-end by using selective forwarding. If there is no interdomain multicast routing protocol, the cable company can configure the router at the regional head-end to send IGMP messages upstream. When both the CMTS and the CM are in passive mode, redundant transmissions can occur. The cable company can address this inefficiency problem via Selective Forwarding at the regional head-end. In networks where the CMTS is passive and CMs are active, multicast is always efficient.

As discussed in Section 5.1, a cable company has incentive to provide an efficient multicast service with both usage-based and non-usage-based pricing. Thus, regardless of whether the CMTS and the cable modems are in active or passive mode, all multicast is likely to be efficient unless the cable company blocks. However, as described in Section 5.1, cable companies have incentive to block competing services when there is usage-based pricing, thereby forcing ISPs to carry these competing services using the multiple unicast technique. With non-usage-based pricing, cable companies will not block, and all multicast will be efficient.

5.3 Encapsulation

As described in Section 3.2, users cannot receive multicast unless either multiple unicast or discriminatory encapsulation are adopted; both require protocols modifications. Whoever provides encapsulation software to the consumer decides whether to use discriminatory encapsulation. This could be either the cable company or the ISP. Moreover, discriminatory encapsulation requires either an interdomain multicast routing information exchange, or a change in the way the cable company's router at the regional head-end handles multicast packets, or for cable network and ISPs to use the same multicast routing protocol.

In the case where the cable company provides the encapsulation software, the cable company can choose discriminatory encapsulation, which is more efficient than multiple unicast, without aid of the ISP. As discussed in Section 5.1, a cable company has incentive to offer an efficient multicast service with either usage-based or non-usage-based pricing. Thus, cable companies would choose discriminatory encapsulation.

With non-usage-based pricing, where cable companies would not block, it is likely that all multicast traffic will use this efficient approach. In contrast, with usage-based pricing, cable companies have reason to block multicast from competing services, so an ISP has reason to use multiple unicast for this traffic. An ISP can do this without aid from the cable network. Consequently, multicast traffic from competing services will be carried inefficiently, and other multicast traffic will be carried efficiently.

In the case where the ISP provides the encapsulation software, cable companies acting alone cannot make discriminatory encapsulation work. Neither can ISPs, except in the case where the cable network and ISPs happen to use the same routing protocols. With non-usage-based pricing, cable companies prefer efficient solutions, but an ISP may or may not be willing to help. ISPs would accept the inefficiency of multiple unicast because inefficiency does not affect their cost, and ISPs would accept the vulnerability of discriminatory encapsulation to blocking, because a smart cable operator would not block in this case. The result may therefore depend on outside factors. For example, if a network has already been deployed without good support for multicast, it is easier to upgrade using inefficient multiple unicast, because this requires no change to software at the customer premises.

With usage-based pricing, a cable company would have incentive to block competing multicast services, so ISPs have incentive to deploy multiple unicast. To avoid the higher cost of multiple unicast due to redundant packet transmissions, the ISP may even choose to adopt two solutions: multiple unicast for streams that are blocked by the cable company, and the more efficient discriminatory encapsulation for other multicast traffic. The cable network does have a little incentive to support discriminatory encapsulation, perhaps with an interdomain exchange of multicast routing information. As above, whether this is enough incentive may depend on outside factors.

5.4 When Blocking Is prohibited

It is possible that regulators will prohibit cable companies from blocking competing services (e.g. Internet television), or they will use the threat of possible regulation to deter cable companies from blocking. This section evaluates the impact of such a restriction, assuming that it is the only regulatory restriction. As in Sections 5.1 to 5.3, we assume that the amount of multicast traffic is significant.

As discussed in Section 5.1, cable companies have no incentive to block when there is non-usage-based pricing, because ISPs can always resort to multiple unicast. Thus, with non-usage-based pricing, the results shown in Section 5.2 and 5.3 should be as appropriate with a prohibition on blocking as they were without the prohibition.

With usage-based pricing, there is a difference. Cable companies would prefer that multicast be inefficient and therefore expensive for competing services, but efficient for all other applications. When blocking is prohibited, they can no longer discriminate. If most multicast traffic comes from competing services, a cable company may prefer that all multicast be inefficient. If competing services are not a serious problem, and people are more likely to become cable modem customers if they can get a low price for noncompeting services that use multicast, then a smart cable company would make multicast efficient. A cable company would also be more inclined to provide efficient multicast when blocking is prohibited if the cable company is affiliated with an ISP that would benefit from the efficiency. While cable companies may or may not choose the most efficient option, recall from Section 5.1 that ISPs always prefer efficient multicast when there is usage-based pricing.

With SABR, multicast is always efficient when the CMTS is passive mode and the CMs are active mode. When the CMTS is in active mode, regardless of the mode of the CM devices, multicast is not possible unless the cable company adopts non-

standard solutions, as proposed in Section 3.1. A cable company could choose to do nothing to solve the problem. This would force ISPs to offer a less efficient multicast using multiple unicast. The cable company benefits from choosing an efficient solution. With usage-based pricing, a profit-maximizing cable company might choose either option, depending on demand for competing services. With non-usage-based pricing, the cable company benefits from choosing an efficient solution. The cable company has a similar choice to make if the CMTS and cable modems are both in passive mode. The cable company might choose to support efficient multicast through Selective Forwarding, as described in Section 3.1. The cable company might also choose not to do this, so when customers of different ISPs want to receive the same multicast stream, each ISP may forward the stream. This is inefficient, but not as inefficient as multiple unicast, where the stream is forwarded once for every recipient cable modem. Thus, in this case, a prohibition on blocking improves efficiency for competing services while degrading efficiency for noncompeting services.

With encapsulation, multicast is efficient or it is inefficiently transmitted via multiple unicast. With usage-based pricing, ISPs will always prefer an efficient approach. However, as demonstrated in Sections 3.2 and 5.3, the ISP cannot do this without help from the cable company except in the case where the cable network and ISPs happen to have adopted the same multicast protocols. Once again, the cable company may or may not prefer the efficient solution, depending on the impact of competing services on its core business. Thus, either outcome is possible.

Table 1. Summary of Probable Deployments from Sections 5.1 to 5.4

Forwarding Mechanism	USAGE-BASED PRICING		NON-USAGE-BASED
	No Prohibition	Blocking Prohibited	
Encapsulation (provided by ISPs)	B or D*	A or B*	A or B*
Encapsulation (provided by cable co.)	D	A or B*	A
SABR (Passive CMTS, Passive CMs)	D	A or C*	A
SABR (Passive CMTS, Active CMs)	D	A	A
SABR (Active CMTS)	D	A or B*	A
A. Multicast traffic is transmitted efficiently.			
B. Multicast traffic is inefficiently transmitted via multiple unicast. The number of copies equals the number of customers who want the stream.			
C. Multicast traffic is inefficiently transmitted. The number of copies equals the number of ISPs that have at least one customer who wants the stream.			
D. Competing multicast services are carried inefficiently via multiple unicast and other multicast traffic is carried via efficient mechanisms.			
* Depends on factors other than technical design.			

6 Summary

When multiple ISPs are providing services over cable networks, a forwarding mechanism is needed to deliver the packets traveling upstream to the correct ISP. The principal forwarding mechanisms that are currently under consideration are source address based routing (SABR) and encapsulation [20]-[23]. With these protocols, there are cases where either multicast does not work or works inefficiently.

If one of the widely discussed encapsulation protocols for cable networks (i.e. L2TP, PPPoE) is adopted, subscribers are not able to receive multicast streams, because a cable network cannot observe the presence of multicast membership reports. This problem can be solved in two ways: 1) multiple unicast, where a router, converts multicast packets to multiple unicast packets and then transmits them downstream; 2) discriminatory encapsulation, where unicast packets are encapsulated but multicast packets are not in the upstream and the downstream directions. The latter is less efficient. Both require that the ISP be able to understand the cable company's multicast packets, perhaps through an inter-domain-multicast routing protocol such as BGMP, or by translating upstream multicast routing packets into IGMP messages at the regional head-end.

If source-address based routing is the forwarding mechanism and the IGMP modes of the CMTS devices are active mode, multicast fails because the regional head end drops a received upstream multicast report (e.g. PIM-SM Join or MOSPF Link State Advertisement) that is sent by a CMTS. This problem can be solved either at the regional head-end or at the CMTS. At the regional head end, the router could filter the upstream multicast messages and forward the packet to a randomly selected ISP whenever a new multicast connection is established, which is *Selective Forwarding*. When the problem is solved at the CMTS, the CMTS would use the IP address of a subscribing cable modem as the source of membership reports sent upstream to the regional head-end instead of using its own IP address. Neither approach guarantees fairness to the ISPs; one ISP may consistently carry most of the multicast traffic. We have described mechanisms a cable company could adopt to solve this fairness problem. On the other hand, if both the CMTS and the CM are passive mode, then multicast works but users can receive more than one copy of every downstream multicast packet. This redundancy can be prevented at the cable network by adopting the above Selective Forwarding mechanism for IGMP-Joins. In only one case, where the CMTS is in passive mode and CMs are in active mode, standard multicast works efficiently.

Additionally, we have shown that DOCSIS gives a cable company the ability to block any multicast stream it wishes, including video multicast that competes the company's core business. However, an ISP can provide its customers a multicast service that the cable company cannot block through multiple unicast.

Assuming that multicast becomes quite common, cable companies and ISPs would be motivated to provide multicast services. Nevertheless, cable companies and ISPs are not always motivated to offer an efficient mechanism. We first consider the case where cable companies are free to block. Motivation depends on whether pricing is usage-based or non-usage-based. When pricing is *usage-based*, cable companies would be motivated to block multicast traffic associated with competing services, causing this traffic to be carried inefficiently via multiple unicast. All other multicast traffic would be carried with an efficient multicast service, except possibly in the case

where encapsulation is used and ISPs provide the encapsulation software. In this case, it is possible that all traffic will be via multiple unicast. When pricing is *non-usage-based*, cable companies have no incentive to block because ISPs can adopt multiple unicast, which hurts cable companies. Traffic would be carried over an efficient multicast service, with the possible exception of the same case where the ISP provides encapsulation software. In this case, the ISP and cable company must work together to provide efficient multicast, which may or may not happen.

A cable company may be precluded from blocking by regulation or fear of regulation. Such a prohibition would be unnecessary with non-usage-based pricing, because cable companies would not block anyway. With usage-based pricing, there are architectures where companies provide efficient multicast. However, there are also architectures where the cable company may or may not be motivated to do so; it depends on whether cable companies profit more from making competing services expensive or from making other services inexpensive. If a cable company chooses not to support efficient approach, then inefficient and expensive multicast will not be limited to competing services, as was the case without the prohibition. Thus, this regulation could actually increase the cost of multicast applications to consumers.

References

1. H. Jinzenji and K. Hagishima, "Real-time audio and video broadcasting of IEEE GLOBECOM '96 over the Internet using new software," *IEEE Communications Magazine*, Vol.: 35, Issue: 4, April 1997 pp. 34–38.
2. "MediaChannel: Your Daily Guide to Internet Television & Video," *MediaChannel*, www.mediachannel.com
3. D. Gunzerath, "Radio and the Internet," *The National Association of Broadcasters*, May 5, 2000, www.nab.org/Research/Reports/RadioandInternet/sld015.htm
4. J. Careless, "Digital Television - JumpTV Fights For Retransmission," *Digital Television.com* Jun 1, 2001, www.digitaltelevision.com/2001/webcast/0801_1.shtml
5. "Broadband Subscriber Count Nears 20 Million - December 2002," *Kinetic Strategies Inc. Cable Datacom News* Dec. 1, 2002, www.cabledatcomnews.com/dec02/dec02-1.html
6. "FCC Approves AOL-Time Warner Merger," *The Federal Communications Commission Media Bureau* Oct. 25, 2001, www.fcc.gov/transaction/aol-tw-decision.html
7. Federal Communications Commission, "Annual Assessment of the Status of Competition in the Market for the Delivery of Video Programming," Ninth Annual Report, FCC Docket MB 02-145, Dec. 31, 2002.
8. R. Mark, "Cox Begins Its First Open Access Broadband Trials," *Internetnews.com* Nov. 6, 2001, www.internetnews.com/isp-news/article.php/8_917471
9. C. Grice, "AT&T to test open access to ISPs," *CNET News.com* June 7, 2000, news.com.com/2100-1033-241559.html?legacy=cnet
10. E. Joyce, "AT&T, Comcast Commit to ISP Access," *Internetnews.com* April 1, 2002, www.internetnews.com/isp-news/article.php/8_1001211
11. M. Ingebretsen and M. Siegel, "Cable slouches toward open access," *IEEE Spectrum*, Vol. 38, No. 4, April 2001, pp. 74–79.
12. D. Menard and T. Denton, "Third Party Access to Cable Modems in Canada," Written for the FCC Open Access NOI, Docket No. 00-185. Version 2.0 Dec. 1, 2000
13. "CableLabs[®] Certified[™] Cable Modems Technical Specifications," *The CableLabs[®] Certified[™] Cable Modem project*, DOCSIS, May 2002, www.cablemodem.com/specifications

14. S. Olsen, "Battle brews over Web streaming," *CNET News.com* May 31, 2001, news.com.com/2100-1023-267638.html
15. M. A Lemley and L. Lessig, "Written Ex Parte of Professor Mark. A. Lemley and Professor Lawrence Lessig In the Matter of Application of consent to the transfer of control of licenses MediaOne Group, Inc. to AT&T Corp," FCC CS Docket No. 99-251, Nov. 10, 1999
16. J. B. Speta, "The Vertical Dimension of Cable Open Access," *University of Colorado Law Review*, Vol. 71, No. 4, 2000, pp. 975-1010
17. T. W. Hazlett and G. Bittlingmayer, "'Open access:' the ideal and the real," *Telecommunications Policy*, Vol. 26, No. 5-6, 7 June 2002, pp 219-359
18. D. Fellows and D. Jones, "DOCSISTM cable modem technology," *IEEE Communications Magazine*, Vol. 39, No. 3, March 2001, pp. 202-209.
19. V. C. Majeti (Ed.), "Cable modems: current technologies and applications," *Intl. Engineering Consortium and Institute of Electrical and Electronic Engineers*, IEEE Press, 1999.
20. D. Farber, S. D. Dukes, and M. E. Laubach, "Delivering Internet Connections over Cable, Breaking the Access Barrier," *Wiley Computer Publishing, John Wiley & Sons Inc.* 2001
21. S. O'Donnell, "Broadband Architectures, ISP Business Plans, and Open Access," *Telecommunications Policy Research Conference (TPRC)* Alexandria, VA. Sept. 23-25, 2000.
22. M. E. Laubach, "Comments on the Technical Ability to Implement Open Access Provisioning via High-Speed Data over Hybrid Fiber-Coaxial Cable Television Systems in the United States," Prepared for the White House National Economic Council, May 30, 1999, www.matmos.com/whitepapers/openaccess990530.pdf
23. E. Tseng, "A Capital Cost Model For Open Access Cable Networks," *MIT Program on Internet & Telecoms Convergence Consortium (ITC)*, MIT June 7, 2000, itc.mit.edu/itel
24. E. Tseng, and Sharon Gillett, "Open Access to Cable Data Networks," *MIT Program on Internet & Telecoms Convergence (ITC)*, 2000, itel.mit.edu/itel/publications.html
25. "AOL OK With Policy-Based Routing - FEB. 2001," *Kinetic Strategies Inc. Cable Datacom News* Feb. 1, 2001, www.cabledatcomnews.com/feb01/feb 01-5.html
26. J. Fijolek, "Access Methods for Support of Multiple Internet Service Providers Over DOCSIS," *CommWorks Corporation*, a 3com Company, White Paper, June 4, 2000.
27. L. Mamakos *et al.*, "A Method for Transmitting PPP Over Ethernet (PPPoE)," *Internet Engineering Task Force Request for Comments 2516*, Feb. 1999, www.ietf.org/rfc/rfc2516.txt
28. W. Townsley *et al.*, "Layer 2 Tunneling Protocol (L2TP)," *Internet Engineering Task Force Request for Comments 2661*, August 1999, www.ietf.org/rfc/ rfc2661.txt
29. U. D. Black, "PPP and L2TP: remote access communications," *Prentice Hall*, 2000.
30. S. Deering *et al.*, "Internet Group Management Protocol," *IGMP*, version 3, *Internet Engineering Task Force Request for Comments 3376*, Oct. 2002, Proposed Standard
31. "DOCSIS Radio Frequency Interface (RFI) Specification," *The CableLabs Certified Cable Modem project* Dec. 31, 2001, www.cablemodem.com/Specs/SP-RFIV2.0-101-011231.pdf
32. B. Mukherjee and L. H. Sahasrabudde, "Multicast Routing Algorithms and Protocols: A Tutorial" *IEEE Network*, vol. 14, no. 1, pp. 90-102, Jan.-Feb. 2000.
33. M. Ramalho, "Intra- and Inter-Domain Multicast Routing Protocols: A Survey and Taxonomy," *IEEE Communications Surveys & Tutorials*, First Quarter 2000, Vol. 3, www.comsoc.org/livepubs/surveys/public/1q00issue/ramalho.html
34. D. Thaler, D. Estrin, and D. Meyer, "Border Gateway Multicast Protocol (BGMP): Protocol Specification," *IETF Draft*, Nov. 22, 2000, work in progress, netweb.usc.edu/bgmp/draft-ietf-bgmp-spec-02.txt, www.ietf.org/internet-drafts/draft-ietf-bgmp-spec-03.txt
35. "DOCSIS Baseline Privacy Plus Interface (BPI+) Specification," *CableLabs Certified Cable Modem project*, Mar. 1, 2002, www.cablemodem.com/Specs/SP-BPI+-I08-020301.pdf
36. Q. Wang, J. M. Peha, M. Sirbu, "Optimal Pricing for Integrated-Services Networks," in *Internet Economics*, MIT Press, 1997, pp. 353-376, www.ece.cmu.edu/~peha/papers.html