

AN IMPROVEMENT PROPOSAL FOR THE TIMING AND SCALING OF DOCSIS IP MULTICAST SERVICES

William T. Hanks, Tom Cloonan, Amit Eshet, Jeff Howe,
John Ulm, Ian Wheelock
ARRIS

Abstract

This paper will present a proposal for an extension to the DOCSIS protocol to facilitate the use of a “DOCSIS Multicast Carousel” table which periodically transmits DOCSIS-specific IP multicast forwarding information for each of the currently active multicast group addresses within a MAC Domain Downstream Service Group (MD-DS-SG). The CM may utilize this table of information to rapidly change its own multicast filters (and, if necessary, quickly retune one or more pre-designated downstream channel receivers) in response to subscriber IP video channel change requests without requiring additional control signaling to or from the CMTS. When the subscriber has settled on a program selection for a provisioned timeframe, the final selection and a list of interim selections may be signaled to the CMTS for metrics collection. This proposal is being submitted in an effort to mitigate several issues that have been identified as MSOs and vendors continue to explore IP Video Multicast deployment scenarios. These issues include the large bandwidth requirements that may be needed to support unicast bursts for fast-channel change assistance. The issues also include the high processing and message exchange rates (between CMs and CMTSs) that may be required when large numbers of subscribers initiate simultaneous channel change requests,

PROBLEM BACKGROUND

Viewer Behavior

Television program selection menu systems have been around for many years - allowing a direct selection of a single channel. In spite of this, many television viewers have the habit of using a television or set-top box remote control to rapidly switch between multiple channels – using either the channel up/down buttons or a “next favorite channel” button - before settling on a program. This behavior tends to occur en-mass at half-hour intervals as a previous program ends and a new one begins. This behavior can also occur at other moments during the hour when different network broadcasters tend to have un-planned synchronization between their respective commercial breaks.

IP Multicast Transport Technique

Modern Internet Protocol (IP) Video technology utilizes the concept of IP multicast to distribute a linear (continuous, sequential programming) video stream transmission from a single video IP source to a potentially large number of subscribers over a wide area in an efficient manner. The use of IP multicast allows a transmission to only be carried to the subscribers which have requested (and are entitled to receive) the video program using only the transmission links which are necessary to reach them. Furthermore, if multiple endpoint clients have access to the same transmission link (as they do in a shared medium technology such as Ethernet or DOCSIS), then these endpoint clients can each receive (and share) the same transmission simultaneously. This sharing of

multicast transmissions can result in great bandwidth savings within the DOCSIS network, because it eliminates the wasted bandwidth that would otherwise be created by the many simultaneous unicast transmissions of identical video content to the multiple endpoint clients within a common Service Group. In some cases, up to 80% bandwidth savings can be realized in IP Video distribution networks through the use of IP multicast transmissions (in place of the less-efficient IP unicast transmissions). These dramatic improvements in bandwidth efficiency have made IP multicast a very important technology within most of the MSO IP Video distribution networks that are currently being planned for the future.

DOCSIS IP Multicast

The majority of modern cable high speed data transmission systems follow the Data Over Cable System Interface Specification (DOCSIS) protocol [1] [2] [3]. The cable modems (CMs) that predate the release of version 3.0 of DOCSIS (pre-3.0 DOCSIS) relied on the snooping of IGMPv2 messaging by the CM. The DOCSIS 3.0 protocol, in an attempt to make the cable modem be multicast protocol agnostic, moved the multicast control plane hooks from the CM device to the Cable Modem Termination System (CMTS). This approach was taken in an attempt to simplify the cable modem operation and to reduce the overall cost of deploying multicast solutions.

This newer DOCSIS 3.0 multicast architecture requires that the CMTS label all packets which are part of a particular multicast group transmission with a 20-bit Downstream Service ID (DSID). The CM uses this DSID to identify the packets which must be forwarded as part of this multicast transmission. The architecture relies on the successful runtime execution of a heavyweight, fault-tolerant, three-way Dynamic Bonding Change (DBC) signaling message transaction between the CMTS and

the CM for each new multicast group service flow that the individual subscriber device wishes to receive. This exchange communicates the DSID, the downstream channel set over which the multicast transmission will be sent, and optionally, some packet resequencing parameters and forwarding directives to the CM.

The Perfect Storm

The scenario intersection of the en-mass rapid channel-changing behavior of video viewers and the overhead associated with the DOCSIS 3.0 multicast DSID forwarding architecture may cause problems both in viewer expectations and system capabilities. The signaling associated with a sequence of rapid channel changes may cause previously unexperienced latency for each single subscriber viewer. The signaling overhead for a large number of subscriber viewers performing rapid channel changes at the same time may also cause scaling issues for the DOCSIS CMTS system.

In addition, many MSOs are proposing to use unicast bursts of Adaptive Bit-Rate (ABR) video traffic to rapidly fill the video buffers within their video client devices to provide Fast Channel Change operations. The simultaneous transmission of these unicast bursts of ABR video traffic when many clients request channel changes at the same time can aggregate together to create very high transient bandwidth requirements within the DOCSIS network. Architecting the network to support the bandwidth needs of these transient bursts may yield higher-cost, higher-bandwidth systems with very low average utilization levels- which may be viewed as cost-prohibitive. Architecting the DOCSIS networks with lower bandwidth levels that ignore these unicast ABR bandwidth bursts may, however, lead to periodic reductions on subscriber Quality of Experience levels for all services, as the simultaneous unicast ABR bandwidth bursts

will undoubtedly lead to periodic congestion within the DOCSIS network.

In a recent study of channel change behavior of more than 100,000 tuners (active video clients + digital video recorders) over several weeks with an available catalog of between 300 and 600 video programs across multiple operators and multiple video technologies, we observed that the worst case service group with 500 tuners had as many as 104 channel changes over an 8 second window, 95 channel changes over a 4-second window, and a peak of 50 channel changes within a one second window. While channel change statistics for larger populations of tuners (multiple service groups serviced by the same device) do not scale linearly, a population of almost 30,000 tuners can be expected to require a maximum of about 700 channel change operations in a one second interval.

If the servicing device has DOCSIS 3.0 CMTS functionality, then each of these channel change operations may result in a multicast group membership LEAVE operation (to leave the channel previously viewed) followed by a multicast group membership JOIN operation. In turn, each LEAVE or JOIN operation requires a successful three-way DBC signaling transaction between the CMTS device and the requester's CM. Thus, each single channel change event within a subscriber's home produces a total of eight protocol exchanges between the CM and the CMTS, so the 50 channel changes per second that are expected for a small 500-tuner Service Group can result in a total of 400 protocol exchanges per second between the CM and the CMTS (which equates to a protocol exchange every 2.5 milliseconds on average). A CMTS with many Service Groups would therefore experience much higher protocol exchange rates. The CMTSs must also perform multicast routing protocol exchanges with

their Northbound network to initiate the IP multicast flows.

Shell Game for Protocol-agnostic CMs

As mentioned previously, the DOCSIS 3.0 MAC protocol moved the multicast control plane hooks from the CM device to the Cable Modem Termination System (CMTS) device in an effort to simplify the cable modem operation and to reduce the overall cost of deploying multicast solutions. These are very reasonable goals to strive for. However, the architectural solution adopted by the DOCSIS 3.0 protocol moved the requirement for multicast protocol snooping (IGMPv2 at the time but now extending to IGMPv3 and MLD v1/v2) from the CM device to the CMTS. In doing so, the CMTS must not only look for every possible IP multicast signaling packet in the hardware-based data stream and send them all to a control-plane processor (an activity which is not at all desirable when vendors are being asked to lower infrastructure costs on a multiple-gigabit router such as the CMTS device), but the processor must then also ensure that the requested multicast group media stream is available at the CMTS NSI (possibly issuing a multicast routing protocol message to the northbound router cloud) and then execute a heavyweight, fault-tolerant (possibly requiring message retransmissions), three-way Dynamic Bonding Change (DBC) signaling message transaction over a cable upstream transmission medium which is prone to noise.

What this multicast processing shell game has done is concentrate the vast majority of the multicast processing requirements onto a single device- the CMTS. Furthermore, the MSO demands for increased scaling on this device means that this single CMTS device must be engineered to address the peak scaling requirements of all multicast requestors across the entire CMTS; not just the much lower average processing requirements. This need for engineering to

peak rates seems to be in conflict with demands for decreased infrastructure costs.

While IGMP (for IPv4 networks) and MLD (for IPv6 networks) are the protocols typically used for signaling multicast group membership requests, these are not the only ways that the CMTS might be signaled to add a particular device to an IP multicast group. IP Multicast signaling extensions have been added to the PacketCable Multimedia protocol [4]. Additionally, vendor-proprietary extensions may exist whereby a CMTS or other device might snoop a different type of media request (perhaps HTTP-Get request from a unicast adaptive bitrate (ABR) request?) from a customer video device and then signal the CMTS in another way to initiate a multicast flow.

Regardless of the signaling mechanism utilized to initiate the multicast stream transmission to each multicast group client, the same heavyweight DBC transaction must be used in order to communicate the DOCSIS MAC-level details for the IP multicast stream from the CMTS to the CM.

In addition, the simultaneous transmission of unicast ABR bursts for Fast Channel Change operations can lead to increased bandwidth requirements within the DOCSIS network.

A single, simple technique for mitigating these two fundamental IP Video problems (heavy CMTS processing loads and high DOCSIS bandwidth requirements) may be beneficial for the cable industry to consider at this point in time.

THE PROPOSED SOLUTION

SDV Inspiration

The authors propose borrowing a solution from the Switched Digital Video (SDV) solution space by implementing a DOCSIS

Multicast Carousel (DMC) which is periodically transmitted and is received by all cable modems within a service group. In keeping with a protocol layering architecture, the DMC may well be published in concert with and in support of any application-layer IPTV signaling carousel. The contents of each Carousel type-length-value structure (described later) consists of the DOCSIS data which is necessary for the cable modem to receive and forward the contents of exactly one IP multicast group.

Architecture Assumptions

The authors have made the following assumptions about the DOCSIS IP Video system:

- 1) In order to eliminate the need for extra group replications and the need for channel reassignment for group reception, all Multicast Group sessions will be carried together on as few DOCSIS downstream channels as possible. For this discussion, we will call the number of downstream channels N.
- 2) All embedded or non-embedded cable modem devices that are to carry multicast IP Video are compatible with the GMAC Promiscuous DSID mode of Multicast DSID forwarding of the DOCSIS specifications and are capable of receiving at least N downstream channels. Ideally, these devices are capable of bonding all multicast group flows over the N channels as well.
- 3) The IP video application layer signaling will somehow cause the subscriber's CM device (either standalone CM or embedded within a home gateway) to attempt to join a particular IP multicast group session. This application layer signaling may be triggered by the IP Video player device, by a video application running on a home gateway,

some application sending directives from the network cloud, or any other possible method.

- 4) The IP video application layer signaling is NOT commingled with the contents of the DOCSIS Multicast Carousel. This way the DOCSIS Multicast Carousel can be legitimately defined as a DOCSIS MAC layer extension.
- 5) The Receive Channel Configuration (RCC) mechanism of DOCSIS by which a CMTS assigns channels to CMs can and will be augmented in a manner to designate a number of receivers to be under the control of the CM (i.e. not directly assigned by the CMTS) for the purposes of tuning to IP Multicast groups in a manner to be described later.

MAC Subinterface Access Control List

The Multiple Services Operator (MSO) may choose to make different IP Video multicast group sessions available to different sets of CMs based upon cable-side topology. Specifically, the CMs that are associated with one MAC Domain Downstream Service Group (MD-DS-SG) might be allowed to receive a slightly different set of multicast sessions from the CMs associated with a different MD-CM-SG. An Access Control List (ACL) with permit/deny directives per group IP address might be assigned to each MD-DS-SG. Ideally, this ACL might be assigned with directives per source IP, Group IP (S, G) tuple.

In order to facilitate this ability, a set of configuration directives may be required within the CMTS to map an MD-CM-SG to a MAC Domain subinterface. Furthermore, each MAC Domain subinterface might present a different cable-helper address to the Dynamic Host Configuration Protocol (DHCP) server for the CMs of a MD-DS-SG once the CMTS determines (or is told by the

CM as a result of CM topology resolution) the MD-DS-SG identifier of the CMs.

The DOCSIS Multicast Carousel

In order to define the parameters necessary for a DOCSIS Multicast Carousel, we begin by studying the parameters of the communications that take place when a television subscriber tunes an IP Video Player device to receive program content data over an IP multicast group using a DOCSIS 3.0 CM device.

Table 1 contains the structure of the payload of a DOCSIS Dynamic Bonding Change Request (DBC-REQ) message which was used by a CMTS to successfully install a resequencing DSID for a DSID-forwarded multicast group. The payload is in the Type-Length-Value (TLV) format that is specified in [1]. Items which are not necessary for the DMC are stricken with the explanation to follow.

Type	Len	Type	Len	Type	Len	Type	Len	Value
50	N							DSID Encodings
		1	3					DSID value (20 bits)
		2	4					Downstream Service Identifier Action: Add
		3	N					Downstream Reseq. Encodings
				4	4			Reseq. DSID flag (1=Resequencing-DSID)
				2	n			DS Channel ID (byte) array [DCID1, DCID2, ..., DCIDn]
				3	4			DSID-reseq-wait time (1-180) x 100 μ s
				4	4			DSID Reseq. Warning Threshold (1-179) x 100 μ s

Type	Len	Type	Len	Type	Len	Type	Len	Value
				5	2			CM STATUS Hold-off Timer for Out-of-range events (in 20 ms units)
		4	N					Multicast Encodings
				4	N			Client MAC Addr Encodings
						4	4	Action (0=add; 1=delete)
						2	6	Client MAC Address joining or leaving
				2	4			Meast CM Interface Mask bitmap
				3	N			Meast Group MAC Address array {GMAC0, GMAC1,...,GMAC n}
31	1							Key Sequence Number
27	20							HMAC Digest

Table 1: Contents of a DBC-REQ message to install a multicast DSID

Note: A typical shorthand used when discussing TLV subtypes is to list the cascading types (ignoring the intervening lengths) with dotted notation beginning with the main type. For example, the DSID Reseq. Warning Threshold (Table 1) would be known as TLV type 50.3.4

No Resequencing

If we assume that all multicast flows are UDP-based and are therefore resequenced by an upper layer application protocol then TLV type 50.3.1 becomes unnecessary. Optional parameters (TLV Types 50.3.3 and 50.3.4 in light grey) are also resequencing parameters and will not be included in the DMC.

No Individual CM Directives

Individual CM directives (TLV Types 50.4.2 and 50.4.3 in medium grey) are sent from the CMTS to the individual CM to tell the CM upon which interface to forward the multicast group packets. The information in these TLVs can be worked around by the CM

noting upon which CM interface the video client is requesting the JOIN operation so these messages are not appropriate for the DOCSIS Multicast Carousel.

No Action Control

If we assume that the carousel message would include information for multicast groups which are available and would not include information for multicast groups which are not, then the “action” parameters of TLV Type 50.2 and 50.4.1.1 become unnecessary.

Since the goal of this feature is to have the CM receive the multicast request and, provided that the multicast group stream is available, to have enough information to process the JOIN; the information about the Client MAC can be determined from the interface upon which the request itself was received and the entire 50.4 TLV branch becomes unnecessary.

Since the CMTS is the only device which can issue a DMC message on the downstream, authentication of the sender via the HMAC-Digest is not necessary. Remember, IP video content itself is expected to be protected via end-to-end encryption.

Finally, since the TLV 50.3.2 is the only piece of information remaining under 50.3, TLV 50.3.2 can be promoted one level for the DOCSIS Multicast Carousel and 50.3 is not needed.

The remaining fields (in shaded rows) are necessary information that will form the basis for the information carried within an entry of the Multicast Carousel.

An IGMP request will have a Group IP address with an optional Group Source IP address (for Source-Specific Multicast or SSM). If the multicast stream is not source specific then the Group Source IP will be 0.

These are the database keys (not to be confused with the security key) that the CM will use to find the correct DSID and Key Sequence Number in the Carousel.

In the end, the DOCSIS Multicast Carousel message will consist of multiple TLV Type 1 messages as shown in Table 2.

Type	Len	Type	Len	Type	Len	Value
1	N					Multicast Group Carousel Encoding
		1				Group ID (GIP)
		2				Group Source IP (optional for SSM)
		3	N			DSID Encodings
				1	3	DSID value (20 bits)
				2	n	DS Channel ID (byte) array [DCID1, DCID2, ..., DCIDn]
	4	1				Key Sequence Number

Table 2: Proposed contents of a Multicast Group Carousel Encoding entry for one multicast group session

The Multicast Group Carousel Encoding TLV Type 1 will be repeated within the carousel messages for each multicast group to be described. It is anticipated that all multicast groups which are active (or statically provisioned) within the CMTS and available (per ACL configuration) to CMs within the MD-DS-SG will be reflected in the DOCSIS Multicast Group Carousel message.

The multi-part Multicast Group Carousel (MGC) message will consist of one or more numbered fragments consisting of DOCSIS frames; each containing a number of Multicast Group Carousel Encoding TLVs. The MGC message fragments will include a configuration change count so that CMs monitoring the stream know when the stream has been updated. All MGC fragments from the same MAC Domain IP Address containing the same Configuration Change Count and MD-DS-SG-ID belong to the same MGC message. The CM MUST successfully

receive all fragments of an MGC message before using the contents of the message. Fragment Sequence Numbers begin with fragment number zero and increase by one for each successive fragment in the message.

Priority of CMTS-Assigned DSIDs

Use of Multicast DSID Forwarding (MDF) for multicast IP Video requires the use of processor and/or memory resources on the CM that are identified with a Downstream Service Identifier (DSID). The number of DSIDs (and their accompanying processor and/or memory resources) are limited by the CM's design and the maximum number of DSIDs is communicated from the CM to the CMTS as part of the modem registration process. Per the DOCSIS protocol, the CMTS assumes that each DSID resource is managed and can be assigned by the CMTS at any time of its choosing.

For this reason, the CM may only use DSIDs and associated resources which have not been explicitly assigned by the CMTS. Furthermore, if all such resources are in use (explicitly assigned + Multicast Group Carousel applications) and the CMTS explicitly assigns a new DSID to the CM, the CM MUST immediately cease using one of the DSIDs for Multicast Group Carousel use (thereby losing its subscription to one of the Carousel-defined multicast groups) and immediately assign the DSID and resources as directed by the CMTS.

Keeping the CMTS Apprised of Multicast Group Subscription

The CM MUST inform the CMTS of its use of multicast DSIDs as part of Multicast Group Carousel processing. The CMTS MAY retain this information for accounting or other use but there will no longer be an intent to track all multicast membership in real time. In fact, periodic multicast polling of general and group-specific queries will be disabled to

prevent further CM query response message avalanches at the CMTS. A new unicast MAC Management Messaging mechanism should be created to allow the CMTS to more slowly poll each single CM for its current multicast membership and recent multicast membership changes. The CM would respond in a unicast message with all of its current multicast DSID resource usage and a listing of (S,G) membership changes with an associated timestamp for each. The CMTS might use the returned information to update standard multicast management information as well as to determine when a dynamically added multicast group stream might be pruned from the MD-CM-SG. However, once the CMTS believes that a multicast stream may be pruned, the CMTS should then send a multicast group-specific membership query to make sure that no CM has dynamically JOINed since its last unicast poll.

Adding a New Multicast Group to the Carousel

A multicast group session may be created by any of the following ways:

1. Multicast group protocol membership request (JOIN)
2. CMTS static multicast session CM configuration file encoding (TLV Type 64, see [1] – Section C.1.1.27)
3. Statically-provisioned multicast group sessions
4. Other methods? (PCMM, etc.)

The Multicast group protocol membership request is the most common, and is typically the method used today on IP Video networks. This method uses a multicast protocol such as IGMP (version 2 or version 3) for devices running over IPv4 and the MLD protocol (usually version 2) of IPv6. This processing could proceed exactly as it does with the DOCSIS 3.0 protocol (using DBC messages to install DSIDs in the requesting CM(s)) and then the CMTS could add the group session

information to the Multicast Group Carousel message for any MD-DS-SGs that might be allowed (via ACL provisioning) to carry the session. Care must be taken to handle race conditions between multiple requestors for the same session and the resultant Carousel update.

The static multicast session CM configuration file encoding is used by the MSO to provide the CMTS with the static ASM or SSM multicast sessions to which the CM should be configured to forward multicast traffic at registration time. The CMTS Static Multicast Session Encoding contains the Static Multicast Group Encoding and, if SSM, Static Multicast Source Encoding.

Statically provisioned multicast sessions might be created via the CMTS command line interface (CLI). The CMTS might be responsible for initiating a network-side JOIN using a multicast protocol (PIM SSM or similar) and then the session information might be added to the Multicast Group Carousel message for the assigned MD-DS-SGs.

Backward Compatibility with non-Carousel-Capable CMs

In order for this DOCSIS Multicast Carousel mechanism to work harmoniously with the general population of CMs, the CMTS must understand which CMs must use DBC messaging for each MDF multicast session and which ones can use the Carousel. This requirement implies the need for a CM capability exchange similar to the ones that are used for other DOCSIS features.

A USE CASE

First Client Joins to Initiate Multicast Stream

To illustrate how the DOCSIS Multicast Carousel might be used, different variations of a possible IP Video use case will be

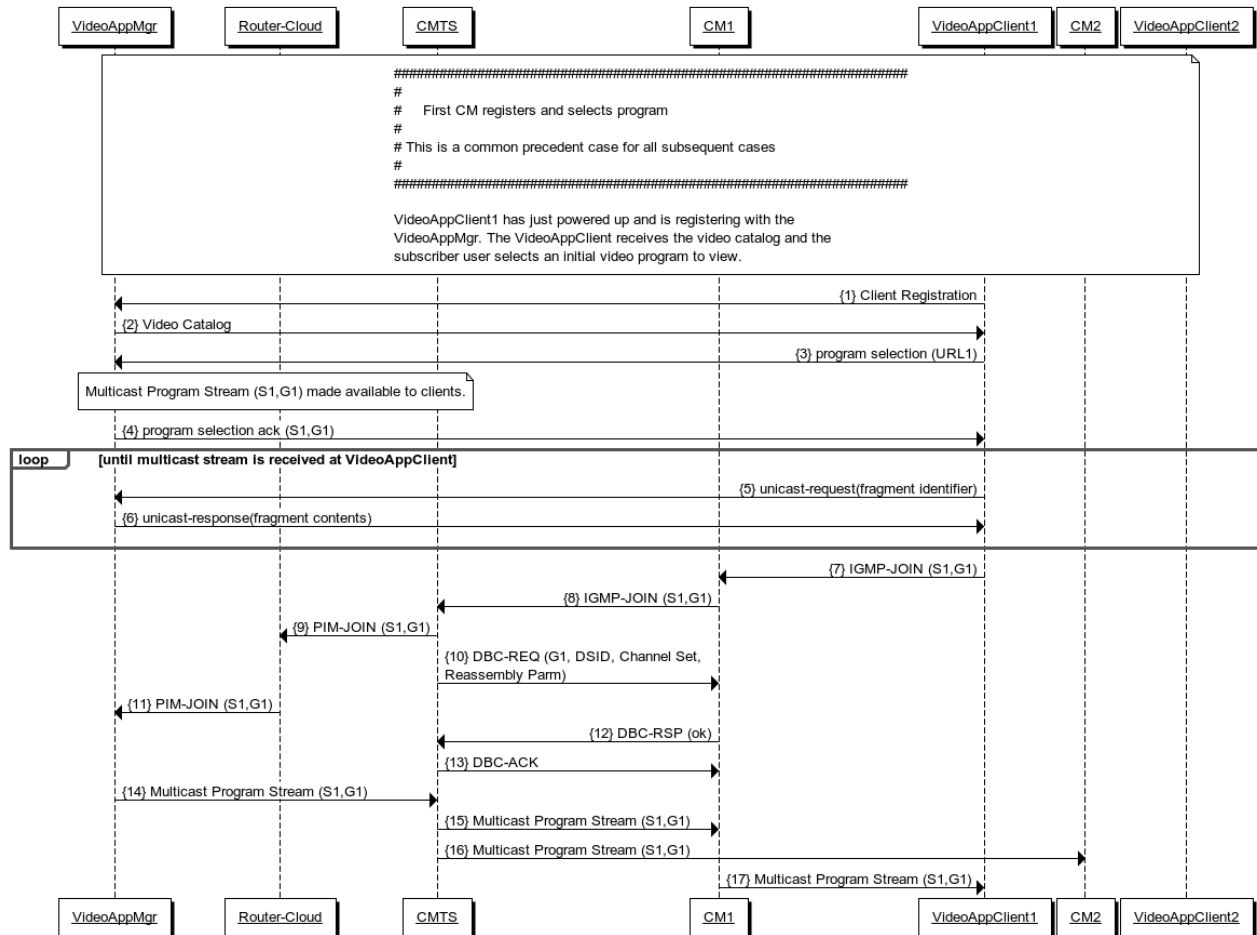


Figure 1: First client joins the IP Multicast group (common to all scenarios)

presented, beginning with the case of the very first client selecting a program which corresponds to a multicast group stream. Then the case of a second viewer of the same program within the same MD-CM-SG will be presented – using only the toolset as published in [1]. Finally, the case of a second viewer of the same program within the same MD-CM-SG will be presented with the use of the DOCSIS Multicast Carousel. This example presents just one of many possible scenarios that involve multiple CMs within the same MD-CM-SG becoming involved in the same IP multicast group session. Similar efficiency benefits can be gained through the use of the DMC in nearly all of the alternate scenarios.

In the following descriptions, each message number in braces {} refers to the numbered message in Figure 1.

A VideoAppMgr (which may include both video control and video media in this example) resides somewhere within an MSO network and a VideoAppClient1 behind a CM powers up and registers with the VideoAppMgr {1}. The VideoAppMgr provides a Video Catalog {2} (perhaps in the form of a browser menu screen) and the user makes a program selection. In this scenario, the VideoAppClient1 communicates this selection via the use of a URL {3}. The VideoAppMgr acknowledges the selection and responds that the video program can be found on a multicast video stream denoted by

network source IP address S1 and multicast group IP address G1 {4}.

As is common in IP Video systems, the VideoAppClient1 uses unicast requests to the server {5, 6} to request video to fill its video jitter buffer before playing. VideoAppClient1 also initiates an IGMP-JOIN request for the multicast stream denoted by (S1, G1) {7}. The VideoAppClient1 will continue to request unicast video fragments until it begins to receive the multicast stream packets.

The CMTS, upon receiving the JOIN request{8}, must request the proper multicast stream from the router cloud using a protocol such as PIM-SSM {9}. The CMTS must also initiate a DBC transaction {10,12,13} to cause the CM1 serving VideoAppClient1 to begin forwarding multicast packets tagged with the DSID on the specified downstream channel set to the VideoAppClient1.

At some point the packets of the Multicast Program Stream identified by (S1, G1) begin to flow through the network, to the CMTS {14} and over the downstream channel set indicated in the DBC-REQ message {15,16}. The CM, having successfully completed its three-way DBC transaction, strips the DSID and forwards the multicast packets to VideoAppClient1 {17}. At this point, VideoAppClient1 ceases its unicast requesting cycle and switches to play the packets received via the multicast stream.

Subsequent Client(s) Join the Multicast Stream

At this point, the packets of the multicast program stream are being tagged with the DSID by the CMTS and are being sent down a downstream channel set that is within CM1's Receive Channel Set. At this point, if other subscribers (single instance using VideoAppClient2 behind CM2 is shown but the example can be repeated for multiple instances of JOINers at the same service

group) wishes to join the same multicast group service flow denoted by (S1, G1), then a comparison can be made between the message flow using the DOCSIS 3.0 protocol as defined in [1] and the message flow as it might appear using the proposed DOCSIS Multicast Carousel.

NOTE: The message flows of the following use case diagrams are intentionally not monotonically sequential. Instead, corresponding messages are numbered the same on both diagrams and unique flows are labeled with either a "-a" or "-b", depending on which case to which they belong.

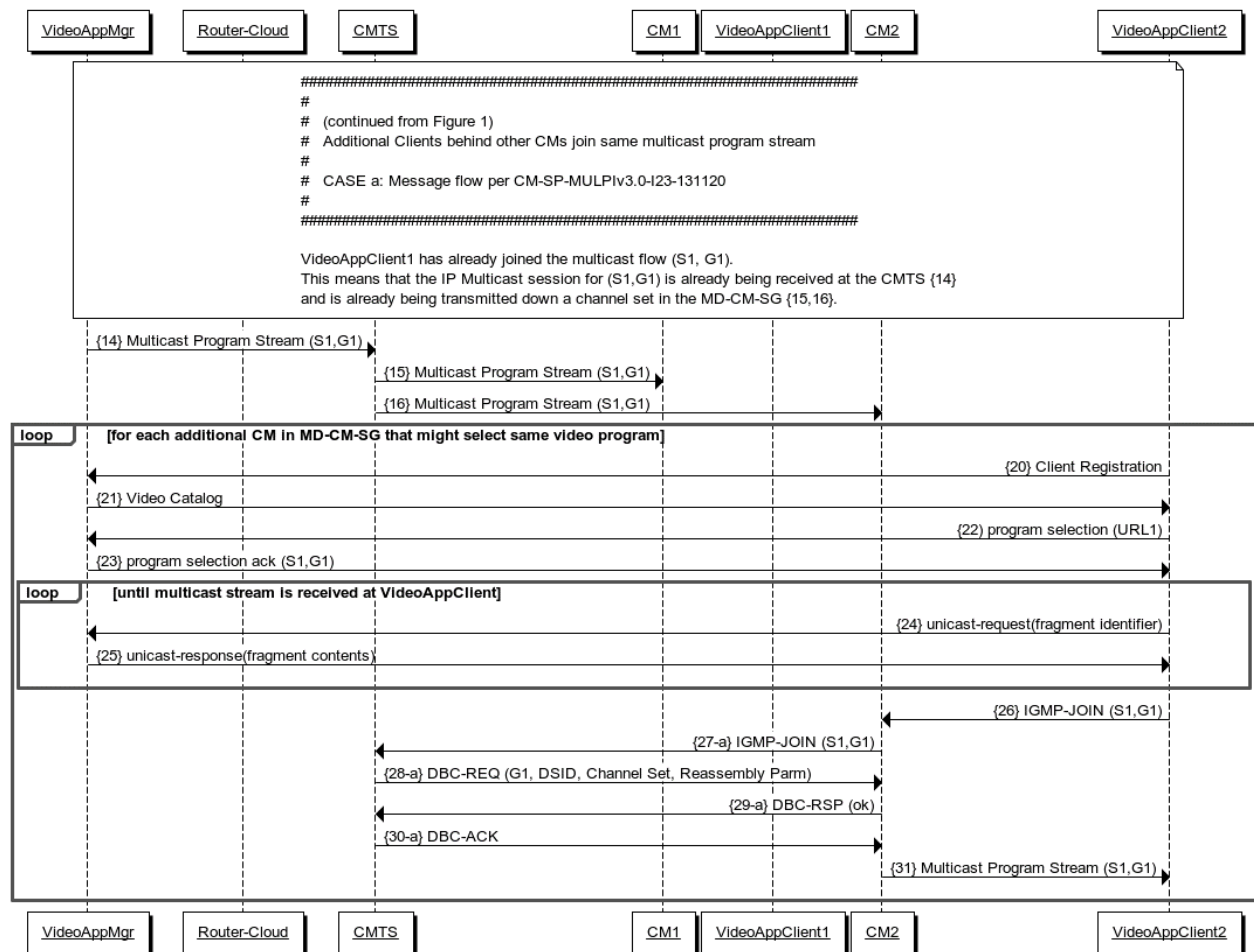


Figure 2: CASE a: Additional JOIN using DOCSIS 3.0 CMTS Control

CASE a: Subsequent Clients using original DOCSIS 3.0 signaling

Using the DOCSIS 3.0 protocol as defined in [1], the message flow would look something like Figure 2. Being a new VideoAppClient, the Client would need to register {20}, receive the video catalog {21}, select a program {22}, and receive a program selection acknowledgement {23}. Also, in order to fill the video jitter buffer, the Client would issue unicast requests {24} and receive fragment contents {25}. Upon receiving the IGMP-JOIN request {26} from the Client, CM2 blindly forwards the request to the CMTS {27-a}. Since the CMTS is already receiving the multicast program stream {14}, it does not need to issue any northbound

requests to the Router-Cloud to get it. The CMTS does still need to initiate the three-way DBC transaction {28-a, 29-a, 30-a} to install the DSID onto CM2. Once the CM installs the DSID, it can begin forwarding the Multicast Program Stream to the VideoAppClient {31}.

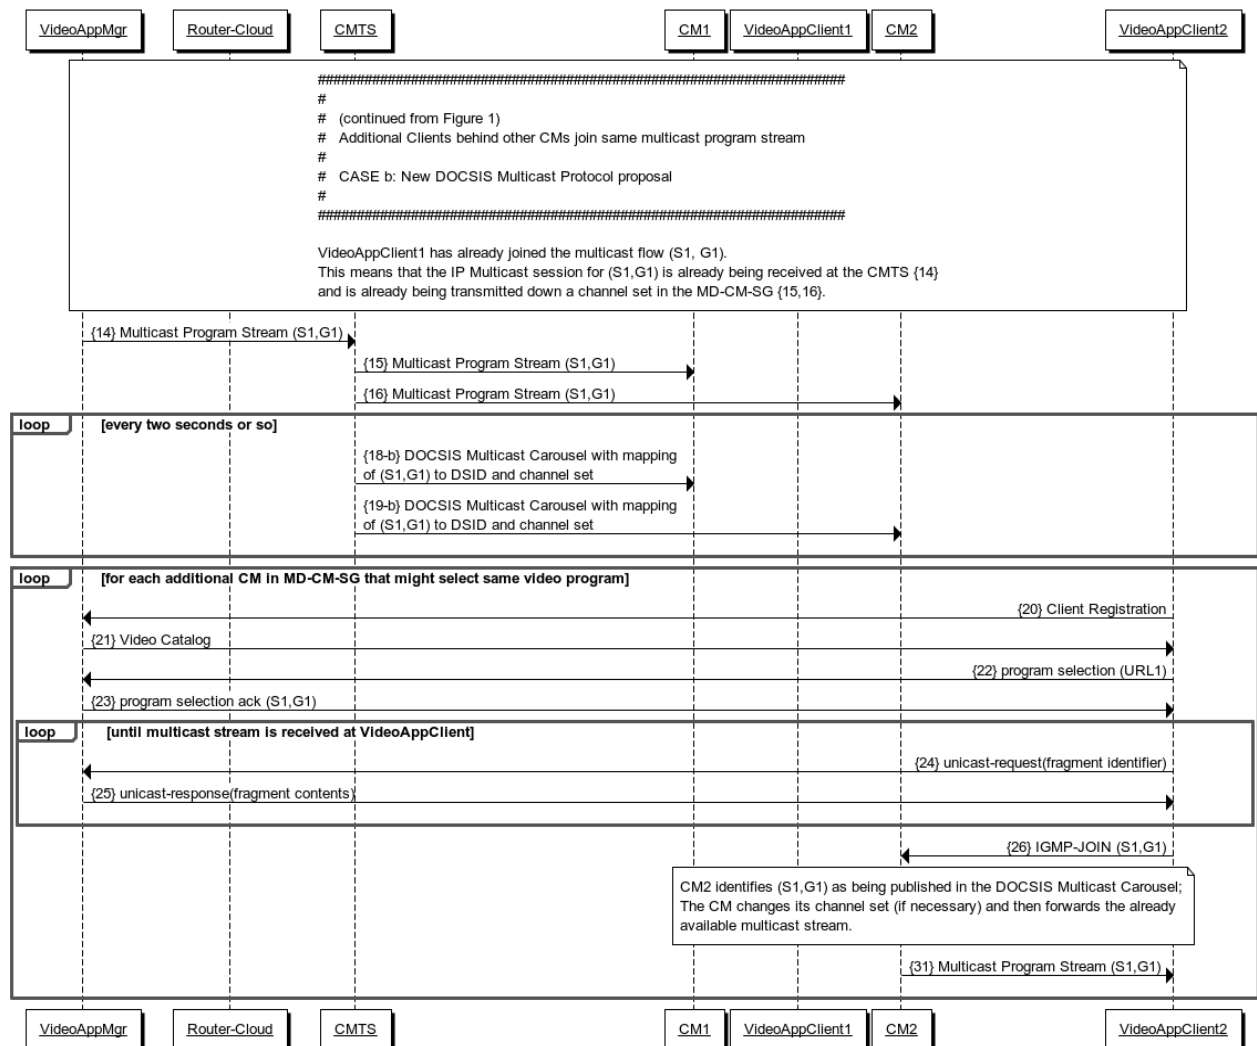


Figure 3: CASE b: Additional JOIN using DOCSIS Multicast Carousel

CASE b: Subsequent Clients using proposed DOCSIS Multicast Carousel

Using the proposed DOCSIS Multicast Carousel protocol, the message flow would look something like Figure 3. With the DOCSIS Multicast Carousel, the CMTS would periodically broadcast the mapping of the multicast stream corresponding to (S1, G1) to the DSID, channel set, and security information within the Multicast Carousel {18-b & 19-b; although this is really only one message that gets to all clients simultaneously}. Once again, being a new VideoAppClient, the Client would need to register {20}, receive the video catalog {21},

select a program {22}, and receive a program selection acknowledgement {23}. Also, in order to fill the video jitter buffer, the Client would issue unicast requests {24} and receive fragment contents {25}.

The interesting operation occurs when the Client sends the IGMP-JOIN request to the CM {26}. Since the CM is already receiving the (DSID, channel set, security info) mapping for the requested multicast group, it can retune one or more pre-designated (perhaps via a new parameter in the DOCSIS RCC encodings?) CM multicast IPTV receivers if necessary and install the required DSID and security information to

immediately begin forwarding the Multicast Program Stream to VideoAppClient2 {31}.

Silent Multicast Group Leaves

With the DMC approach, multicast session LEAVES can be processed in a similar manner as the JOINS without informing the CMTS about the LEAVE. As mentioned previously, a new unicast polling mechanism is being proposed to allow a CMTS-controlled reporting of channel changes. This mechanism can be used to determine when a multicast session might be a candidate to be pruned from the MD-CM-SG. Again, a group-specific query should be issued before the session is actually pruned to catch any CM which have silently JOINed since they were last polled.

Statistics Collection for Subscriber Viewing Activities

In many set-top box environments and/or media gateway environments today, MSOs have the ability to monitor and observe the fine-grained viewing activity of every subscriber, knowing when and how every channel change event was initiated. The use of the DMC approach eliminates (by desire) most of the transmission of multicast JOINS and LEAVES between the CMs and the CMTS, so the CMTS is no longer cognizant of every channel change event. At first glance, this may appear to reduce the level of observability that the MSO may have into the channel change activities of their subscribers. This does not have to be the case.

Slight augmentations to the DMC proposal can restore the fine-grained observability of subscriber channel change activity to the MSO. In particular, the IP Video client application could be required to archive every channel change event, which could then be periodically collected by a Subscriber Management system to provide the MSO with a detailed view of each subscriber's channel

change activities. Alternatively, the CM could be required to archive every channel change event, which could then be periodically collected by a Subscriber Management system or could be periodically collected by the CMTS to provide the MSO with a detailed view of each subscriber's channel change activities.

Benefits of DOCSIS Multicast Carousel

Looking at Figure 1, Figure 2, and Figure 3; the following desirable benefits result:

- Only the first JOINer to new multicast group will require direct CMTS action {8, 9, 10, 12, 13, 15, 16} as part of the channel change; steps {27-a, 28-a, 29-a, and 30-a} can all be avoided for every subsequent JOINer. In addition, even the delay for some of the first JOINers to multicast groups might be avoided if the most popular Multicast Program Streams are statically forwarded in each MD-CM-SG at system startup time.
- Every time that the CM can process the request on its own, significant time and message processing is saved in the channel changing process. In turn, this means that the unicast load imposed by fast-channel-change algorithms on the entire system will be lower than when using a CMTS-controlled multicast model.
- Since the DOCSIS Multicast Carousel is broadcast to all CMs at once, rather than being unicast to each CM like the three-way DBC transaction is, the CMTS can adjust the channel set of one or more multicast streams dynamically and all CMs can retune very quickly if one of the channels experiences a failure.
- Since the DOCSIS Multicast Carousel permits CMs to begin passing any accessible IP Video multicast packets

directly through to the IP Video clients without incurring the delays of the IGMP/PCMM and DBC protocol exchanges, it is feasible that (with appropriate GOP selections) the IP Video multicast streams may be able to be rendered quite rapidly, resulting in relatively Fast Channel Change times without a unicast ABR burst. (Note: Only the first JOINers to an IP Video multicast stream would typically be forced to make use of unicast ABR bursts to expedite their channel change times). If this is the case, then the DOCSIS Multicast Carousel may greatly reduce the need for transient unicast ABR bursts to provide for Fast Channel Changes. As a result, the DOCSIS Multicast Carousel may also help to solve the aforementioned DOCSIS bandwidth problems associated with the aggregated bandwidth requirements of many simultaneously-transmitted unicast ABR bursts.

CONCLUSION

The effects of rapid changes amongst multicast program streams may pose a challenge for many DOCSIS-based IPTV systems in terms of multicast processing power, membership messaging latency, and unicast bandwidth needs for fast-channel-change techniques. This challenge can be overcome by using the DOCSIS Multicast Carousel - a technique borrowed from SDV technology - to transmit the information necessary for a CM to make its own tuning and multicast forwarding decisions, thereby relocating most of the time-critical multicast processing from the CMTS back amongst the distributed processors of the CM population at the very edge of the DOCSIS network.

Author Contacts

Tom Cloonan
Tom.Cloonan@arrisi.com

Amit Eshet
Amit.Eshet@arrisi.com

William T. Hanks
Bill.Hanks@arrisi.com

Jeff Howe
Jeff.Howe@arrisi.com

John Ulm
John.Ulm@arrisi.com

Ian Wheelock
Ian.Wheelock@arrisi.com