



INTERNATIONAL TELECOMMUNICATION UNION

ITU-T

TELECOMMUNICATION
STANDARDIZATION SECTOR
OF ITU

FS-VDSL
White Paper

Full-Service VDSL

Focus Group
White Paper

Channel Change Protocol

Version 1.00
29 November 2002

ITU-T FS-VDSL Focus Group White Paper

CHANNEL CHANGE PROTOCOL

Summary

This White Paper describes how the broadcast channel selection works in an FS-VDSL compliant system architecture.

Source

This White Paper was produced by the CPE-SA Working Group of the ITU-T FS-VDSL Focus Group. Please refer to the FS-VDSL web site at <http://www.fs-vdsl.net> for more information.

This document contains general overview information and should not be construed as a technical specification.

As the FS-VDSL Specifications are revisited, a revised version of this White Paper may be issued.

This White Paper is part of a set of White Papers, published by the ITU-T FS-VDSL Focus Group; for a complete and updated list of published White Papers please refer to the FS-VDSL Focus Group web pages at www.fs-vdsl.net/whitepapers and at <http://www.itu.int/ITU-T/studygroups/com16/fs-vdsl/wps.html>

Acknowledgments

This White Paper is based on contributions of many FS-VDSL Committee and Focus Group members produced during the normal activities of these groups.

However special thanks go to the following companies and individuals who have collected and edited the information here reported.:

Next Level Communications - David Malloy

Telecom Italia Lab - Maurizio Valvo

© ITU 2002

All rights reserved. No part of this publication may be reproduced in any form or by any means, electronic or mechanical, including photocopying and microfilm, without permission in writing from the ITU.

CONTENTS

Page

1. INTRODUCTION.....	1
2. BASICS OF DIGITAL VIDEO DELIVERY OVER THE FS-VDSL NETWORK.....	2
3. CHANNEL CHANGE REQUIREMENTS	3
4. USE OF IGMP VERSUS DSM-CC.....	4
5. APPLICATION OF IGMP IN THE FS-VDSL AN.....	5
6. IGMP OPTIMIZATIONS.....	7
7. CONCLUSIONS	7
8. REFERENCES.....	8
9. GLOSSARY OF TERMS.....	8

FS-VDSL WHITE PAPERS

CHANNEL CHANGE PROTOCOL

1. INTRODUCTION

The FS-VDSL architecture, reported in Figure 1, has been designed to support, besides data and voice services, also broadcast and on demand video services. The way video services are delivered is dictated by the need to provide a user experience similar to that provided by other common means (cable, terrestrial, satellite), on which the video services distribution is based today. The purpose of this white paper is to provide an insight on the way the broadcast channel selection works in an FS-VDSL compliant system architecture.

First, a short introduction is given on the techniques and formats used to transport digital video over fixed networks, then the paper addresses the channel change mechanism requirements, the proposed solutions, the rationales behind them, the pros and cons of the two allowed solutions and finally discusses optimisation and future developments issues.

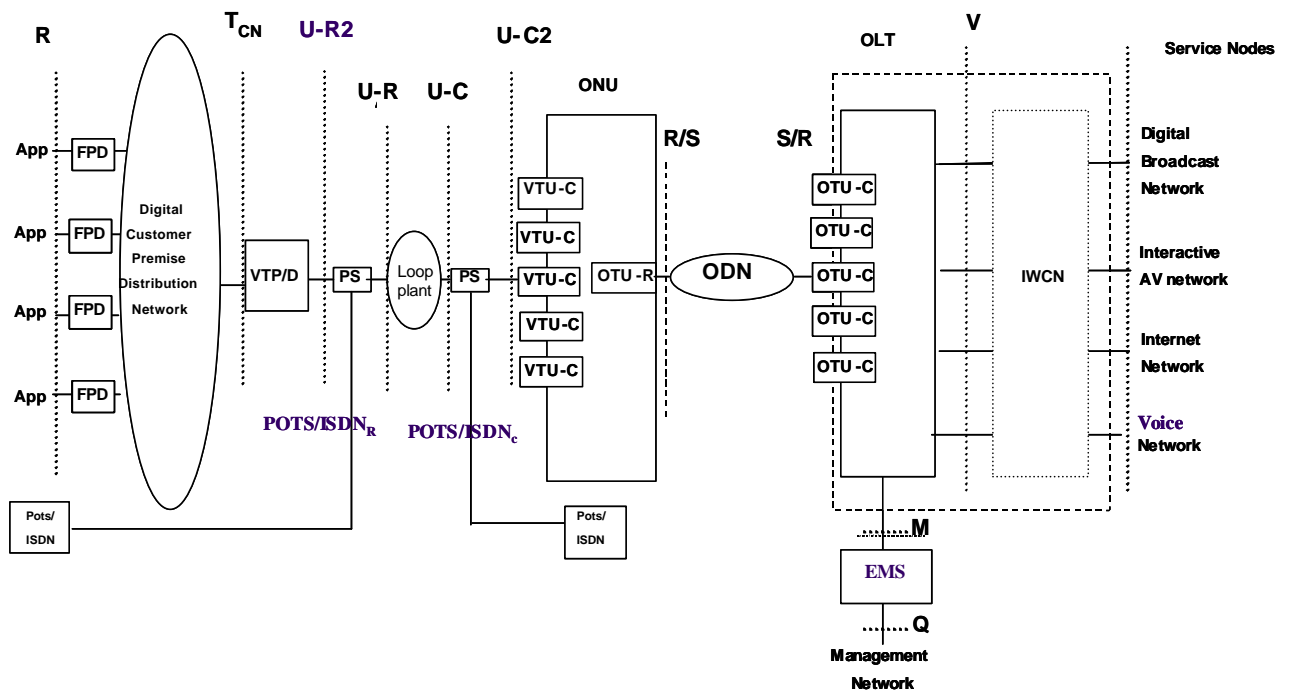


Figure 1: FS-VDSL architecture

2. BASICS OF DIGITAL VIDEO DELIVERY OVER THE FS-VDSL NETWORK

Although the FS-VDSL architecture is transparent to the video coding technique and transport format used, reference is made here to the MPEG-2 coding and format because this is currently the most widely deployed. The MPEG-2 standard, a truly integrated audio-visual standard developed by the International Organization for Standards (ISO), is capable of compressing NTSC or PAL video into an average bit rate of 3 to 6 Mb/s.

Audio and video are coded and packetized into fixed-length MPEG-2 Transport Stream (MTS) packets. Details about the MPEG-2 packets format and the ways to transport synchronization information are outside the scope of this paper.

Although a single MPEG-2 Transport Stream is capable of carrying multiple video and audio channels, the FS-VDSL architecture is based on the assumption that a single channel is transported over one MTS (Single Program Transport Stream – SPTS) at the V reference point, hence there will be as many SPTS as there are channels to be offered. Also it is assumed that all SPTSs are permanently available at this reference point.

Two options for the delivery of digital video over the FS-VDSL network are allowed. Option one is an MPEG over ATM implementation: MTS packets are mapped directly into ATM cells via ATM Adaptation Layer 5 (AAL5) in a process known as “straight eight mapping”. Straight eight mapping gets its name because two MTS packets are segmented into eight ATM cells forming a single AAL5 PDU.

Option two is based on the encapsulation of the MTS in UDP/IP (one or more MTS packets per UDP/IP packet). This second option has the advantage that transport in the core network can occur over any Layer 2 supporting IP packets (e.g. Ethernet, SDH, ATM). The FS-VDSL architecture assumes that the UDP/IP based MTSs are further encapsulated over the Ethernet layer and presented over ATM AAL5 at the V reference point. Ethernet encapsulation simplifies operation of the CPE when MTSs must be delivered to the user’s STB over the Ethernet based home network (distributed home network architecture). Multicast IP and Ethernet addresses are used in this case.

Channels, present at the V reference point, are replicated in the access network (at the OLT or ONU or both) and delivered to a specific drop based on the user’s channel selection commands.

The protocol stack for delivery of DBTV is described in Figure 2 and Figure 3.

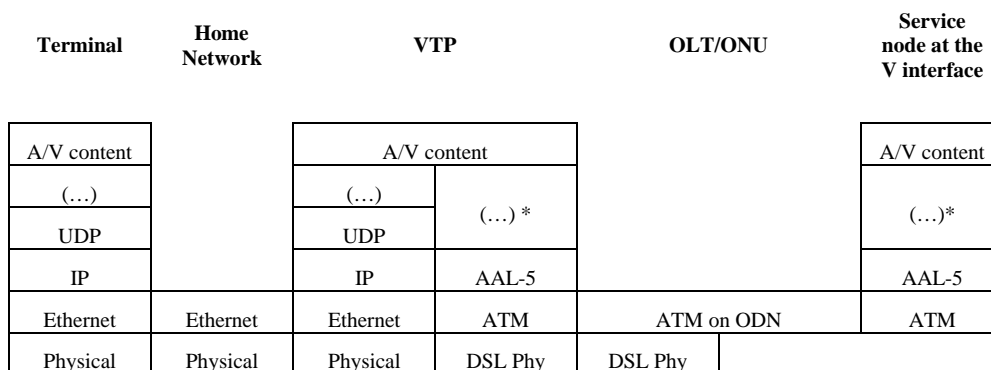


Figure 2: Digital Broadcast Connections Protocol Stack with VTP

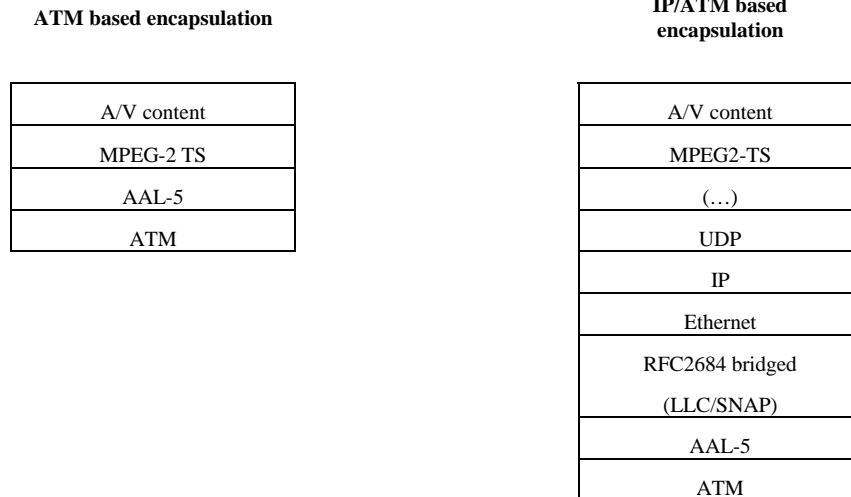


Figure 3: MPEG2-TS/ATM and MPEG2-TS/IP/ATM protocol stacks

3. CHANNEL CHANGE REQUIREMENTS

Channel zapping delays in switched digital broadcast systems result primarily from two sources: the end-to-end associated with protocol processing and connection switching, and the MPEG-2 decoder reframe time.

To mitigate these delays, “human factors” considerations are important. It is important to give quick on-screen feedback when the user presses a button to initiate channel zapping. And, secondly, it is important that the channel zapping delay be as uniform as possible across the channel lineup. That said, the performance target for channel zapping in a switched digital broadcast system should be “as good as possible”. The better the performance, the greater the user satisfaction; the worse the performance, the larger number of unsatisfied and frustrated customers.

In a deployed system, channel surfing represents a varying load on the multicast fabric, and an important performance metric is “sustained performance”. Best-case performance on a lightly loaded system, although interesting, is of little real value.

In addition, resilience to peak load (e.g., “half time at the Superbowl”) needs to be considered in a practical design. There are doomsday and maintenance scenarios that can result in very high peak load, such as neighborhood power outage or a mass reboot. Especially important is the ability for the protocol to tolerate and recover from cell or packet loss during these overload conditions.

4. USE OF IGMP VERSUS DSM-CC

(Note: the reader not familiar with IGMP and DSM-CC protocols is invited to read the Appendix at the end of this White Paper before continuing)

FS-VDSL has specified IGMP v2 as the zapping protocol in the home network.

There has been wide support for the use of the IGMP [1, 2, 3] as the “channel zapping” protocol in the industry. IGMP falls in the “IP family” of protocols, is well known and, in the case of IGMP v2, is readily available. Most Ethernet/IP settops in video over DSL trials use IGMP v2 as the zapping protocol.

Video multicasting may be performed in a multicasting router in the core (in which case, the Access is just a “dumb pipe”). Alternately, multicasting may be performed in the Access Network, and in many cases this approach scales better and saves on transport costs.

FS-VDSL has specified that IGMP v2 is a valid option as the zapping protocol between the VTP and the Network. However, despite its widespread support, there are several robustness and performance issues with IGMP which have been addressed in two ways in FS-VDSL. Firstly, the FS-VDSL committee has specified an IGMP proxy operation in the VTP and alternate timeout values and retry counts than those recommended by IETF in order to attain better performance and more robust operation. Secondly, the committee has allowed DSM-CC [4], a zapping protocol optimised for use in ATM networks, as a valid option at this interface point.

Channel zapping almost always entails “leaving” one stream and “joining” a new stream. In IGMP v2 these two operations require two separate messages. Further, IGMP, since it has no explicit reply, normally requires the host to send each message twice (e.g., “Leave”, “Leave”, “Join”, “Join”). DSM-CC performs this operation with a single message. The result is that normally one fourth of the message traffic is required by DSM-CC to perform the same function. However note that in the home network the additional traffic is inconsequential and, with the optimisations described later, no redundant IGMP messages are sent between the VTP and the AN, relying on other mechanisms to recover from abnormal situations.

The second issue with IGMP has to do with the protocol delays. In IGMP, the “Leave” is separate from the “Join”. To guard against message loss and to verify that there are no other hosts on the home network receiving the same multicast stream, it is necessary for the VTP to perform an “audit” at the time of a “Leave” request. The operation of the audit is described later; with the timeouts recommended by FS-VDSL, this step requires 100 ms to complete. DSM-CC has no such delays, and a well-done implementation can complete the protocol work in 20-30ms. It is recognized, however, that this additional speed is only available if DSM-CC protocol is terminated in the settop (by the centralized model).

Since IGMP has no explicit reply, it is not possible to signal either “success” or, in the case of failure, a “failure code” to the client. With IGMP, if a channel zap fails, for whatever reason, the host gets nothing. With DSM-CC it is possible to indicate the reason for the failure. Importantly, the error code can indicate whether the problem is upstream from the Access System (e.g., channel out of service, no channel defined), or in the Access System (e.g., insufficient bandwidth).

Pros and cons of the two mechanisms are summarized in Table 1.

5. APPLICATION OF IGMP IN THE FS-VDSL AN

Two main options have been studied concerning the functionality required by the different FS-VDSL network elements to support channel change actions in an efficient and robust way. First note that STBs on the home network behave like multicast hosts: they generate unsolicited MRs to request a new channel and LG messages to release a channel; also they respond to MQ messages, generated by the querier, with MR messages. The simplest solution would be to terminate the IGMP v.2 channel change signalling originated by the STB directly in the AN (DSLAM or OLT/ONU); the VTP/D would in this case simply need to filter the IGMP messages and forward them onto a dedicated control connection towards the AN. However, considering that the AN may support even thousands of customers, each of which may have more than one STB, this solution implies an unaffordable processing load onto the AN elements. In fact, a complete state machine must be run for each multicast group (video channel) and for each customer drop, with all associated timers and with stringent reaction times.

A different approach has therefore been followed by including an IGMP “proxy” functionality in the VTP/D. According to this model, the VTP/D assumes the role of a multicast querier with respect to the home network and forwards channel change/leave requests to the AN only if needed. For example, no new request is sent to the AN if the channel requested by a STB is already being received by another STB, or no leave message is forwarded when a STB leaves a channel which is still being received by another STB of the same customer drop. Especially this last situation is very important because it means that the AN does not need to issue group specific query messages on each channel change (when the VTP/D sends a leave message, for sure no other STB is viewing that channel). The AN only needs to react to unsolicited report and to leave messages; it also needs to issue general query messages, but with no stringent time requirements, to recover from possible message losses in the AN to VTP/D segment. The query mechanism may even be started only when abnormal situations occur (e.g. when a new channel is requested and there is no free video connection on the customer drop).

#	Item	IGMP	DSM-CC	Remarks and Explanations
1	Part of the IP suite	+		IGMP is a "living" standard that is still being developed and enhanced by the IETF. IGMP is part of most commercial (real time and non-real time) operating systems. Note that off the shelf IGMP stack should be probably tuned for TV broadcast implementation
2	Number of packets required for every channel switching (i.e. leave + join)		+	DSM-CC needs only one packet in order to switch between channels. IGMP needs at least 2 packets for the same purpose ('Leave' and 'Report' messages). See also item 3.
3	Reliable protocol (Acknowledge messages)		+	A DSM-CC server sends a response for each request it receives. With IGMP, there are no acknowledgements
4	Smooth migration to IP networks	+		An OLT/ONU implementing IGMP can interact with any multicast capable IP platform by enhancing its functionality from IGMP server to IGMP proxy. An OLT/ONU Implementing DSM-CC should, for this purpose, implement an IGMP client front –end towards the network and a translation function.
5	Supports querying	+		The query functionality is important in 2 aspects: 1. It enables discovering disconnected NTs/STBs. This is important in scenarios where the ONU/OLT requests broadcast channels from the network upon demand. In that case a discovery of disconnected NT/STB can cause release of unnecessary CBR bandwidth for the favor of UBR applications. 2. It enables a rebooted ONU/OLT to discover the channels required by the active STBs without user intervention. IGMP servers should be designed to absorb IGMP report bursts generated as a result of queries
6	Simple protocol stack at the NT	+		If an NT has to support DSM-CC on the network side and IGMP on the customer premises side, then both DSM-CC and IGMP state machines are required at the NT. A translation function between the protocols is also required.
7	Supports zapping status codes		+	DSM-CC messages carry various status codes. One use of this feature is, for example, giving the user an indication why her zapping request failed (e.g. channel not available, user not authorized, etc.). Note that if the DSM-CC protocol terminates at the NT, there is no standard way to forward such information back to the STB. It is important to forward the status information back to the STB in order to be able to display an appropriate message on the TV screen and let the user know why the system is unable to provide the requested service.
8	Enables network side zapping		+	A DSM-CC server can instruct a user to switch to a specified channel. One application of this feature is to move the user to/from a PPV channel when the purchased viewing time starts/ends. Note that if the DSM-CC protocol terminates at the NT, there is no standard way to forward such an instruction back to the STB.

Table 1: IGMP/DSM-CC comparison

6. IGMP OPTIMIZATIONS

In the FS-VDSL Specifications, several IGMP parameters, related to the STB and VTP/D querier operation, have been modified, compared to the RFC2236 defaults, to greatly reduce the zapping time. First, to reduce the time to join a group in case the first unsolicited report message is lost, the Unsolicited MR Interval is set to 100 ms, whereas the RFC2236 default is 10 s. But a more effective optimization is achieved through the reduction of the Last MQ Interval and Last MQ Count parameters, which are set respectively to 100 ms and 1, whereas the RFC2236 defaults are 1 s and 2. This optimization reduces the time to drop a group and hence the time needed to send a new channel when the customer drop bandwidth is fully utilized.

All other parameters are supposed to take their default values as indicated in RFC2236. However other IGMP parameters optimisations are possible. For example, the general MQ mechanism parameters can be changed, both in the VTP/D and in the AN, to allow for a faster recovery from abnormal situations originated by a loss of messages.

In the most general service scenario, the AN, the VTP/D and the STB may be managed by different operators. Hence the assignment of IP addresses to be used in IGMP messages is an issue. However note that IGMP destination addresses are only related to the multicast groups being joined or left, which in turn represent broadcast video channels; these multicast group addresses must always be used consistently by all network elements.

The situation is different concerning IP source addresses used in IGMP messages. In fact, considering that user requests can be identified on the basis of the drops they are received from, at least as far as the AN operation is concerned, and that IP source addresses are not relevant to the IGMP operation, there is no need to assign these addresses in a consistent way on all network segments, i.e. they can be assigned independently from each other.

This is summarized by saying that the STB, the VTP/D, and the AN may use any source IP address in transmitted IGMP messages, provided that they do not care about the source IP address found in received IGMP messages.

7. CONCLUSIONS

This paper has described the channel change mechanisms adopted in the FS-VDSL Specifications, illustrating the rationale behind their selection and pros and cons of the two allowed solutions.

8. REFERENCES

- [1] RFC 1112 – *Host Extensions for IP Multicasting*. © 1989, The Internet Society
- [2] RFC 2236 – Internet Group Management Protocol, Version 2. © 1997, The Internet Society
- [3] <draft-ietf-idmr-igmp-v3-11.txt> - Internet Group Management Protocol, Version 3. October 2002
- [4] ISO/IEC 13818-6 – Information Technology – Generic Coding of moving pictures and associated audio information – Part 6: Extensions for DSM-CC. © ISO/IEC. 1998. See Chapter 10, U-N Switched Digital Broadcast – Channel Change Protocol, and Annex H (informative), Switched Digital Broadcast Service. Note that within this standard, 13818-1 is “Systems”, 13818-2 is “Video”, 13818-3 is “Audio”, ... 13818-6 is “Extensions for DSM-CC”

9. GLOSSARY OF TERMS

A glossary of terms is available on the FS-VDSL Focus Group White Papers web pages at www.fs-vdsl.net/whitepapers and at <http://www.itu.int/ITU-T/studygroups/com16/fs-vdsl/wps.html>.

APPENDIX 1

Basics of IGMP v.2

IGMP is used by IP hosts to report their multicast group memberships to any immediately-neighboring multicast routers. IGMP v.2 is specified in IETF RFC2236; a short summary is given here in which IGMP v.1 backward compatibility is not addressed and only with reference to the operation of a querier router.

IGMP messages are encapsulated in IP datagrams and have the format shown in Figure 4. The basic IGMP message types are: Membership Query (MQ), Membership Report (MR) and Leave Group (LG). Membership Queries can be generic (General MQ) or specific (Group Specific MQ).




Figure 4: IGMP message format

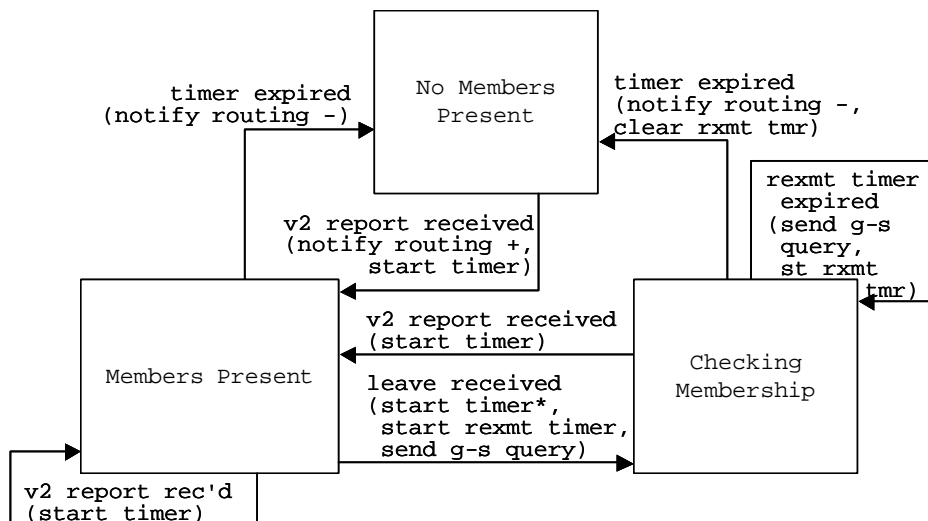


Figure 5: Simplified IGMP querier router state machine diagram

A multicast router keeps a list of multicast group memberships and a timer for each membership; querier routers periodically send a General MQ, to solicit membership information. Hosts respond to this General MQ to report their membership status for each multicast group. Group specific MQs may also be sent, for example when a router

needs to check whether there are more members of a group for which a LG message has been received. If no MRs are received for a certain multicast group during a predefined period of time, the router assumes that there are no more members and stops forwarding traffic for that group. The querier router state machine diagram, simplified for the IGMP v.2 only case, is shown in Figure 5.

When a host joins a multicast group, it sends an unsolicited MR for that group. To cover the possibility of the initial MR being lost or damaged, RFC2236 recommends that it be repeated once or twice after short delays [Unsolicited Report Interval]. When a host leaves a multicast group, it sends a LG message.

When a Querier receives a LG message, it sends Group-Specific MQs to the group being left. If no MRs are received in response to these MQs, the router assumes that there are no more members and stops forwarding traffic for that group.

APPENDIX 2

Basics of IGMP v.3

The draft IGMP v3¹ proposes some enhancements to IGMP v2. In particular, it adds “Source Filtering”, which supports an atomic leave/join. This helps recovery from lost “leave” messages and simplifies some of the error recovery scenarios. This capability also halves the message traffic, as now only a single message is needed to “leave” one stream and “join” another.

An IGMP v2 PDU fits in a single ATM cell (20 byte IP header, 8 bytes of IGMP). The source filtering capability of the IGMP requires at least 52 bytes, and, in turn, two ATM cells (20 byte IP header, 8 bytes of IGMP header, 12 bytes for the first group record, 12 bytes for the second group record). This doubles both the ATM bandwidth and buffer memory required to process channel zapping messages. In turn, this requires a more complicated ATM SAR, and increases the amount of software processing in the multicast signaling termination function within the Access Network.

IGMP v3 has no confirming reply. Like IGMP v2, messages must be sent twice to guard against loss, and there is no mechanism to return an error code to the client.

IGMP v3 is backward compatible to IGMP v2, so IGMP v3 ready terminals will correctly interoperate with the FS-VDSL AN where IGMP v2 is used.

IGMP v3 is only in draft form and is not widely available.

¹ <draft-ietf-idmr-igmp-v3-11.txt> - *Internet Group Management Protocol, Version 3*. October 2002

APPENDIX 3

Basics of DSM-CC

The DSM-CC Switched Digital Broadcast (SDB) Protocol is defined in Chapter 10, Part 6 of ISO/IEC 13818 (e.g., the MPEG standard). This protocol was designed specifically for efficient channel surfing in a switched digital environment, and although it could be used in any digital transport network, it was optimized for use with ATM transport networks.

The protocol is a simple message / reply protocol. Both client initiated and network initiated zapping are supported. Four messages are defined: SDBProgramSelectRequest and SDBProgramSelectResponse for client initiated zapping, and SDBProgramSelectIndication and SDBProgramSelectConfirm for network initiated zapping.

All messages are less than 48 bytes in length and therefore fit into a single (AAL5) ATM cell.

An example of a Client initiated channel zapping request is given below.

```
msgId           = 0x0001           (SdbProgramSelectRequest)
transactionId   = 0x00112233
sessionId       = 0x0000 0000 0000 0000 0001
broadcastProgramId = 0x0000 0009
```

```
msgId           = 0x0002           (SdbProgramSelectConfirm)
transactionId   = 0x00112233
sessionId       = 0x0000 0000 0000 0000 0001
broadcastProgramId = 0x0000 0009
response        = 0x0000 (OK)
vpi             = 0x0001
vci             = 0x0065
```

messageId: SdbProgramSelectRequest

SdbProgramSelectResponse

SdbProgramSelectIndication

SdbProgramSelectConfirm

transactionId: used to correlate replies to requests

sessionId: indicates which "stream" is being "zapped". Each session can receive only one channel at a time

broadcastProgramId: arbitrary identifier to identify the stream. For FS-VDSL, this should be the IP Multicast id

response: 0 - success; non-zero values indicate errors

vpi/vci: the ATM address of the stream.