

IPDVB Working Group
Internet-Draft
Expires: June 9, 2007

J. Cantillo
ENSICA/TESA/Alcatel Alenia Space
J. Lacan
ENSICA/LAAS-CNRS
December 6, 2006

draft-cantillo-ipdvv-s2encaps-04
A Design Rationale for Providing IP Services Over DVB-S2 Links

Status of this Memo

By submitting this Internet-Draft, each author represents that any applicable patent or other IPR claims of which he or she is aware have been or will be disclosed, and any of which he or she becomes aware will be disclosed, in accordance with [Section 6 of BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/1id-abstracts.txt>.

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

This Internet-Draft will expire on June 9, 2007.

Copyright Notice

Copyright (C) The Internet Society (2006).

Abstract

This document describes a framework for the transmission of IP datagrams over DVB-S2, the second generation standard for Digital Video Broadcasting over Satellite. The new standard features an improved and adaptive physical layer, as well as a new framing structure at link level, the Generic Streams. Combined use of adaptability and Generic Streams is expected to offer throughputs never achieved for IP services up to now, but no standard way to

carry IP data using the specific features of DVB-S2 has been published up to date. The present document analyzes these issues, and it identifies the requirements for the definition of a standard interface between the DVB-S2 link layer and an IP subnetwork.

Table of Contents

1.	Introduction	3
2.	Conventions Used in This Document	4
3.	DVB-S2 Architecture	6
3.1.	Functional Blocks and Framing Overview	6
3.2.	BBHEADER Fields	8
4.	Providing IP services over DVB-S2	9
4.1.	Basic Architecture Considerations	10
4.1.1.	Protocols Supported	10
4.1.2.	Encapsulation	10
4.1.3.	DVB-S2 Network Topologies	10
4.2.	DVB-S2 Logical Channels and ISI	11
4.3.	Address Resolution Issues	11
4.4.	QoS Mapping and MODCOD Selection	12
4.4.1.	Cross-Layer Optimizations for QoS Scheduling Policies	12
4.4.2.	Differentiated QoS over Logical Channels	13
5.	Motivations for a New Approach	13
5.1.	ACM and DVB-S2 Framing	13
5.2.	IP over Generic Streams	14
6.	General Framework Requirements	14
6.1.	Use of Existing Encapsulation Capabilities	15
6.2.	Fragmentation Issues. Adaptive Packing and Scheduling Optimization	15
6.2.1.	Fragmentation Schemes to be Supported	16
6.2.2.	Packing Optimization	17
6.3.	Next Level Protocol Type	18
6.4.	Precise Payload Unit Delineation and Reassembly	19
6.5.	Integrity Check	19
6.6.	Link Layer Addressing Capabilities	20
6.7.	Flexibility for Future Extension	20
6.8.	Security Considerations	21
7.	Summary	21
8.	Acknowledgements	21
9.	Document History	22
10.	References	22
10.1.	Normative References	22
10.2.	Informative References	23
	Authors' Addresses	24
	Intellectual Property and Copyright Statements	25

1. Introduction

This document defines a framework that may be used to send/receive IP datagrams and packets of other network protocols using DVB-S2, the new physical layer and framing specification for satellite links defined by the Digital Video Broadcasting project [1]. Recently approved by the European Telecommunications Standards Institute (ETSI), this architecture is designed for broadband satellite applications such as digital television or radio, as well as interactive services such as Internet access or content distribution.

Compared to its predecessor DVB-S [2], DVB-S2 features two main enhancements. First, higher order modulations and stronger FEC are teamed up with return channels to achieve real-time adaptability to the link and propagation conditions. This novelty called Adaptive Coding and Modulation (ACM) might allow for an enhanced throughput by 30% to 150%, which explains in part the genuine interest for this new standard. DVB-S2 supports also Variable Coding and Modulation and of course, Constant Coding and Modulation (CCM) modes.

Second, DVB-S2 offers a new link-layer framing structure called Generic Streams (GS) in addition to the classical packetized Transport Streams (TS). GS can be packetized or continuous, and are intended to address the non-native way in which IP services are carried today over MPEG2-TS using the Multi Protocol Encapsulation (MPE) [5] or the Unidirectional Lightweight Encapsulation (ULE) [4]. Indeed, MPEG2-TS constraints such as constant bit-rate and constant end-to-end delay are not a must for IP services, which added to the accumulation of multiple overheads undermine IP carriage efficiency. Since the use of TS is no longer mandatory in DVB-S2 for carrying IP data, IP over GS seems to offer new possibilities for satellite-based IP networks.

Up to now, no standard procedure to carry IP datagrams over Generic Streams has been published yet. The proposal named Generic Stream Encapsulation (GSE) however, is the most likely to achieve convergence among the different solutions proposed. In order to take advantage of the new facilities provided by DVB-S2, the previously mentioned solutions to encapsulate IP over DVB-S should be adapted or new solutions should be proposed. The key goals are to reduce complexity when using the system while improving performance, increasing flexibility for IP services and providing opportunities for better integration of IP-based networks.

After a detailed introduction to DVB-S2 architecture, [Section 4](#) describes how a DVB-S2 link can be used to provide the Internet service. Finally, guidelines for the definition of a standard interface between the new link layer and an IP subnetwork are

exposed. When defined, the proposed interface should minimize overhead and be simple enough to reduce processing while ensuring adequate network services, as well as be robust to errors and security threats while providing enough flexibility for future extension, as exposed in [RFC 3819](#) [6].

2. Conventions Used in This Document

ACM: Adaptive Coding and Modulation. Dynamic adjustment of transmission parameters (FEC coding rate and modulation scheme) on a BBFRAME-by-BBFRAME and Receiver-per-Receiver basis, according to link and propagation conditions. In order to implement ACM, feedback from each Receiver has to be provided by a return channel such as DVB-RCS [8].

b: bit. For example, one byte consists of 8b.

B: Byte. Groups of bytes are represented in Internet byte order.

BBFRAME: Main framing unit of the DVB-S2 protocol stack, consisting of a 10B BBHEADER, a variable size DATAFIELD and Padding (when necessary). It may carry either Generic Streams (GS) or Transport Streams (TS).

BBHEADER: 10B header of a BBFRAME.

CCM: Constant Coding and Modulation. In CCM transmission mode, the system uses a single type of pre-defined waveform for every Receiver. DVB-S is a CCM system.

DATAFIELD: Payload transported by each BBFRAME. Its maximum size determines the length of the BBFRAME that contains it. For a given Receiver, this maximum size is allowed dynamically on a BBFRAME-by-BBFRAME basis among 21 possible sizes (ranging from 372B to 7264B) by the VCM/ACM command. Each size is related to the FEC coding rate to be used for encoding each particular BBFRAME. Lower BBFRAME sizes are used when low coding rates (i.e. stronger protection against channel errors) are needed, whereas longer sizes (i.e. higher coding rates) are used under good channel conditions.

DFL: One of the BBHEADER fields. Length in bits of the DATAFIELD.

DVB-S2: Second Generation of the Digital Video Broadcast for Satellite applications standard [1]. A framework and set of associated standards published by the European Telecommunications standards Institute (ETSI) for the transmission of video, audio, and data, intended to replace DVB-S [2].

FEC: Forward Error Correction. The scheme used in DVB-S2 relies upon the concatenation of Bose-Chaudhuri-Hocquenghem (BCH) and Low Density Parity Check (LDPC) codes. For a given Receiver, its overall coding rate is adjusted on a BBFRAME-by-BBFRAME basis according to the needs specified for each BBFRAME by the VCM/ACM command.

FECFRAME: Encoded BBFRAME, as seen at the output of the FEC encoder. FECFRAMEs have only two possible sizes, 2025B and 8100B, no matter the size of the BBFRAME to code.

Generic Stream: One of the 2 possible input streams in DVB-S2, the other one being Transport Streams. Generic Streams can be packetized or continuous, and are intended to be used especially for carrying data services such as IP distribution. An example of packetized Generic Stream could be a flow of MPEG2 packets.

GS: See Generic Stream.

ISI: Input Stream Identifier, second byte of the BBHEADER field for multiple input streams. It provides a way to separate different BBFRAMEs within a single multiplex, defining logical channels for BBFRAMEs.

MPEG2: A set of standards specified by the Motion Picture Experts Group (MPEG), and standardized by the International Standards Organization (ISO) [3].

MODCOD: Information provided by the VCM/ACM command to the BBFRAME assembler, specifying the coding rate (therefore the size of the maximum size of DATAFIELD to be allowed) and the modulation scheme to be used for encoding and modulating each BBFRAME.

NPA: Network Point of Attachment. Addresses primarily used for station (Receiver) identification within a local network (e.g. IEEE MAC address). An address may identify individual Receivers or groups of Receivers.

PID: Packet Identifier [3]. A 13 bit field carried in the header of TS Packets. This is used to identify the TS Logical Channel to which a TS Packet belongs [3]. The TS Packets forming the parts of a Table Section, PES, or other Payload Unit must all carry the same PID value. The all 1s PID value indicates a Null TS Packet introduced to maintain a constant bit rate of a TS Multiplex. There is no required relationship between the PID values used for TS Logical Channels transmitted using different TS Multiplexes.

PDU: Protocol Data Unit. Examples of a PDU include Ethernet frames, IPv4 or IPv6 datagrams, and other network packets.

PLFRAME: Last framing unit of the Physical Layer of DVB-S2.

PLHEADER: Header of a PLFRAME. The PLHEADER contains synchronization information coded over 2b, as well as the MODCOD used for the current frame (5b). Since it has to be demodulated and decoded for every Receiver without a priori knowledge of the carried MODCOD, it features an unique modulation and coding, no matter the payload's MODCOD.

Receiver: An equipment that processes the signal from the satellite and performs filtering and forwarding of encapsulated PDUs to the network-layer service (or bridging module when operating at the link layer).

Return Channel: A way for end-users to interact with the transmitting Gateway, in order to establish a bi-directional communication. Such technologies can make use of two-way satellite links [8] and/or terrestrial links.

SAR: Segmentation And Reassembly, generally for encapsulated SNDUs.

SNDU: Sub Network Data Unit. A framing entity created on the basis of a network-layer PDU by an adaptation layer protocol, allowing this PDU to travel along a L2 subnetwork.

TS: Transport Stream [3], a method of transmission at the MPEG2 level using MPEG2 Packets.

UPL: One of the BBHEADER fields. It carries packets lengths in bits, in the case of packetized input streams.

VCN: Variable Coding and Modulation. Differentiated optimization of transmission parameters according to the physical characteristics of every Receiver, such as geographical position, size etc.

[3. DVB-S2 Architecture](#)

[3.1. Functional Blocks and Framing Overview](#)

DVB-S2 is organized as a sequence of functional blocks.

The Mode Adaptation block processes input data structured either as TS or GS, single or multiple. Input streams are sliced into DATAFIELDS to which a 10B BBHEADER is appended. The maximum length of every DATAFIELD is chosen dynamically among 21 allowed sizes in the range [374B; 7264B] by the VCN/ACM command, according to the protection required for everyone of them. Basically the shorter they

are, the more space there is for FEC redundancy protection. Actual systems may only implement a subset of those 21 sizes.

The Stream Adaptation block may complete every DATAFIELD with Padding in order to match the length of a valid BBFRAME. BBFRAMEs therefore have one of 21 possible pre-defined sizes in the range [384B; 7274B]. This is a major difference with DVB-S, since at this level there are only 188 Byte MPEG2 containers. Note that DATAFIELDs sizes are not multiples of 188B: Transport Streams, as well as Generic Streams, are mapped asynchronously over BBFRAMEs.

Adaptive FEC encoding constitutes the third block. A set of coding schemes based on a concatenation of LDPC and BCH codes ensures a very efficient error protection, only 0.7 to 1 dB away from the Shannon limit (DVB-S FEC is around 2.5 dB from that margin). In ACM mode, the ACM command dictates dynamically the coding rate to be used for every BBFRAME in order to provide a Quasi-Error Free (QEF) quality target at the input of the Receiver's DEMUX (roughly equivalent to $PER < 1e-7$ for a MPEG2 transmission). FEC parity bits are calculated and appended systematically to each BBFRAME in order to provide fixed-length FECFRAMEs of 2025B or 8100B. This implies that shorter BBFRAMEs are completed with more redundancy than long BBFRAMEs, and are therefore more protected.

Finally, FECFRAME bits are modulated and raised-cosine filtered, to provide the body of a PLFRAME. 4 different modulations with spectral efficiencies ranging from 2bits/s/Hz to 5 bits/s/Hz are available in DVB-S2. Finally, information about the FEC coding rate and modulation used for every frame (MODCOD) is stored in a PLHEADER and appended to every PLFRAME. Of course, DVB-S2 provides mechanisms to ensure proper reading of every PLHEADER for every Receiver without a priori knowledge of the contained MODCOD, so all PLHEADERS use a pre-determined coding and modulation. The final PLFRAME is finally sent over the RF carrier using classical TDM and/or FDM techniques.




Figure 1: DVB-S2 architecture and framing structure overview

3.2. BBHEADER Fields

Several statements made in the following sections will refer to fields present in the 10B BBHEADER, so a very short description of this entity is presented here:



Figure 2: A BBHEADER

The first byte of the MATYPE field specifies whether TS or GS are used, and whether they are single or multiple. In the multiple case, the second byte is an "Input Stream Identifier" (ISI).

UPL specifies packet lengths in bits, in the case of packetized input streams. As an example, a value of 0x05E0 (188*8 hexadecimal) is characteristic of MPEG2 packets. A value of 0x0000 means a continuous GS.

DFL specifies the length of the DATAFIELD actually used in bits, in the range [0b; 58112b] (58112 = 7264*8, 7264B being the maximum DATAFIELD length allowed).

SYNC stores the synchronization byte carried by all the packets of a packetized stream, when there is one (e.g. if MPEG2 packets are transported, SYNC=0x47). Since the synchronization byte is carried by BBHEADERS, there is no need for every packet to carry it anymore. A CRC-8 calculated for every packet replaces therefore the synchronization byte in every packet : it is used to validate Segmentation And Reassembly (SAR) applied on them. SYNC is not relevant for continuous Generic Streams.

SYNCD is the distance in bits, for a packetized stream, from the beginning of the DATAFIELD to the first start of packet contained in this DATAFIELD. Its use is therefore similar to a Payload Pointer, as defined in ULE [4]. SYNCD is not relevant for continuous Generic Streams.

Finally, a CRC-8 is calculated from the previous 9B of the BBHEADER.

Note that BBHEADER fields natively support SAR applied to MPEG2 packets or any other fixed-length packets asynchronously mapped over a BBFRAME flow. Indeed, perfect delineation and reassembly can be achieved by the exclusive use of UPL, DFL and SYNCD for packetized Generic Streams. Finally, the CRC-8 stored in the 1B slot liberated by SYNC in every packet provides an end-to-end integrity check, achieving thus an encapsulation that does not produce any overhead at all (except when Padding is necessary). In DVB-S2, a flow of MPEG2 packets can therefore be sent over a packetized Generic Stream using UPL=0x05E0 and SYNC=0x47.

[4.](#) Providing IP services over DVB-S2

The purpose of this section is to describe how a DVB-S2 link can be used to deliver the Internet service. DVB-S2 not only provides all the tools necessary for a reliable and efficient transmission/reception of network layer datagrams, but it also allows to improve the way their delivery is classically done over satellite links.

4.1. Basic Architecture Considerations

4.1.1. Protocols Supported

This document focuses mainly on the transmission and reception of IPv4 and IPv6 unicast, multicast and broadcast/anycast datagrams. This includes datagrams with compressed or encrypted IPv4/IPv6 headers (e.g. [RFC 1144](#) [9], [RFC 2507](#) [10], [RFC 3095](#) [11] and [RFC 2401](#) [12]). However, the scope of the proposed framework is general enough to support a wide range of other network layer protocols, such as the ones recorded in the IANA EtherType registry.

4.1.2. Encapsulation

PDUs arrive to the transmission Gateway using the existing infrastructure, as e.g. IP datagrams, Ethernet frames etc. They are then shaped by an Encapsulator into SNDUs, fragmented when necessary and packed into BBFRAMES, either directly (Generic Streams) or via an additional layer (MPEG2-TS). BBFRAMES are coded, modulated and sent through the RF channel to the satellite.

Encapsulation allows PDUs to travel along an L2 subnetwork, and provides mapping of network-level functionalities over L2 entities.

4.1.3. DVB-S2 Network Topologies

As a successor to DVB-S, some compatibility at least, as well as enhancement of network capabilities are expected for DVB-S2. Transmission networks to be supported by DVB-S2 contain therefore basically those specified in [RFC 4259](#) [7]. Those are:

Uni-directionnal scenarios such as:

- Digital radio and TV broadcast
- Broadcast networks used by an ISP
- Uni-directionnal star IP scenarios
- Datacast overlay

Bi-directionnal scenarios such as:

- Point-to-Point links
- Two-way IP networks

The growing demand for IP services and the increased availability of return channels are bound to play an important role in the development of the last 2 scenarios. In addition, the enhanced throughput capabilities of the new standard will most likely influence positively this development.

4.2. DVB-S2 Logical Channels and ISI

In the same way the PID field allowed to distinguish TS logical Channels for MPEG2, the ISI byte of every BBHEADER (see comment on the MATYPE field in [Section 3.2](#)) can be used to support logical channels for BBFRAMES. This would provide a powerful discriminating tool within a single multiplex of BBFRAMES, that could be used e.g. for efficient address resolution, QoS differentiation or to provide other services. However, up to now there is no standardized signalling associated to ISI (such as tables or reserved values), and further work has to be done in this direction in order to set a complete framework. This important point includes precise mapping of upper layer QoS functions, or standards to associate the capabilities of these potential logical channels to upper layer flows.

Note that other methods could be imagined to set up logical channels in DVB-S2. However, the use of ISI, although not described in detail in the DVB-S2 specification, seems naturally adapted to this particular task.

If MPEG2 packets are to be mapped into BBFRAMES, there are no indications either in DVB-S2 about the way PIDs and ISI have to be related to each other. Since PID already defines logical channels at MPEG2 level, the simplest solution would be to keep the ISI field unused, and filter at MPEG2 level. In another figure case, ISI could be used to define meta-channels of PIDs. The last possibility is ISI being just a copy of the transported PIDs, which supposes that only MPEG2 packets sharing the same PID are allowed to travel together. Note that even though the original PID is 13b long and ISI is only 8b, actual hardware filters limit the maximum number of active PIDs per multiplex (e.g. to 32), thus making this mapping possible in some cases.

If Generic Streams are to be mapped into BBFRAMES, ISI allocation and use will have to be defined from scratch. For this, static or dynamic tables must be specified, using when possible the basis and the experience of those in use over MPEG2. ISI signalling for Generic Streams is a key issue for building an IP sub-network over a DVB-S2 link, since GS are a very strong candidate to replace MPEG2 bearers for the carriage of IP datagrams over satellite links.

4.3. Address Resolution Issues

Logical channels allow differentiation of upper layer flows within a BBFRAME multiplex. A mapping function -to be defined- can provide the means to associate dynamically a network flow to a particular logical channel of a single multiplex. This mapping function will be the main basis over which Address Resolution will be done in IP/

DVB-S2 networks.

Precise address resolution considerations are out of the scope of the present document, but previous work done on DVB-S encapsulation procedures can provide valuable input here. Furthermore, since BBHEADERS provide similar functionalities to the ones given by MPEG2 headers, it is highly probable that existing considerations for PIDs can be transposed directly to ISI.

4.4. QoS Mapping and MODCOD Selection

Under ACM mode, every Receiver will be able to adjust its reception parameters by the choice of a minimum protection MODCOD, in order to cope with rapid link fading (due to e.g. rain or interference) and ensure data reception under any condition.

From an IP point of view, MODCOD variations will only change the available bandwidth for the overall transmission. The ACM command will therefore have to schedule packets according to QoS considerations, filling in an optimal way the available BBFRAMES. This might imply rapid changes of packets queues, delaying and/or dropping PDUs.

Under DVB-S2 links, QoS will be a mix of 2 aspects: minimum protection required (MODCOD), and priority of transmission over other PDUs. Proper mapping of QoS requirements over MODCODs has not been done yet and this is a topic in which comments and suggestions are most welcome. In particular, there are many ways in which QoS requirements for every PDU can be passed to the scheduler. This information could be added to every encapsulation header, or IP flows could be distributed over different queues (each one having different QoS requirements) prior to encapsulation and packing. A less classical method could consist on making instant scheduling decisions at link layer with basis on information directly provided by every IP header.

Other basic and non-exhaustive ideas are presented here:

4.4.1. Cross-Layer Optimizations for QoS Scheduling Policies

The value of the TOS (Type Of Service) field in the IP header will certainly be taken in account for QoS mapping, but QoS differentiation procedures can also rely on complementary information.

A promising research possibility is to add specific application-level information to every single network packet (e.g. in the encapsulation header, complementing or overriding TOS), to determine the particular

BBFRAME in which it will be packed (that is, the particular MODCOD assigned to it) or its priority over other PDUs. It has been pointed out that using application-level information for link-level scheduling and MODCOD protection might prove interesting. An example of this could be an error-tolerant application (e.g. streaming), that simply may not need the maximal protection available at a given time for its particular Receiver: lowering its protection might free needed resources elsewhere and allow for processing and overhead saves. Research lead on the relation between link-level FEC and application-level FEC will certainly provide interesting input to this particular issue.

4.4.2. Differentiated QoS over Logical Channels

Instead of assigning an optimal MODCOD to every single PDU, a simpler idea is to associate a given QoS level with a particular Logical channel, in a static or dynamic way. Equivalent QoS levels could therefore have different instances according to the MODCOD they use.

Again, input concerning these particular issues is most welcome.

5. Motivations for a New Approach

Even though many existing solutions used in DVB-S can be extended or adapted to the new standard, the new features of DVB-S2 raise a number of important and interesting issues worth consideration.

5.1. ACM and DVB-S2 Framing

Through the use of ACM, the dynamic variations of the RF channel will have a direct impact over the physical waveform to be used for every piece of data to be transmitted. Analysis of MODCOD allocation policies due to channel variations might open field for IP carriage efficiency improvement, and several competing resource allocation algorithms should be tested.

On the other hand, the presence of BBFRAMEs measuring up to 7 kB (around 37 times the size of a MPEG2 packet) suggests that in many cases several full PDUs will be able to fit together in a single DATAFIELD. Making reasonable assumptions about the fragmentation complexity allowed by a future encapsulation, Segmentation and Reassembly (SAR) should therefore occur less often than in DVB-S, which raises other kind of questions and risks. In particular, a partially filled and sent BBFRAME will make worse use of the RF resource than a partially filled and sent MPEG2 packet. In contrast, optimal BBFRAME filling should allow optimal resource use as it minimizes redundant overhead. Indeed, available PDU encapsulations

such as MPE/MPEG2-TS, ULE/MPEG2-TS or AAL5/ATM were designed for a relatively high probability of SAR use during SNDU processing. The definition of a scheduling algorithm ensuring optimal BBFRAME filling will certainly be a key point for improving IP carriage efficiency.

5.2. IP over Generic Streams

TS constant end-to-end delay and constant bit-rate are not a must for IP services, and could be overridden if a GS solution were considered for IP carriage. On top of that, the mandatory asynchronous mapping of MPEG2 over the BBFRAMEs directly constitutes a triple drawback from an IP point of view. First, it adds significant complexity and processing to the overall process. Second, the eventual overhead (here, padding) generated by this asynchronous mapping will add to the overall overhead of the IP encapsulation over MPEG2-TS, decreasing global efficiency. Finally, unexpected hardware/software functioning that may affect proper reassembly of fragmented MPEG2 packets will directly impact PER at IP level.

The use of MPE and recently ULE to convey IP data over MPEG2-TS has contributed over the past years to its wide acceptance as a IP subnetwork technology. However, despite the unquestionable services done by MPEG2-TS in the DVB-S context, the use of GS in a DVB-S2 system for conveying IP data seems to offer better perspectives. In order to take full advantage of the handy features of GS and ACM, and stick to the key goals specified in [Section 1](#), new solutions have to be considered. Those should rely on already available proved concepts when possible, but should also innovate where existing mechanisms do not offer a satisfactory solution.

6. General Framework Requirements

Detailed requirements for transmission of IP datagrams over MPEG2-TS networks have been defined in [RFC 4259](#) [7]. The present section will therefore focus on the requirements for transmission of IP datagrams over continuous Generic Streams in ACM mode. Note however, as stated in [Section 3.2](#), that seen from a BBFRAME, there is no difference between a MPEG2-TS and a packetized Generic Stream with packets of length 188B. From this perspective, MPE or ULE could be used for encapsulation of upper layer datagrams over a packetized Generic Stream, although this would be a very inefficient solution. Indeed, in addition of artificially introducing the complexity and overhead of the MPEG2 layer, advanced fragmentation and adaptive scheduling could not be used, due to the flexibility limitations of the stream-based MPEG2 approach.

This section is concerned with the efficient carriage of IP datagrams

over continuous Generic Streams relying on an adaptive link/physical layer. Fundamental differences with a TS approach will therefore be pointed out when possible.

6.1. Use of Existing Encapsulation Capabilities

In the general encapsulation case, PDUs are formatted one by one into SNDUs, by receiving a particular header and an integrity check trailer such as a CRC. When required, SNDUs are fragmented and their fragments are sent over the link using one or more bearers (e.g. MPEG2 packets or BBFRAMES). The encapsulation header provides delineation information to the Receiver, so that received SNDU fragments can be located and properly assembled. The end-to-end integrity check stands as a protection against re-assembly errors.

However, MPEG2 packets are encapsulated into BBFRAMES without the need of additional encapsulation headers, since BBHEADERS provide all the functionalities necessary to delineate and reassemble fragmented packetized streams. BBHEADERS have therefore some inherent encapsulation capabilities, and a future framework intended to standardize an IP over GS interface could take advantage of this handy feature.

Of course, IP streams are not composed of fixed-length packets and the above-mentioned encapsulation capabilities of BBHEADERS do not directly apply. However, several concepts used in classical encapsulation headers (e.g. Payload Pointer or Length Field for ULE) could be transposed if IP packets were to be mapped over Generic Streams, using available fields in BBHEADERS (e.g. SYNC and UPL/DFL).

Note finally that among the 10 bytes of BBHEADERS, at least 3 (SYNC and SYNCN) are not relevant for continuous GS. Their re-definition and use in an "IP over continuous GS" context might prove useful as this would allow reducing the header length of a future encapsulation.

6.2. Fragmentation Issues. Adaptive Packing and Scheduling Optimization

In order to ensure optimum use of the available resources, it is required that BBFRAMES addressed to a single NPA carry as much data as possible, and therefore SNDU fragmentation is important. Furthermore, since BBFRAMES are considerably longer than classical bearers, optimal packing of the SNDU fragments according to QoS or size criteria is a key point for achieving efficient PDU carriage. This is a major difference with DVB-S, in which SNDU fragments are sent over the next MPEG2 bearer available, regardless of their sizes or required QoS.

6.2.1. Fragmentation Schemes to be Supported

Several fragmentation schemes supported by the encapsulator with increasing complexity can be imagined. For the authors, the 4 examples described below represent the maximum fragmentation levels that should be implemented in an encapsulator in order to avoid excessive Receiver complexity while ensuring optimum scheduling management. Note that cases b. c. and d. do not occur in the MPEG2 stream-based approach, and therefore cannot be managed by simple ULE-like SAR procedures.

The following BBFRAMES belong to the same logical channel, and are therefore seen by all the Receivers correctly tuned. For the examples, we focus only on fragmentation of SNDU2 into 2 fragments SNDU2a and SNDU2b, carried by 2 different BBFRAMES. Those 2 BBFRAMES do not necessarily use the same MODCODs, but we suppose that the Receiver assembling SNDU2 can at least decode/demodulate them correctly. Note that fragmentation of an SNDU over more than 2 fragments is of course possible.

a. Classical consecutive fragmentation

```
+-----+-----+ +-----+-----+
|  SNDU1  | SNDU2a | |SNDU2b|   SNDU3   |
+-----+-----+ +-----+-----+
```

This is the most natural and intuitive one, since it is classically used over MPEG2 bearers. In this case, SNDU2 is sent over 2 consecutive BBFRAMES, using the end of the first BBFRAME and the start of the second one.

b. Consecutive fragmentation with arbitrary SNDU fragment placement

```
+-----+-----+ +-----+-----+
|  SNDU1  | SNDU2a | |  SNDU3   |SNDU2b| SNDU4 |
+-----+-----+ +-----+-----+
```

After sending the beginning of SNDU2, the scheduler decides to use the beginning of the following BBFRAME to send SNDU3. This latter SNDU could be addressed to another Receiver for example, or be addressed to the recipient of SNDU2 but bear more priority. This is not necessarily related to a change of MODCOD after the first BBFRAME, although a narrowing of the bandwidth of the channel can motivate this kind of decisions. SNDU2 is resumed using available space in the following bearer, which implies basically that SNDU2b does not start at the beginning of the second BBFRAME.

c. Non-consecutive fragmentation

```

+-----+-----+ +----- // -----+ +-----+-----+
|  SNDU1  | SNDU2a | |  SNDU3  -->  SNDU7 | |SNDU2b|      SNDU8      |
+-----+-----+ +----- // -----+ +-----+-----+

```

After placing fragment SNDU2a, the scheduler decides to delay placement of SNDU2b in order to send 5 SNDUs, labelled 3 to 7. After their transmission has finished a few BBFRAMEs later, SNDU2b is placed and sent in the beginning of the next available bearer. In this case, the Receiver needs to identify the BBFRAME carrying SNDU2b. Note that SNDUs #3 to #7 might have been fragmented as well over the intermediate BBFRAMEs, so that there might be more than 1 SNDU with suspended fragments within the same BBFRAME flow.

Such situation will happen e.g. because of a MODCOD change: intermediary BBFRAMEs might be using a very efficient MODCOD, so that the Receiver assembling SNDU2 would not be able to decode/demodulate them. In another scenario, they could have the same MODCOD as the previous ones, but be addressed to another Receiver of the same logical channel bearing more priority than the one assembling SNDU2.

d. Non-consecutive fragmentation with arbitrary SNDU fragment placement

```

+-----+-----+ +----- // -----+ +-----+-----+-----+
|  SNDU1  | SNDU2a | |  SNDU3    -->  SNDU7a| | SNDU7b |SNDU2b| SNDU8 |
+-----+-----+ +----- // -----+ +-----+-----+-----+

```

This case is a combination of cases b. and c. The Receiver needs to identify the BBFRAME carrying SNDU2b as well as its position inside the BBFRAME. This solution seems to offer the greatest flexibility vs. reasonable complexity trade-off in the DVB-S2 context.

Note finally that in all previous schemes, SNDU2a is always at the end of a given BBFRAME. A more generic fragmentation scheme could allow for a free placement of SNDU2a in the BBFRAME flow (not necessarily at the end of a BBFRAME), although this solution does not seem much more performant than d.

6.2.2. Packing Optimization

MODCOD allocation by the ACM command is closely related to packing optimization, since available DATAFIELD sizes will vary according to the dynamics of the channel. A scheduling algorithm is required to optimize filling and minimize BBFRAME Padding, that may be up to

7264B for an empty DATAFIELD. It should provide ways to fragment, re-order SNDUs and delay them when necessary for the sake of optimal filling, but always in the limits of an admissible complexity. In particular, packet re-ordering between different IP flows to optimize BBFRAME filling is encouraged, while fragment reordering within a single flow of IP packets (that is between 2 fixed ports of 2 end hosts) should be avoided, according to [RFC 3819](#) [6].

A scheduling algorithm should take in account the statistical characteristics of the carried IP traffic, and its functioning should not be independent from the ACM command. It should also provide BBFRAME Padding when necessary (when no PDU is ready to be encapsulated).

6.3. Next Level Protocol Type

A key feature of the required encapsulation is to identify the payload type being transported. Such requirement is not specific to DVB-S2: most protocols use a type field to identify a specific process at the next higher layer that is the originator or the recipient of the payload.

Given the length of BBFRAMES, several PDUs will often be packed within the same BBFRAME. Possible ways to differentiate protocol types to which PDUs belong are:

- ISI channels. This requires no overhead and demands that only PDUs from (or to) the same protocol can be sent together in a single BBFRAME. The use of ISI for this purpose can interfere with its use for address resolution or QoS mapping.

- A single Type field per BBFRAME (ex: appended to the BBHEADER or inside it) in an homogeneous traffic environment (e.g. an IPv4-only network). Only homogeneous PDUs (that is, originated or going to a same protocol) will be packed together. This solution produces very small overhead but offers low flexibility for future evolution of the traffic mix.

- A Type field per PDU. In an heterogeneous traffic environment (e.g. a mix of IPv4 and IPv6 packets), it is required that every single PDU is labelled with a proper Type field. This solution produces an overhead proportional to the number of transported PDUs but offers no limits in its flexibility, since the detailed composition of the traffic mix do not affect the encapsulation procedures.

In a context of IPv4 to IPv6 migration and of increased use of the

Internet by new applications and users, the last solution seems to be the most adapted. It is also the choice done in ULE.

6.4. Precise Payload Unit Delineation and Reassembly

Accurate delineation and identification of scattered SNDU fragments must be done by every Receiver. As an example, ULE achieves delineation with the joint use of MPEG2's PUSI, a Payload Pointer and a Length field.

Precise PDU delineation is also required for a future encapsulation over Generic Streams. The implemented solution may define a ULE-like header for this, but it may also re-use (partially at least) BBHEADER fields that already provide similar functionalities. It should also be robust to synchronization losses, for which a "Payload Pointer plus Length field" approach could prove desirable.

On the other hand, a future encapsulation must provide ways to ensure reassembly of a scattered SNDUs even in the case that its fragments are not "adjacent" within 2 consecutive BBFRAMES, which could happen if advanced SNDU scheduling/fragmentation procedures are used. In the classical ULE/MPEG2-TS/DVB-S scenario, SNDU fragmentation is done over MPEG2 packets of the same flow (same multiplex and PID) with Continuity Counters differing only by 1 (modulo 16). This means that a ULE Receiver knows in advance the size and position in the flow of the next SNDU chunk needed for proper reassembly. However, in a DVB-S2 context, the scheduling algorithm may chose to send SNDU fragments over non-consecutive BBFRAMES, or place SNDU fragments in the middle of a given BBFRAME (cases b., c. and d. of [Section 6.2.1](#)). Since ULE does not provide tools to locate scattered SNDU fragments with a priori unknown positions and lengths in a BBFRAMES multiplex, it cannot be used directly in a GS context. More advanced reassembly procedures will certainly have to be studied.

6.5. Integrity Check

For the IP service, the probability of undetected packet error should be small or negligible, as stated in [RFC 3819](#) [6]. There is therefore a need for a strong error control, either provided by FEC or by other means such as CRCs. FEC has been greatly improved in DVB-S2, and it provides means to re-evaluate the way integrity check can be done. The following considerations apply also for packetized streams.

The CRC-32 used in ULE relies on the use of a 32-bit redundancy for each SNDU, intended to stand as a protection against reassembly errors following corruption or loss of SNDU bearers, due to transmission errors or unexpected hardware/software operation.

Concerning resilient channel errors, DVB-S2 FEC has reduced the probability of such event by several orders of magnitude compared to DVB-S. Indeed, the implemented LDPC and BCH concatenated encoding provide error protection close to the theoretical limits established by Shannon in [13]. On top of that, and in addition to their correction capabilities, outer BCH decoders can provide very accurate error-detection information that could be used to detect and discard corrupted BBFRAMES. DVB-S2 FEC offers therefore the possibility to manage globally the SNDU error-detection problem, done classically on a SNDU-by-SNDU basis with CRCs. Details concerning these particular issues are fully analyzed in [14].

As for BBFRAME losses, only SNDUs with fragments in lost BBFRAMES will face reassembly problems. A non-fragmented SNDU within a lost BBFRAME will be simply lost, even if it had a CRC. In this context it seems adequate to apply CRC integrity checks to the PDUs that may suffer SAR.

Accurate estimation of PER at IP level with or without CRCs should drive the design decisions concerning this issue, and not legacy considerations based on different or less efficient systems.

6.6. Link Layer Addressing Capabilities

Individual Receivers are not addressable at a BBFRAME level. MPEG2-TS addressing considerations exposed in [RFC 4259](#) [7] apply therefore to BBFRAMES too and should be used as guidelines for future work on this key topic.

These considerations imply the use of an optional NPA field appended to every PDU or group of PDUs sharing the same BBFRAME. There are indeed cases where the use of a NPA is required (e.g. when link layer filtering is desired) and if present, it should be carried in a way to allow Receivers to pre-filter and discard unwanted PDUs. There are other cases where an NPA is not required (e.g. when a Receiver is the end host), and network layer filtering may be used.

An IP over GS interface should therefore support an optional NPA, as current encapsulations for IP over MPEG2-TS do. The interactions and synergies that can be achieved with the use of BBFRAMES' logical channels defined in [Section 4.2](#) are to be analyzed.

6.7. Flexibility for Future Extension

The evolution of the Internet service may in the future require additional functions. A flexible encapsulation protocol should therefore provide a way to introduce new features when required, without having to provide additional out-of-band configuration. This

is not a difference with TS, and a native way to signal header extensions (like the Next-Header protocol type in ULE) should be implemented.

6.8. Security Considerations

Security considerations are basically the same for GS and TS, and are based on those concerning wireless links, see [RFC 3819](#) [6]. Although an analysis of specific security issues concerning GS is out of the scope of this document, it will provide helpful input for addressing this important topic in future work.

Current work on security issues for ULE carried out at the IPDVB working group should certainly integrate considerations regarding DVB-S2.

7. Summary

DVB-S2 offers new possibilities for deploying IP services over satellite links, as a successor of DVB-S that offers many IP-friendly capabilities.

The new standard's physical adaptability and new framing structure motivate therefore a new encapsulation for IP, much more efficient in terms of complexity and overhead than the classical MPEG2-TS approach.

For this, some solutions developed for IP/MPEG2 can be simply transposed to DVB-S2, like the use of Type fields or the definition of logical channels. However, the full use of DVB-S2 specificities will also need new procedures -like datagram scheduling optimization and advanced fragmentation- to be defined from scratch. Finally, other issues like integrity check management might be re-evaluated in a DVB-S2 context, taking in account the enhanced error-control achieved by the new FEC.

Finally, DVB-S2 BBFRAMES are defined in such a way that BBHEADERS present some inherent encapsulation capabilities. The definition of a new IP over DVB-S2 adaptation layer could take advantage of this handy feature, and open the way for other cross-layer synergies in order to improve overall system efficiency.

8. Acknowledgements

This document is a result of discussions at IETF 63 and 64, as well as on the ipdvb mailing list. Special thanks to the guidance of

ipdvb chairman G. Fairhurst, and to all those who, by their curiosity, questions, critics and remarks have made the scientific debate concerning DVB-S2 grow.

This draft is intended as a study item for proposed future work by the IETF in this area. Comments relating to this document will be gratefully received by the authors and the ipdvb mailing list at: ipdvb@erg.abdn.ac.uk

9. Document History

- 00 This draft is intended as a study item for proposed future work by the IETF in this area.
- 01 Draft re-written from scratch. Added a comprehensive glossary and a complete introduction to DVB-S2.
- 02 Focused the draft on the delivery of IP service over DVB-S2, rather than on a requirements specification. Integrated a converged view of possible fragmentations, and introduced important axes for future cross-layer optimization of the IP/DVB-S2 protocol stack. References updated.
- 03 and 04: Minor modifications: withdrawn one author and updated one reference.

10. References

10.1. Normative References

- [1] "EN 302 307, Digital Video Broadcasting (DVB); Second generation framing structure, channel coding and modulation systems for Broadcasting, Interactive Services, News Gathering and other broadband satellite applications. European Telecommunications Standards Institute (ETSI)".
- [2] "EN 301 421, Digital Video Broadcasting (DVB); Modulation and Coding for DBS satellite systems at 11/12 GHz, European Telecommunications Standards Institute (ETSI)".
- [3] "ISO/IEC DIS 13818-1: Information Technology; Generic Coding of Moving Pictures and Associated Audio information: Systems, International Standards Organisation (ISO)".
- [4] Fairhurst, G. and B. Collini-Nocker, "Unidirectional Lightweight Encapsulation (ULE) for transmission of IP datagrams over an MPEG-2 Transport Stream", [RFC 4326](#), December 2005.

10.2. Informative References

- [5] "EN 301 192 Specifications for Data Broadcasting, European Telecommunications Standards Institute (ETSI)".
- [6] Karn, P., Borman, C., Fairhurst, G., Grossman, D., Ludwig, R., Mahdavi, J., Montenegro, G., Touch, J., and L. Wood, "Advice for Internet Subnetwork Designers", [RFC 3819](#).
- [7] Montpetit, M., Fairhurst, G., Clausen, H., and H. Linder, "A Framework for Transmission of IP Datagrams over MPEG-2 Networks", [RFC 4259](#), November 2005.
- [8] "EN 301 790, Digital Video Broadcasting (DVB); Interaction Channel for Satellite Distribution Systems. European Telecommunications Standards Institute (ETSI)".
- [9] Jacobson, V., "Compressing TCP/IP Headers for Low-Speed Serial Links", [RFC 1144](#).
- [10] Degermark, M., Nordgren, B., and S. Pink, "IP Header Compression", [RFC 2507](#).
- [11] Bormann et al., C., "RObust Header Compression (ROHC): Framework and four profiles: RTP, UDP ESP and uncompressed", [RFC 3095](#).
- [12] Kent, S. and R. Atkinson, "Security Architecture for the Internet Protocol", [RFC 2401](#).
- [13] Shannon, C., "A Mathematical Theory of Information", 1948.
- [14] Cantillo, J., Lacan, J., and I. Buret, "Cross-layer enhancement of error control techniques for adaptation layers of DVB satellites", International Journal of Satellite Communications and Networking, special issue on Cross-Layer Protocols for Satellite Communication Networks, Volume 24, Issue 6, Pages 579-590, November/December 2006.

Authors' Addresses

Juan Cantillo
ENSICA/TESA/Alcatel Alenia Space
1, place Emile Blouin
Toulouse 31056
France

Email: juan.cantillo@ensica.fr

URI: http://dmi.ensica.fr/auteur.php3?id_auteur=61

Jerome Lacan
ENSICA/LAAS-CNRS
1, place Emile Blouin
Toulouse 31056
France

Email: jerome.lacan@ensica.fr

URI: http://dmi.ensica.fr/auteur.php3?id_auteur=5

Intellectual Property Statement

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in [BCP 78](#) and [BCP 79](#).

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

Disclaimer of Validity

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Copyright Statement

Copyright (C) The Internet Society (2006). This document is subject to the rights, licenses and restrictions contained in [BCP 78](#), and except as set forth therein, the authors retain all their rights.

Acknowledgment

Funding for the RFC Editor function is currently provided by the Internet Society.

