Design Issues for the Generic Stream Encapsulation (GSE) of IP Datagrams over DVB-S2

Juan Cantillo* Jérôme Lacan† Isabelle Buret‡ Fabrice Arnal‡

*TéSA/Thales Alenia Space, Univ. of Toulouse, France
†ENSICA/TéSA/LAAS, Univ. of Toulouse, 1, place E. Blouin, 31056 Toulouse Cedex, France
‡Thales Alenia Space, 26, avenue J.F. Champollion, B.P. 1187, 31037 Toulouse Cedex 1 - France
Email: juan.cantillo@supaero.org, jerome.lacan@ensica.fr, {Isabelle.Buret, Fabrice.Arnal}@thalesaleniaspace.com

***Abstract*—The DVB-S2 standard has brought an unprecedented degree of novelty and flexibility in the way IP datagrams or other network level packets can be transmitted over DVB satellite links, with the introduction of an IP-friendly link layer -the continuous Generic Streams- and the adaptive combination of advanced error coding, modulation and spectrum management techniques. Recently approved by the DVB, the Generic Stream Encapsulation (GSE) used for carrying IP datagrams over DVB-S2 implements solutions stemmed from a design rationale quite different from the one behind IP encapsulation schemes over its predecessor DVB-S. This paper highlights GSE's original design choices under the perspective of DVB-S2's innovative features and possibilities.**

I. General Introduction

The uses and performances of the Multi Protocol Encapsulation (MPE) [1] and the Unidirectional Lightweight Encapsulation (ULE) [2] have been widely analyzed in the literature, and they are commonly accepted as the standard ways to carry IP datagrams over DVB satellites. Truth is, their design was constrained by the imperatives of using already deployed DVB satellite architectures built over the MPEG2-TS [3] link layer, a technology optimized for media broadcasting and not for IP services delivery. 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.

Recently approved by the European Telecommunications Standards Institute (ETSI), the DVB-S2 [4] architecture uses the most recent advances in physical layer technology with the unprecedented possibility in DVB networks to carry network layer datagrams without the use of the MPEG2-TS link layer paving the way to efficient and more flexible IP carriage over satellite links. It appeared soon that the existing mechanisms to encapsulate IP datagrams or Protocol Data Units (PDUs) over DVB-S offered could not fully exploit the innovative features of the new standard, for which a novel encapsulation had to be proposed. The resulting Generic Stream Encapsulation (GSE) [5] has been designed with the specific characteristics of DVB-S2 in mind, providing all the necessary methods to fully exploit its enhanced capacity, reliability and flexibility.

This paper summarizes and updates the contents of [10] for informative purposes

The purpose of this paper is to expose the rationale behind the original design choices made for GSE under the lights of DVB-S2's new features, explaining GSE's new approach for IP datagrams transmission over DVB satellite links. Rather than presenting a detailed description of GSE or DVB-S2, it highlights the way GSE fully allows DVB-S2 to keep its promises at the network layer, stressing the points where it brings originality where previous solutions would fail. For this, after a review of the essential points that make DVB-S2 original from its predecessor from a network layer point of view, the second part of this paper focuses on the specific characteristics and originalities of GSE. Finally, a last part will discus possible extensions for GSE and its evolution in the context of digital communications.

II. Innovative Features in DVB-S2 Affecting the Design of an Encapsulation Layer

Just as any other adaptation layer, GSE provides basic functions related to network protocol identification, addressing and hardware filtering, flow delineation, PDU fragmentation and integrity checks. However, compared to its predecessor, DVB-S2 features enhancements both in its physical and link layers that have a major impact on the way such classical functions are to be implemented to achieve efficient PDU transmission.

A. Physical Layer Enhancements

DVB-S2 implements the most recent developments in modulation and channel coding, with the use of QPSK, 8-PSK, 16-APSK, 32-APSK and especially, the use of concatenated Bose-Chaudhuri-Hocquenghem (BCH) and Low Density Parity Check (LDPC) codes. The LDPC code rate can be chosen among 11 values: $1/4$, $1/3$, $2/5$, $1/2$, $3/5$, $2/3$, $3/4$, $4/5$, $5/6$, $8/9$ and $9/10$, for a resulting family of concatenated Forward Error Coding schemes (FEC) only 0.6 to 0.8 dB away from the Shannon limit [6], easily ensuring the Quasi Error Free (QEF) target of $BER < 10^{-10}$ at the input of the network layer. On the other hand, it has been pointed out that in DVB-S2 link layer frames are either correctly decoded or totally messed up in the functioning E_s/N_0

domain, and that the probability of resilient or undetected errors after FEC decoding in DVB-S2 has been lowered by more than 3 orders of magnitude [7] compared to DVB-S. As a major consequence for IP, encapsulated datagrams in DVB-S2 should be more protected and less exposed to resilient channel errors than in DVB-S.
Combined use of higher order modulations and powerful channel coding allows covering a wide range of E_s/N_0 values from -2.35 dB to 16.05 dB, enlarging considerably the functional domain of the new standard over DVB-S, and increasing *de facto* its raw transmission capacity over more than 35% in terms of spectral efficiency [6] [8]. When used for interactive point-to-point applications like IP unicast, theoretical analyses and simulations point out that DVB-S2 performs even better, providing an increase in transmission capacity by a remarkable 150% [9].

In order to take full advantage of this flexibility, the new standard provides richer alternatives to the classical Constant Coding and Modulation (CCM) approach. The new Variable Coding and Modulation (VCM) functionality allows 28 different combinations of modulations and error protection levels, labeled as MODCODs to be used and changed on a frame-by-frame basis. This may be combined with the use of a return link -either satellite, such as DVB-RCS, or terrestrial- to achieve dynamic closed-loop Adaptive Coding and Modulation (ACM), thus allowing the transmission parameters to be optimized for each individual user, on a frame-by-frame basis, according to individual link conditions. This means that the physical layer can provide differentiated Quality of Service (QoS) levels, a major difference with DVB-S where a single CCM mode existed for every receiver. QoS requirements from the upper layers (e.g. DiffServ) could therefore be mapped into physical layer MODCODs, e.g. making use of new cross-layer techniques. Although the definition of those mechanisms -including a packet scheduling policy- are out of scope of the design of an encapsulation scheme, an acceptable adaptation layer for DVB-S2 should clearly provide methods to implement QoS-related scheduling decisions, and to allow for flexible PDU placement and enhanced fragmentation in the flow in order to fully exploit DVB-S2 adaptability. For this reason, since MPE and ULE-like encapsulations provide PDU fragmentation over consecutive frames exclusively, their use, although possible, would be suboptimal in the DVB-S2 context.

B. Link Layer Enhancements

In addition to the classical packetized MPEG2-Transport Streams, DVB-S2 introduces the new *Generic Streams* (GS) above its physical layer, intended to address the non-native way in which network services -such as IP- are carried today over MPEG2-TS using MPE or ULE. Generic Streams can be packetized or continuous: the former are particularly suited for carrying fixed-length Protocol Data Units (PDU) such as MPEG2 packets or ATM cells, whereas the latter have been designed to accommodate smoothly any kind of input stream format, including continuous bit-streams and PDUs of variable size such as IP datagrams.
As an important addition over DVB-S, before FEC coding both Generic Streams and Transport Streams are tailored into a series of 21 possible *BBFRAMEs* offering the different efficiency vs. error protection tradeoffs used in the adaptive transmission modes, with predefined sizes in the range $[384B; 1779B]$ (*short* BBFRAMEs) and $[2001B; 7274B]$ (*long* BBFRAMEs). BBFRAMEs' sizes match the input block lengths of the outer BCH codes in DVB-S2, which make them the true basic link-level units of any DVB-S2 stream. Furthermore, BBFRAMEs' headers provide some inherent encapsulation capabilities, that allow to fragment and transmit MPEG2 packets or any other fixed-size packets asynchronously mapped over BBFRAMEs without incurring into additional overhead [10]. Conceptually, in DVB-S2, MPEG2 packets are dealt with as simple network-level packets and no longer as link-level bearers, as it was the case in DVB-S.
The choice of continuous Generic Streams for IP datagrams transmissions presents obvious advantages over MPEG2-TS: first, non relevant constraints for interactive services such as constant bit-rate and end-to-end delay can be totally by-passed, allowing for faster and better datagram delivery at reduced overhead and processing complexity. Second, QoS-related rapid changes in the flow structure taken aside, packet fragmentation should occur rather seldom given the large BBFRAMEs payload size, up to 40 times broader than a single MPEG2 packet. Measurements on the Internet backbone point out that the frequency-weighted average size of an IP datagram is around 500 bytes, so a rough average of $7000/500 \simeq 14$ full IP datagrams should in principle be carried in long BBFRAMEs, whereas in average every single IP datagram suffers 2 to 3 fragmentations on top of the MPEG2-TS layer and up to 10 when using ATM. Current integrity check policies based on frequent fragmentation use (in particular, the use of a systematic CRC per PDU in MPE and ULE) could therefore be reassessed in the sake of resource optimization, leaning partially on the system's increased FEC power. As an obvious and direct consequence, BBFRAMEs using continuous Generic Streams can be expected to accommodate several datagrams in their payloads simultaneously paving the way for several interesting optimization choices; a quite rare situation today where short link layer payloads are most commonly used.

III. Specificities of the GSE Protocol

This section focuses on the particular specificities that make GSE fit to take the maximum advantage of the innovative features of DVB-S2. For a detailed description of the protocol itself, please refer to [5] and [11].

A. GSE Basic Principles

The basic principles of GSE are similar to those of ULE or MPE. 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. As shown in Figure

1, PDUs scheduled for transmission are encapsulated in one or more GSE Packets. The encapsulation process uses an encapsulation header providing control information for fragmentation and link layer filtering, and adds an overall PDU integrity check in the form of a CRC-32, when needed as discussed in the following paragraphs.

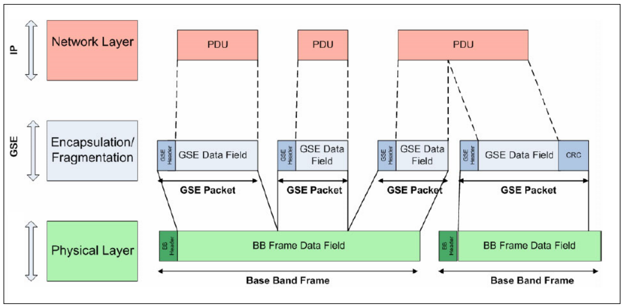

Fig. 1. GSE encapsulation within the DVB-S2 protocol stack [5]

Each GSE Packet is composed of a GSE Header followed by a GSE payload, where the (fragment of the) encapsulated PDU is located. The GSE Header is composed of the fields shown in Figure 2. The unshaded fields are always present, while the shaded fields may be omitted depending on the preceding control fields in the first 4 bits of the GSE Header. The presence of possible Extension Headers is determined by the Protocol Type value, following ULE-like rules and IANA allocations. The minimum GSE Header length is therefore 2 Bytes.

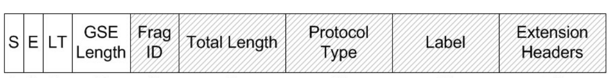

Fig. 2. GSE Header Format [5]

Rather than providing a precise description of the protocol, the following paragraphs highlight the major functionalities provided by the GSE Header under the lights of the new features of DVB-S2, providing comparisons with ULE and MPE and their shortcomings in dealing with them.

B. Flexible Fragmentation and PDU fragments placing

The full exploitation of the adaptive features of DVB-S2 is the main challenge in the design of GSE. Contrary to MPE or ULE where PDUs had to be simply mapped into an organized series of identical frames (namely, MPEG2 packets), MODCOD allocation on a frame-by-frame basis introduces additional difficulties in the manipulation of carried PDUs. Under CCM and even VCM modes, ULE and MPE-like fragmentation and delineation procedures could be used. However, this is not the case under ACM. The non-deterministic nature of the MODCOD sequence allows for cases in which the first fragment of a PDU is being sent over particular MODCOD but its next (or final) fragment cannot be sent immediately in the next frame, e.g. because a sudden change in the transmission conditions needing a more restrictive MODCOD. GSE provides therefore methods to *pause* and *resume* a PDU transmission at any time with its *Start* and *End* flag bits, a complete novelty over ULE and MPE, where fragmentation of a PDU implied the immediate transmission of its remaining pieces in the following available MPEG2 packet. This is the reason why in GSE, encapsulated PDU fragments, rather than full PDUs, are the real Subnetwork Data Units (SNDUs) of the adaptation layer.

Given that such MODCOD changes or scheduling events may occur at any time, situations in which a fragmentation has to be put on hold for a new one -with higher priority- to be done may arise easily, for instance if two successive MODCOD downgradings occur in a small period of time. Thanks to its *Frag ID* field, GSE provides a method to identify different active fragmentations and therefore to allow several parallel fragmentations to be carried simultaneously: combined use of the *Frag ID* field and the *(S)tart* and *(E)nd* bits provides a powerful method to assemble scattered PDU pieces upon reception, provided that some basic rules defined in [5] are followed. Fragments placement in BBFRAMEs can therefore be done with great flexibility, allowing for a great deal of freedom in the scheduler decisions in the sake of system adaptability.

C. Selective Integrity Checks

For the IP service, the probability of undetected packet error should be small or negligible. Physical layer reliability has been greatly enhanced in DVB-S2, meeting the targets defined in the best common practices defined in RFC 3819 [12] under QEF operation [7]. In addition, as stated in the previous section, the frequency of PDU fragmentation -and therefore the frequency of reassembly errors- has been lowered by the use of large payload sizes. For these reasons, the use of a CRC per single PDU following the design legacy of ULE or MPE is suboptimal and unadapted. In these encapsulation methods, every single PDU had a great probability to be fragmented given the comparable size of average PDUs (around 500 bytes for a representative IP datagram) and link layer payloads (184 Bytes for MPEG2 packets). In GSE, a CRC-32 calculated over the overall packet is appended only to the last PDU fragment of a fragmented PDU (identified with the flags SE = 01 in the GSE Header) to verify the correctness of the reassembly operation, allowing up to 10% overhead reduction for short packets such as TCP ACKs.

D. Overhead Considerations

Given the large sizes of BBFRAMEs (up to 7 kB), overhead has become a less important concern in GSE than for ULE or MPE, where link layer frames where about the same sizes of transported PDUs. Design concerns in GSE have focused on fragmentation flexibility and exploitation of the adaptive features of GSE rather than on the minimization of overhead and other unused bits, which explains apparent overhead inefficiencies *from an overhead point of view* such as the encapsulation of PDU fragments.

GSE's design has however made considerable steps towards the reduction of header overhead (which can be minimized to 2 bytes in particular contexts) and the elimination of CRCs for unfragmented PDUs. Given the big payload sizes available, overhead figures are more than ever sensitive to traffic characteristics and especially inter-arrival times obliging a latency vs. tolerable padding compromise to be found. Overall overhead figures calculated as a ratio between useful (non-header, non-padding, non-CRC) and sent bits in GSE have been reduced by a factor 2-3 with respect to MPE over MPEG2-TS, situating them around 2% to 3% with common IP traffic assumptions.

IV. Extensions for GSE

Although GSE's definition has only been achieved in the previous months, recent studies have already started to point out possible modifications and extensions for its use, either within the DVB-S2 context or in other next-generation DVB standards such *e. g.* DVB-SH [13]. Some of them could make use of the *ad-hoc* extensibility capabilities provided in GSE by the use of extension headers, defined by specific and IANA-allocated protocol types. Some others are inspired by emerging cross-layer techniques, and others adapt currently used methods, such as the GSE-FEC proposal of [14].

A. Label Re-use, PDU Concatenation

The *Label* field can be 0, 3 or 6 Bytes long, depending on the values of the mandatory 2-bit *LT* control flag. This field provides link layer addressing by binding a PDU to a Network Point of Attachment (NPA) such as a MAC address, a VCI/VPI in ATM, a DVB-RCS group/log-on etc following ULE-like rules as such defined in RFC 4259 [15]. Current work carried at the Internet Engineering Task Force (IETF) aims to exploit even better the large payload sizes available in BBFRAMEs with the introduction of the "PDU-Concat Extension", which enables a sequence of (usually short) PDUs to be aggregated and carried in a single GSE packet. This extension should be particularly useful in those frequent situations where several consecutive PDUs belong to the same logical flow and their common information such as protocol type, link layer address etc. can be somehow "factorized" according to rules and restrictions defined in [5]. The absence of repeated headers using this method added to the fact that unfragmented PDUs do not carry CRCs allows for a really light PDU encapsulation over Generic Streams.

B. Cross-Layer Enhancements for GSE

Recent studies have shown the interest of using cross-layer mechanisms in digital communications systems. The adaptive features of DVB-S2 make it a natural target for such optimization techniques, some of which could be implemented in future versions of GSE.

One example is using the information that may be available at the output of FEC decoders in DVB-S2 on the accuracy of the FEC decoding, unused today for practical purposes. Theoretical and experimental results show the high probability of detecting frame decoding errors by the BCH decoder [7] of DVB-S2, prior to CRC integrity checks in the adaptation layer. For most of the available FEC configurations, decoding errors that the FEC decoder itself can recognize as such are up to 10^8 times more frequent than undetected errors (i.e. those involving a mis-decoded frame). This suggests that a GSE receiver aware of the presence of errors in a given BBFRAME at the output of the FEC decoder could *immediately* discard it upon reception, saving processing at the upper layers and overhead due to the use of CRCs. Such a feature could off-load frame integrity checking from the adaptation layer with a guaranteed frame error rate better than 10^{-12}. Although not as tight as for the one achieved if CRC per BBFRAME covering the whole BBFRAME was used ($FER < 10^{-16.6}$), this bound is still is several times tighter than the QEF target and that specified in RFC 3819 [12] [11].

C. Bits Re-Use in BBFRAME Headers

Note that among the 10 bytes of BBFRAME headers, at least three (SYNC and SYNCD) are not relevant for continuous Generic Streams. Indeed, since their use has been defined in the DVB-S2 standard for the sole purpose of allowing native fragmentation and encapsulation of over BBFRAMEs for fixed-length PDUs carried over *packetized* Generic Streams. Their re-definition and use in an "IP over continuous Generic Streams" context might prove useful, and pave the way for further optimizations of future versions of the GSE protocol. Possible uses include allowing further flow organization, stamping BBFRAMEs for e.g. Operation and Management (OAM) purposes or adapting MODCOD selection based on network layer QoS signaling [11].

V. Conclusion

This paper presents GSE's originalities over MPE and ULE under the new features of DVB-S2. It particularly highlights the interests of this encapsulation protocol for IP services delivery. Some possible extensions are also proposed. This set of features should make GSE a good candidate for the encapsulation protocol of future DVB protocols.

Acknowledgements

The authors wish to thank the persons that made possible the definition of the GSE protocol, and that supported the scientific debate in the first stages of its definition. Particular thanks to Gorry Fairhurst and Bernhard Collini-Nocker.

References

[1] ETSI, "Specifications for Data Broadcasting, EN 301 192."
[2] G. Fairhurst and B. Collini-Nocker, "Unidirectional Lightweight Encapsulation (ULE) for Transmission of IP Datagrams over an MPEG-2 Transport Stream (TS)," RFC 4326 (Proposed Standard), Dec. 2005. [Online]. Available: http://www.ietf.org/rfc/rfc4326.txt

[3] ISO, "Information Technology; Generic Coding of Moving Pictures and Associated Audio Information Systems, ISO/IEC DIS 13818-1 ," 2000.
[4] ETSI, "Digital Video Broadcasting (DVB); Second generation framing structure, channel coding and modulation systems for Broadcasting, Interactive Services, News Gathering and other broadband satellite applications (DVB-S2), EN 302 307," 2004.
[5] DVB, "Generic Stream Encapsulation (GSE) Protocol," DVB BlueBook A116, May 2007.
[6] M. Eroz, F.-W. Sun, and L.-N. Lee, "DVB-S2 Low Density Parity Check Codes with Near Shannon Limit Performance," *International Journal of Satellite Communications and Networking*, vol. 22, pp. 269–279, 2004.
[7] J. Cantillo, J. Lacan, and I. Buret, "Cross-layer enhancement of error control techniques for adaptation layers of DVB satellites," *International Journal of Satellite Communications and Networking*, vol. 24, pp. 579–590, 2006.
[8] A. Morello and U. Reimers, "DVB-S2, the Second Generation Standard for Satellite Broadcasting and Unicasting," *International Journal of Satellite Communications and Networking*, vol. 22, pp. 249–268, 2004.
[9] R. Rinaldo, M. A. Vazquez-Castro, and A. Morello, "Capacity analysis and system optimization for the forward link of multi-beam satellite broadband systems exploiting adaptive coding and modulation," *International Journal of Satellite Communications and Networking*, vol. 22, pp. 401–423, 2004.
[10] J. Cantillo and J. Lacan, "A Design Rationale for Providing IP Services Over DVB-S2 Links," in *IETF draft, draft-cantillo-ipdvb-s2encaps, expired June 2007*.
[11] J. Cantillo, B. Collini-Nocker, U. D. Bie, O. D. Rio, G. Fairhurst, A. Jahn, and R. Rinaldo, "GSE: A Flexible, yet Efficient Encapsulation for IP over DVB-S2 Continuous Generic Streams," *Submitted*, 2007.
[12] P. Karn, C. Bormann, G. Fairhurst, D. Grossman, R. Ludwig, J. Mahdavi, G. Montenegro, J. Touch, and L. Wood, "Advice for Internet Subnetwork Designers," RFC 3819 (Best Current Practice), July 2004. [Online]. Available: http://www.ietf.org/rfc/rfc3819.txt
[13] DVB, "Framing Structure, channel coding and modulation for Satellite Services to Handheld devices (SH) below 3 GHz," DVB BlueBook A111, March 2007.
[14] J. Lei, G. S. Granados, and M. V. Castro, "MPE/ULE-FEC vs GSE-FEC Efficiency Comparison of IP Datagram Transmission over DVB-S2," in *25th AIAA International Communications Satellite Systems Conference*, May 2007.
[15] M.-J. Montpetit, G. Fairhurst, H. Clausen, B. Collini-Nocker, and H. Linder, "A Framework for Transmission of IP Datagrams over MPEG-2 Networks," RFC 4259 (Informational), Nov. 2005. [Online]. Available: http://www.ietf.org/rfc/rfc4259.txt