Encapsulation Requirements for Return Links and Mesh Systems over Satellite

Fabrice Hobaya, Cédric Baudoin, Emmanuel Dubois, Patrick Gélard, Emmanuel Chaput, André-Luc Beylot, Gorry Fairhurst

Encapsulation Requirements for Return Links and Mesh Systems over Satellite

Fabrice Hobaya[1], Cédric Baudoin[2], Emmanuel Dubois[3], Patrick Gélard[3], Emmanuel Chaput[4], André-Luc Beylot[4], Gorry Fairhurst[5].

[1]fabrice.hobaya@tesa.prd.fr, TéSA, University of Toulouse, France.
[2]cedric.baudoin@thalesaleniaspace.com, Thales Alenia Space, Toulouse, France.
[3]{emmanuel.dubois,patrick.gelard}@cnes.fr, CNES, Toulouse, France.
[4]{emmanuel.chaput,andre-luc.beylot}@enseeiht.fr, IRIT-ENSEEIHT, University of Toulouse, France.
[5]gorry@erg.abdn.ac.uk, Department of Engineering, University of Aberdeen, United Kingdom.

***Abstract*— Initially designed for TV broadcasting, DVB standard families have become mature technologies for Internet communications via satellite access networks. A generic encapsulation protocol is therefore required for the forward link, the return link, and for mesh satellite systems. Moreover, because of the cost and the scarcity of resources, this should be as efficient and light as possible. The GSE (Generic Stream Encapsulation) protocol enables efficient transport of any protocol, originally designed for DVB-S2, that is to say on Quasi Error Free (QEF) forward links of satellite systems. This paper addresses limitations in GSE for the return link context and examines the requirements for an encapsulation mechanism for satellite systems using multiple access. This includes the return link but also mesh satellite systems and any satellite systems that uses multiple access. This paper also analyses approaches to derive requirements and provides a first analysis of these requirements.**

***Keywords*— Encapsulation, multiple access, return link, mesh, OBP, satellite, requirements.**

I. Encapsulation in Satellite Systems

This section introduces the encapsulation issues for satellite systems and defines the main features of satellite systems that employ multiple and random access to resources.

A. Introduction to encapsulation in satellite systems

Initially designed for TV broadcasting, the DVB (Digital Video Broadcasting) satellite systems (DVB-S/DVB-S2) [1][2] are based on the MPEG2-TS frame structure. Satellite systems are no longer used only for TV broadcasting but also for Internet communications and mesh communications, that is communications between satellite terminals. The DVB-RCS standard [3] defines a satellite return-link that may support interaction for end users.

The bi-directional communications and the introduction of encapsulation protocols allows the DVB-RCS standard to support a wide range of services. The standard permits IP encapsulation using either a stack AAL5//ATM//DVB-RCS or a stack MPE//MPEG2-TS//DVB-RCS. In both cases, the overhead is important, mainly because of ATM cell header and the very small size of ATM cells in the first case and because of the MPEG2-TS heaviness (overhead) in the second case. The standards MPE [4] and ULE [5] allow the transport of any protocol (including IP) only over MPEG2-TS. This is not an optimal solution for the transport of IP because of the MPEG2-TS framing, initially designed for TV broadcasting.

B. Main features of satellite systems with a multiple access

Satellite systems support multiple access to allow many satellite terminals to dynamically share the same medium. The DVB-RCS standard is an example of such a network, where the Network Control Center (NCC) manages the resource. The following describes the main features of such a network.

The radio resource sharing relies on MF-TDMA dynamic allocation. Support for multiple and random access results in resource management and allocation algorithms that can be complex. To prevent wastage of resources and to address a large number of satellite terminals (ST), these algorithms require small physical layer (PHY) frames, which are typically smaller than the frame sizes available on the forward link.

In satellite systems a wide range of traffic profiles are possible. A star topology may commonly experience asymmetric traffic matrices, with many small packets (e.g. TCP ACKs) on the return link and larger packet on the forward link. In the professional market, mesh communications may lead to more symmetrical traffic matrices.

The paper is organized as follows. Section 2 introduces the GSE standard and its limitations in the context of multiple access satellite systems. Section 3 derives the principal requirements for design of a new encapsulation protocol suitable for the return link of satellite systems. Section 4 describes mesh communications and On Board Processing (OBP) satellite systems. Finally, section 5 concludes this paper.

II. The Generic Stream Encapsulation (GSE)

This section introduces the GSE [6] standard and its applicability to other physical layer beyond the DVB-S2.

A. Introduction to the GSE standard

GSE allows an "efficient encapsulation of IP over a "generic" physical layer" (e.g. GSE can be used for Generic Streams in a way analogous to ULE over TS over S2). It allows the encapsulation of any protocol using the protocol type field (as described in [7]). The extension header mechanism allows easy evolution, such as the additions described in [8]. Figure 1 shows IP encapsulation over DVB-S2 using GSE.

The design of GSE was optimized for use with the DVB-S2 standard and more generally for Quasi Error Free (QEF) physical layers with large PHY frames, where fragmentation overhead represents a small proportion of total overhead.

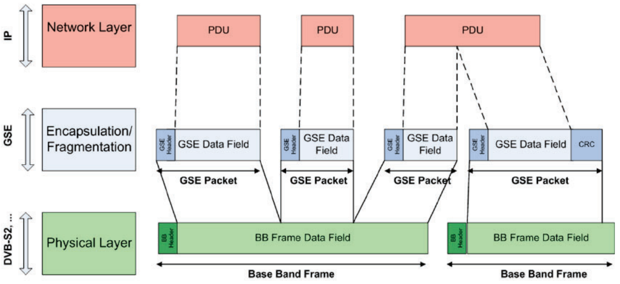

Figure 1: GSE encapsulation by DVB protocol stack [6]

B. Limitations in GSE for the return link context

The GSE overhead is too large for small Protocol Data Unit, PDU, (e.g. TCP ACKs and compressed TCP signaling). The fixed header part of the GSE overhead is 4 bytes (including a protocol type field of 2 bytes) as depicted in Figure 2. Considering a TCP ACK of 40 bytes, this leads to 10% of the overhead at level 2 (even more if fragmentation is required and/or if a higher layer compression mechanism like ROHC [9] is used). This overhead may be proportionally more for L2 control data. In the case of large PHY frames, the concatenation extension [8] allows a reduction of this overhead and can be very effective for some scenarios (e.g. VoIP trunking) [10]. GSE is efficient on the forward link [12][13].

S (1 bit)	E (1 bit)	LT (2 bits)	Length (12 bits)	Protocol Type (16 bits)	DATA

GSE header

Figure 2: GSE packet format

GSE appears not adapted to small PHY frames because of the large overhead that results when the fragmentation mechanism is used. The GSE header is at least 7 bytes for the first fragment (of a PDU) including the protocol type field and the total length field, at least 7 bytes for the last fragment including a CRC-32 used to ensure the correct reassembly, and at least 3 bytes for intermediate fragments including an identifier referring to the initial non fragmented PDU.

The GSE standard does not consider non QEF physical layers. In a non QEF physical layer, the GSE guidelines [11] suggest the use of CRC-32 located at the end of each PHY-frame. However no signaling is provided to indicate this and this size of CRC may be too large for (very) small PHY-frames when the link employs sophisticated error correction and may be redundant with the CRC-32 used for reassembly.

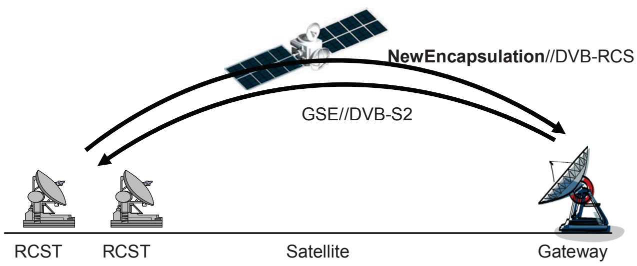

Figure 3: Forward and return link encapsulation

GSE presents several limits for small PHY frames, small PDUs, and non QEF links. A new encapsulation protocol is therefore required to obtain similar performance (and functions) for satellite waveforms using multiple access and random access (e.g. DVB-RCS or DVB-RCS-NG) as depicted in Figure 3.

III. REQUIREMENTS FOR A RETURN LINK ENCAPSULATION

This section derives requirements, in term of performance, flexibility, robustness, addressing, evolution and QoS support, which all should be considered when specifying a suitable encapsulation protocol for multiple access satellite links.

A. Performance

As resources in satellite networks are expensive and scarce, performance in terms of overhead and signaling are strong commercial requirements.

The overhead, that is the per packet signaling should be reduced to the minimum as PHY-frames are generally smaller than the frames sent on the forward link.

The (out of band) signaling required by the encapsulation protocol should be as light as possible. First, this allows to improve the performance in terms of the bandwidth consumed, but also in terms of the exchanges required to configure the stacks. The latter is an important issue in a context with appreciable delay. In addition to performance, a reduction of the signaling can increase interoperability between different implementations and configurations.

Therefore, to meet the performance requirements, the new encapsulation protocol should:

- Reduce the per packet signaling to the minimum.
- Not require out of band signaling (or at most minimal signaling).

These two requirements potentially conflict and a trade of may be needed.

B. Flexibility

The encapsulation protocol should permit the encapsulation of any protocol (e.g. IP, MPLS [14]) and should work over every intended physical layer. This introduces several issues:

First, most return link physical layers are not Quasi Error Free (QEF) (e.g. DVB-RCS, DVB-RCS+M). Therefore, an additional integrity check may be required to ensure a very low probability of transmissions of corrupted PDUs to higher layers. Two approaches are possible: either an additional check associated with the Sub-Network Data Unit (SNDU) that is, a part of the encapsulation protocol, or a check associated with the PHY frame to emulate a QEF physical layer [11]. A cross-layer mechanism is conceivable that performs a CRC assisted by feedback from the FEC decoder, possibly allowing optimization of the use of the CRC (both size and detection capacity). To ensure a simple encapsulation protocol, the physical layer should be QEF or at least be able to detect the errors with a very low probability of no detection.

The encapsulation protocol should then work with a range of PHY frame structures/sizes and be optimized for most of these. Typically, this must consider short PHY frames even if longer PHY frames are also conceivable, for example using a DVB-S2 bidirectional system. Modern PHYs can often support a range of burst sizes resulting from different combinations of coding and modulation (this is not specific to multiple access).

Because the GSE protocol is the encapsulation standard for the forward link, the new encapsulation protocol should also present the same interface to L3 as GSE – or at least a superset – that is any combination of {0,3,6 byte addresses; types; lengths of PDU and QoS priority} so that L3 can evolve to use the most appropriate method for a specific scenario.

Therefore, to meet the flexibility requirements, the new encapsulation protocol should:

- Permit the encapsulation of any protocols.
- Work over any physical layer (either QEF or a link enhanced to become QEF).
- Be efficient with any PHY frame size.
- Should work with deployed encapsulation protocol on the forward link (GSE).

C. *Robustness*

As a PHY layer can not be guaranteed to be error free, residual errors in headers are possible. The encapsulation therefore needs to be designed to be robust to errors, considering the state machine of the fragmentation/reassembly process and encapsulation/decapsulation process. It should also be robust to unintended reordering by the PHY, or at least this should not cause serious side-effects.

One possible solution is to consider adding a HEC (Header Error Correction) to protect the link header from corruption. This introduces an additional overhead. To avoid this overhead, the encapsulation protocol state machine could anticipate all possible errors and avoid unexpected behaviors. This makes the link quality transparent to the upper layers. The control and management plane will benefit from support for OAM (Operations And Maintenance) testing and management that enables the control plane to understand whether uncorrected errors have been found.

Therefore, to meet the robustness requirements, the new encapsulation protocol should:

- Be robust to any residual error in the header.
- Be robust to any loss.
- Be robust to potential reordering.

D. *Addressing*

On the forward DVB-S2 link, GSE is used for point to multipoint communication. Several different return link topologies are possible: multipoint to point (star) and multipoint to multipoint (mesh) communications. These must each be addressed.

In the case of multiple sources for the same destination, the destination needs to identify the L2 source to be able to reassemble received fragments of PDUs. A destination may have to reassemble several fragmented PDUs from several sources at the same time. A way to identify the source is to put explicitly the L2 address in the header of the SNDU. This has the advantage that the identification only relies on the header encapsulation protocol. A drawback is that this adds overhead. Another way to achieve this identification is to use the information of the lower layer (based on the reception interface). In the case of a DVB-RCS system, the destination may deduce the sender from the allocation table, that is the TBTP (Time Burst Table Plan). The advantage of this approach is that it does not involve additional overhead, however, this identification is dependent on the type of system and its use may impact product implementation and interoperability.

Both source and destination addresses may be required. For example the following section discusses use of the destination address for switching in the case of mesh communications.

VPN (Virtual Private Network) addressing may need to be supported to allow a satellite network to be divided into logical networks, although this may be handled by addressing modes and use of the type field. Another possibility is to use the IEEE802.1q [15] standard to tag the SNDU, as in Ethernet networks (or another type of label). The most appropriate method may depend on whether the virtual network extends into the connected terrestrial network (as when bridging Ethernet LANs) or terminates within the satellite terminal.

Therefore, to meet the addressing requirements, the new encapsulation protocol should:

- Provide addressing support for multipoint to point (star) and multipoint to multipoint (mesh) communications.
- Allow configuration of VPNs.

E. *Evolution*

The encapsulation protocol should be designed in a way that facilitates easy integration of potential evolution. A well-known way to support potential evolution is the use of extension headers, as in ULE and GSE. The new encapsulation protocol should allow such extension headers.

F. *QoS support*

Since return/mesh links are capacity-constrained they must support QoS and provide a form of pre-emption scheduling. Links operating at speeds of 2 Mbps or less generally benefit from design that supports priority queuing using a set of traffic classes. 4-8 queues should be sufficient, since this can support a wide range of higher-layer QoS behaviors.

Therefore, to meet the QoS support requirements, the new encapsulation protocol should:

- Provide support for priority queuing (e.g. a 3 bit QoS field, or equivalent method for pre-emption).

IV. MESH COMMUNICATIONS AND OBP SYSTEMS

This section introduces the concept of mesh communications. Both transparent and regenerative mesh satellite systems are considered. On one hand, transparent mesh relies on a transparent payload with time or frequency switching capabilities, such as a Digital Transparent Processor (DTP) to ensure layer 1 (L1) switch. On the other hand, On

Board Processing (OBP) allows to demodulate the signal and to perform both layer 2 (L2) or layer 3 (L3) on board switching. To illustrate the potential issues, the description focuses on systems that use DVB-S2/RCS, but the results may be generalized to other physical layers.

A. Introduction to mesh communications and OBP

Considering satellite systems, two types of communication are possible. One, using the satellite network as an access network allowing communications with the outside (e.g. Internet) as depicted in Figure 4.

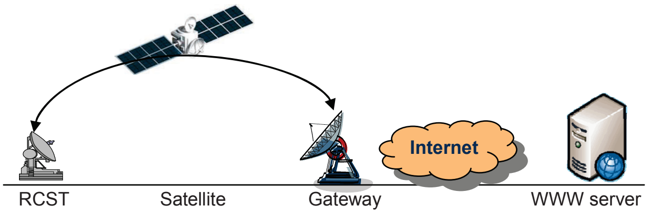

Figure 4: Communication outside the satellite network

Another use employs the satellite network for mesh communications that is for communications between the satellite terminals. Figure 5 depicts case the satellite network is seen as an independent network. The satellite payload allows On board Processing for the switching.

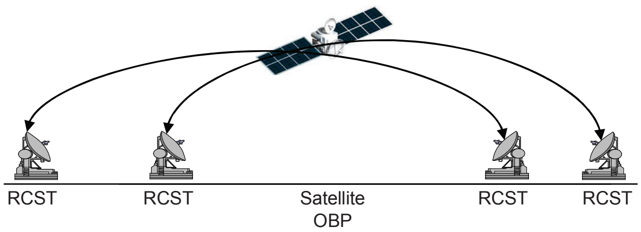

Figure 5: Mesh communications with OBP

The majority of the deployed satellite systems use transparent satellites because of the additional cost of OBP systems (and that OBP systems need to be designed to reflect a specific system). In this type of system, the satellite is fully transparent for upper layers and only frequency switching and burst switching are conceivable.

The following subsections examine the use of on board L1, L2 and L3 switching. OBP switching introduces complexity to the system. This complexity arises in several ways. First, the satellite needs to provide a forwarding method that moves packets from the input to the output. This may need to also provide support for reformatting the packet when the transmit and receive formats differ. Second, the forwarding method needs a way to select the appropriate output. Methods range from timeslot-based multiplexing, to label-based switching. Each has merits in terms of implementation cost/transmission efficiency/flexibility. Finally, the OBP system requires a control and management function.

B. On board L1 (layer 1) switching – Transparent satellite

The switching may be L1 (using burst switching or frequency switching) as depicted in Figure 6.

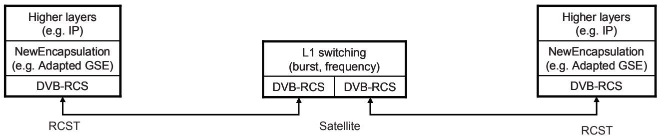

Figure 6: Protocol stacks for on board L1 switching

This architecture allows a lighter satellite payload compared to an OBP, however, as the satellite is transparent, it prevents using a different waveform on the downlink. Typically, the switching configuration is performed by management means (DTP) or through a dedicated link to transport labels (e.g. ULISS [16]) and thus does not impact the addressing requirements. The requirements are therefore those listed in section III.

C. On board L2 (layer 2) switching – OBP

A satellite that employs L2 OBP switching can employ a protocol stack that switches according to the L2 destination address. Figure 7 depicts the protocol stacks for mesh systems with on board L2 switching (e.g. Amheris).

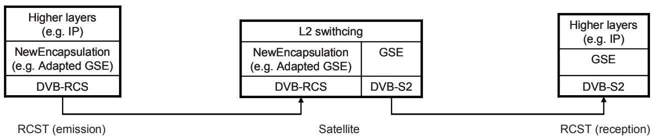

Figure 7: Protocol stacks for on board L2 switching

To switch packets, the satellite needs to identify the destination. The first solution (available in GSE) is to explicitly specify the destination L2 address (or a label associated to this or to a virtual channel) in the first fragment of each SNDU. In this case, the satellite has to reassemble the entire SNDU to switch it. To avoid this on board reassembly, the label may be attached to each fragment, but this then adds overhead (a significant impact if PHY frames are short). The destination address (or the label) can also be used for QoS filtering.

A L2 source identification is also required for reassembly (either on board or at the destination terminal) as detailed in the section III.D. It is possible that the addresses used for switching are associated with the setup of the switch (i.e. a tag or label that identifies the appropriate output port). This approach has similarities to ATM or MPLS.

As the encapsulation protocols used on the uplink and on the downlink need not be the same (and commonly will differ because the up-link is typically multiple access and the down link is multipoint), the satellite has to translate between the two protocols. GSE is expected for the multi-point down link. Therefore, the closer the new up link encapsulation is to GSE, the easier it will be to translate the protocols.

Therefore, to allow on board L2 switching in a mesh system, the new encapsulation protocol should:

- Specify a mechanism to identify the L2 destination.
- Take into account requirements in section III to identify the L2 destination (particularly the performance and flexibility requirements).
- Be as close as possible to the GSE protocol.

D. *On board L3 (layer 3) switching – OBP*

In a L3 design, the satellite switches according to the L3 destination address (e.g. IPv4 or IPv6 destination address). Figure 8 depicts the protocol stacks for mesh systems with on board L3 switching.

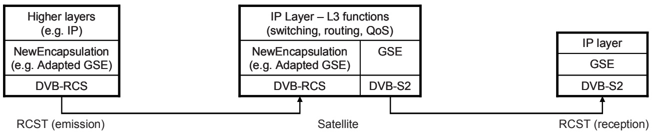

Figure 8: Protocol stacks for on board L3 switching

To perform layer 3 switching, the IP packet must be reassembled and decapsulated. The L2 destination address and the L2 source address are derived from the L3 information in the packet (depending on address resolution protocols as described in [17] for MPEG-2 networks).

The signaling required for the configuration of the encapsulation protocol should be as light as possible because management costs are an important component of the operational expenditure. Among others, this signaling protocol needs to perform address resolutions, routing, connection setup (including QoS and policy maps), resource allocation to the mesh. Therefore, the less complex is this signaling, the lower are the management costs and the better is the interoperability.

Unlike L2 switching, there is no need to translate uplink encapsulation into downlink encapsulation scheme (GSE) since the IP datagram is available at the output of the L3 forward engine and could be encapsulated on the downlink independently from the uplink encapsulation scheme.

Therefore, to allow on board L3 switching in mesh systems, the new encapsulation protocol should:

- Either specify or use signaling protocols for configuration (e.g. address resolution, QoS).
- Take into account deployed system to allow optimize interoperability.
- Not require out of band signaling (or at most minimal signaling).

V. FUTURE WORK AND CONCLUSION

This paper has derived the requirements for an encapsulation protocol in terms of the overhead and network resource, performance in terms of resources needed to execute, flexibility, addressing, robustness, evolution, QoS support, OBP systems, interoperability. These requirements may guide specification of an encapsulation protocol suited to return link and mesh satellite systems.

The main conclusions are that the new encapsulation protocol should:

- Be as light as possible in terms of overhead.
- Require no out of band signaling (or at most minimal signaling to ensure interoperability).
- Be as close as possible of the deployed encapsulation protocol on the forward/down link (that is GSE).

The next step of this work is the specification (e.g. a GSE-like solution) and the evaluation of an encapsulation protocol suitable for the return link and for mesh satellite systems. The brief survey of options available for OBP would suggest that it will not be easy to design a single optimal approach for all intended use of OBP.

ACKNOWLEDGMENT

The authors wish to acknowledge support of the European Space Agency (ESA) contract 22471/09/NL/AD, and the insight provided by many individuals contributing to the next generation of the DVB-RCS system specifications.

REFERENCES

[1] ETSI EN 300 421: “Digital Video Broadcasting (DVB); Framing Structure, Channel Coding and Modulation for 11/12 GHz Satellite Services”, 2008.

[2] ETSI 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”, 2008.

[3] ETSI 301 790, “Digital Video Broadcasting (DVB); Interaction Channel for Satellite Distribution Systems”, European Telecommunications Standards Institute (ETSI), 2009.

[4] ETSI EN 301 192, “Specifications for Data Broadcasting”, European Telecommunications Standards Institute (ETSI), 2004.

[5] Fairhurst, G. and Collini-Nocker, B., “Unidirectional Lightweight Encapsulation (ULE) for Transmission of IP Datagrams over an MPEG-2 Transport Stream (TS) ”, IETF, RFC 4326, 2005 (Standards Track)

[6] ETSI TS 102 606 “Digital Video Broadcasting (DVB); Generic Stream Encapsulation (GSE) Protocol”, European Telecommunication Standards, Institute (ETSI), Technical Specification, 2007.

[7] Montpetit, M.-J., Fairhurst, G., Clausen, H., Collini-Nocker, B., and H. Linder, “A Framework for Transmission of IP Datagrams over MPEG-2 Networks”, RFC 4259, November 2005.

[8] Fairhurst, G. and Collini-Nocker B., “Extension Formats for Unidirectional Lightweight Encapsulation (ULE) and the Generic Stream Encapsulation (GSE) ”, RFC 5163, April 2008.

[9] Borman, C., et al, “RObust Header Compression (ROHC): Framework and four profiles: RTP, UDP, ESP, and uncompressed”, RFC3095, 2001.

[10] Christian Prähauser; Bernhard Collini-Nocker “Experimental Evaluation of PDU Concatenation in ULE with IP Telephony”, IEEE International Workshop on Space and Satellite Communications, Salzburg, 2007.

[11] ETSI, “Digital Video Broadcasting (DVB); Generic Stream Encapsulation (GSE), Implementation Guidelines”, DVB BlueBook A134, February 2009.

[12] Cantillo, J., Collini-Nocker, B., De Bie, U., Del Rio, O., Fairhurst, G., Jahn, A., Rinaldo, R., “GSE: A Flexible, yet Efficient Encapsulation for IP over DVB-S2 Continuous Generic Streams”, International Journal of Satellite Communications and Networking, 2008.

[13] Mayer, A., Collini-Nocker, B., Vieira, F., Lei, J., Vasquez Castro, M. A., “Analytical and experimental IP encapsulation efficiency comparison of GSE, MPE and ULE over DVB-S2”, IEEE International Workshop on Space and Satellite Communications, Salzburg, 2007.

[14] Dubois, E., Baudoin, C., Haardt, C., Chaput, E., Beylot, A. –L., “IP/MPLS satellite network convergence in ULISS project”, International Workshop on Signal Processing for Space Communications, October 2008.

[15] IEEE Std. 802.1Q-2005, Virtual Bridged Local Area Networks.

[16] Haardt, C. and Courville N., “Internet by satellite: a flexible processor with radio burst switching,” in Satellite and Space Communications, 2006 International Workshop on, Sept. 2006, pp. 58–62.

[17] Fairhurst G., and Montpetit, M. –J., “Address Resolution Mechanisms for IP Datagrams over MPEG-2 Networks”, IETF RFC 4947, Jul. 2007.