ACCEPTED FROM OPEN CALL

DVB-S2 Full Cross-Layer Design for QoS Provision

M. Ángeles Vázquez Castro, Universitat Autònoma de Barcelona
Fausto Vieira, Instituto de Telecomunicações, and Universidade do Porto

ABSTRACT

The second generation of the Digital Video Broadcasting standard for Satellite transmission, DVB-S2, is the evolution of the highly successful DVB-S satellite distribution technology. DVB-S2 has benefited from the latest progress in channel coding and modulation such as Low Density Parity Check Codes and higher order constellations to achieve performance that approaches Shannon's theoretical limit. We present a cross-layer design for Quality-of-Service (QoS) provision of interactive services, which is not specified in the standard. Our cross-layer approach exploits the satellite channel characteristics of space-time correlation via a cross-layer queueing architecture and an adaptive cross-layer scheduling policy. We show that our approach not only allows system load control but also rate adaptation to channel conditions and traffic demands on the coverage area. We also present the extension of our cross-layer design for mobile gateways focusing on the railway scenario. We illustrate the trade-off between system-wide and individual throughput by means of simulation, and that this trade-off could be a key metric in measuring the service level of DVB-S2 Broadband Service.

INTRODUCTION

The second generation of the Digital Video Broadcasting standard for Satellite transmission, DVB-S2, was approved in 2004 [1]. In this article, we propose a design of the DVB-S2 optional interface for interactive services, which is left open in the standard. Our solution is based on a full cross-layer design. Furthermore, our design allows for a straightforward extension to support reliable transmission in mobile scenarios. Specifically, we present the case of providing Internet to trains.

Compared to the original DVB-S, the second generation brings fundamental changes related to both physical and access layers promising 30 per cent bandwidth efficiency increase over existing systems. A set of coding rates and modulation formats are provided with performances close to Shannon capacity via sophisticated modulation techniques and Low-Density Parity-Check error correction codes (LDPC). The physical layer is no longer fixed by enabling Variable Coding and Modulation (VCM) and Adaptive Coding and Modulation (ACM). This adaptation requires a return channel to feedback Channel State Information (CSI). The Return Channel Satellite (DVB-RCS) standard [2] provides a satellite-only solution, although hybrid solutions with terrestrial return channels are also specified for DVB-S2 [15]. The CSI enables transmission at the most efficient coding and modulation scheme for the instantaneous Signal-to-Noise-plus-Interference-Ratio (SNIR). This SNIR depends on channel conditions, antenna size and the satellite terminal's location within the satellite's coverage area. The access layer enables native unicast services, which is an evolution from the general broadcast scope of DVB standards. These new capabilities enable satellite networks to support a wider range of IP applications and services.

Nevertheless, the actual adoption of DVB-S2 has faced important challenges that can be summarized as follows:

- Traditional design of DVB-S point-to-point satellite systems is based on link set-up for the worst-case propagation and location conditions. This design methodology is not optimal for DVB-S2 physical layer adaptation and in particular for unicast networks enabled with ACM to best match the user SNIR.
- Designs based on layered protocols hinder joint optimization of different communications layers, thus precluding upper layers from exploiting physical layer adaptation.
- The DVB protocol stack was originally designed to support digital television transmission, including video, audio and Service Information (SI). These are carried on an MPEG2 Transport Stream (TS) that multiplexes the different streams. Later on, encapsulation protocols added basic support for Internet Protocol services. This approach aggregates all data into a single data pipe, not allowing differentiation of unicast traffic.

The three points above point to the fact that a full design of a DVB-S2 system requires a novel methodology based on adaptation and cross-layer optimization. It should also be pointed out that the current version of the DVB-S2 standard only supports transmission to fixed ter-

0163-6804/12/$25.00 © 2012 IEEE

minals. Hence, additional techniques on top of the DVB-S2 physical layer are necessary to support mobility.

This article presents a novel methodology for full cross-layer design of satellite systems using adaptive physical layer. The originality respect to other works [3–5] can be summarized as follows:

- Concrete definition of fairness and design of adaptive/tunable scheduling policies for fairness enforcement. We have addressed this problem by providing a concrete solution that is "tunable" thus allowing the system designer to have full control of the fairness to meet the required service requirements of quality at network level and, as a consequence, of the system-level load and user-level throughput.
- Structured definition of a general cross-layer methodology for the queuing architecture and scheduling policies design, based on a modular approach that admits different policies, and upgrades if needed. Specifically, we prove the scalability and flexibility of our design by showing the easy extension to new encapsulation schemes and communications scenarios (in particular, from a fixed scenario to a mobile scenario).
- Architectural cross-layer design compliant with state-of-the-art standards for broadband satellite systems through a clear definition of satellite-dependent and satellite-independent blocks that admits a neat mapping to standardized modules and signaling. This further allows easy integration with terrestrial networks and hybrid wireless networking.
- Presentation of key results showing how our design allows full control of the system-level load and user-level throughput independently of the channel conditions throughout the entire coverage region.

The remainder of this article is organized as follows. We present the time-space channel model. We describe the proposed cross-layer design of the architectures for both fixed and mobile gateways. We present a practical implementation compliant with the ETSI BSM (Broadband Satellite Multimedia) Satellite Independent — Service Access Point (SI-SAP) model. We provide simulation results to illustrate the performance of our scheme and present the conclusions.

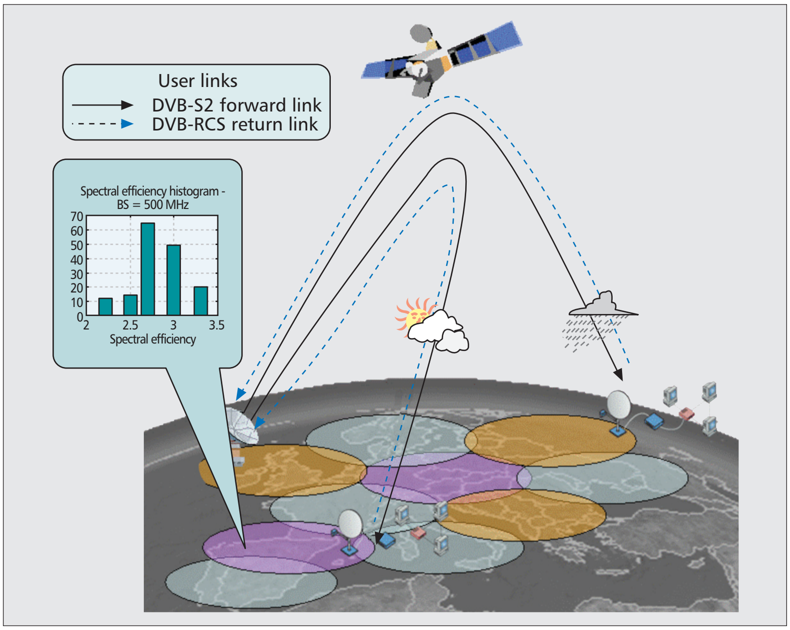

Figure 1. *System model of a transparent satellite in a star topology.*

DVB-S2 SYSTEM AND CHANNEL MODELS

The reference DVB-S2 system scenario considered is a broadband satellite system in the Ka-band (30/20GHz) with a multi-beam architecture [6]. The system aims at providing IP services ranging from real-time multimedia content (VoIP, video-conferencing), to best-effort services (e.g., web browsing, e-mail, data transfers, …). The system is composed of a number of gateways, delivering service to user terminals distributed over different spot-beams. A gateway is an earth station that bridges the terrestrial and satellite networks, namely providing Internet access and other networking services to satellite terminals. The satellite is assumed transparent in a star topology. As a typical commercial system, two types of users are considered: corporate and mass-market customers. We focus on a single Time Division Multiplexed (TDM) carrier of the forward downlink (i.e., from the satellite to the user terminal) transmitted over one beam implementing DVB-S2. This model is depicted in Fig. 1 together with an example of long-term values of DVB-S2 spectral efficiencies for average weather statistics, represented via a histogram. Within the beam coverage, we consider *N* fixed users experiencing location-based time-variant channel conditions. We assume that end-user terminals send their CSI to the transmitter, and physical layer adaptation is performed on a DVB-S2 physical layer frame-by-frame basis. We assume a constant spectral bandwidth and constant transmitted power per beam. Hence, an individual user receives the total beam power in the scheduled TDM frame. As a consequence, only bit rate adaptation is performed and not symbol rate adaptation.

The Ka-band channel attenuation is mostly conditioned by the effects of rain events, with attenuations ranging from a few dBs up to more than 20 dB. However, channel variations are extremely slow (1 dB per half a second maximum). A key aspect of this channel is therefore that the SNIR variations throughout the satellite coverage are highly correlated both in time and space dimensions. This is clarified in Fig. 2, which presents a simulation of a cell rain over a given satellite beam coverage. All users under the rain cell can be assumed to have similar values of channel attenuation, i.e., they are correlated both in time and space. We model this channel by defining a number of "Correlated Areas" (CA), where users will have similar channel conditions [6]. Such correlated areas are not fixed but move within the coverage area at the wind speed. Hence, logic traffic aggregation is possible within a CA, since all users have similar channel conditions and thus will undergo similar rate adaptation at the transmission. This physical

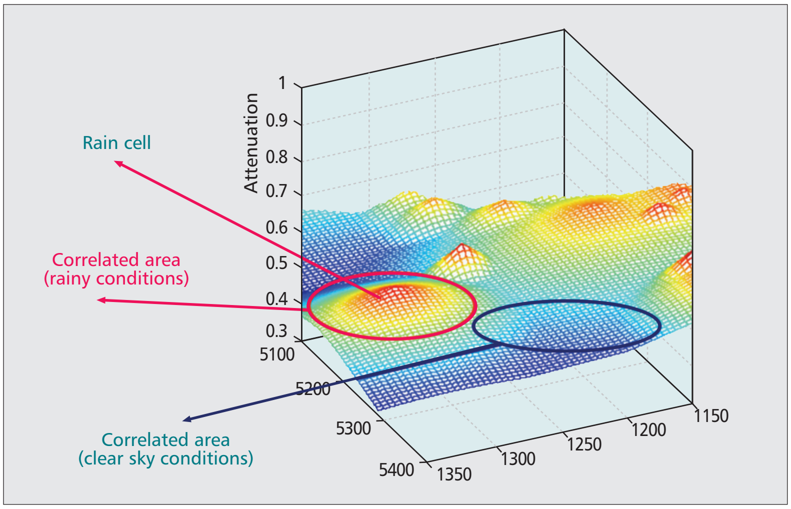

Figure 2. *Correlated area concept.*

modeling is used to drive the cross-layer design and scheduling policies. We explain our cross-layer design and scheduling policies later.

DVB-S2 CROSS-LAYER DESIGN

In this section, we present a cross-layer design methodology for DVB-S2 as an alternative to a traditional satellite system's design methodologies. Later, this methodology is applied to the fixed and mobile gateways scenarios.

The fundamental objective of our cross-layer optimization is to exploit physical layer adaptation as much as possible across layers. In particular, our methodology aims to exploit physical layer adaptation in order to support the maximum stable system load while guaranteeing Quality of Service (QoS) requirements, in terms of network metrics such as bandwidth, delay, jitter and packet loss. As for the QoS model, we assume an IETF Differentiated Services (Diffserv) model at IP level with three Classes-of-Service (CoS): Best Effort (BE), Assured Forward (AF), and Expedited Forward (EF). The CoS concept is at the basis of the Diffserv model, where packets are classified into a limited number of traffic classes, which can be used to provide service differentiation between multiple flows aggregated into these classes. With this assumption, we have developed the cross-layer architecture shown in Fig. 3. We also define QoS in terms of the following parameters per CoS:

- Delay constraints (T_{BE}, T_{AF}, T_{EF}).
- Jitter constrains (J_{BE}, J_{AF}, J_{EF}).
- Fairness among users under different channel conditions.

The methodology presented here has the following three essential steps:

Define a cross-layer queuing architecture. Such architecture should allow adaptation across layers. In particular, our architecture will consist of two levels of queuing: IP level and a joint PHY/MAC level.

Obtain adaptive scheduling policies that enforce QoS. In particular, we propose a time-variant adaptation of the scheduling parameters to the dynamics of both system load and channel conditions. System load is kept within the maximum stable system capacity. QoS is enforced by scheduling the traffic within each Class-of-Service (CoS) followed by a per-CoS scheduling for transmission.

Practical implementation. Several practical cross-layer implementations are possible. Our preferred proposal is to enclose a shared database and a coordinator of the cross-layer functionalities. In addition, the design should be compliant with state-of-the art implementation standards.

CROSS-LAYER QUEUING ARCHITECTURE FOR FIXED TERMINALS

Our cross-layer queuing design takes advantage of the time-space characteristics of the channel (explained earlier). The overall cross-layer architecture is shown in Fig. 3 and is composed of two queuing modules and three scheduling modules. As the IP packets enter the system, they are identified and sent to the *IP unicast flow aggregator* according to CA and CoS. The *IP scheduler* takes the IP packets from the queues and into the *Cross-layer DVB encapsulator*. Within each CoS, the encapsulation is performed according to the DVB-S2 coding and modulation combination (MODCOD) supported by current SNIR of the destination terminal. The encapsulated packets are then sent to the physical layer buffers by the *Cross-layer DVB hierarchical scheduler*, which implements a Tunable-Fairness Weighted Round Robin (TF-WRR) at a CoS level, followed by an Adaptive Weighted Round Robin (AWRR) between the different CoS. The set of bits buffered at the corresponding physical layer (MODCOD) buffers are directed to the corresponding MODCOD interface of the DVB-S2 modem.

Following, each module and policy are both described and justified in detail.

IP Unicast Flow Aggregator — This module aggregates traffic in two dimensions, CA and CoS, with a corresponding number of IP queues. As each IP packet arrives either from the local network with Diffserv tagged traffic or from the Internet (classified as BE), they are directed by a packet classifier to the corresponding CoS according to the Differentiated Services Code Point (DSCP). The DSCP allows routers to apply a different set of rules and algorithms to packets. The packet classifier also identifies the CA to which the destination terminal belongs in order to discriminate between different weather conditions within the beam. Each input packet is queued into a First-In-First-Out (FIFO) queue that corresponds to the CoS and the CA of the packet. This way, the traffic within each queue will experience similar service rates, i.e., similar QoS characteristics in terms of network metrics. Note that the granularity of rate adaptation (i.e., number of CAs implemented) is a system design parameter. The queued packets are no longer identified by the flow they belong to within the satellite sub-network, since they are aggregated by QoS and channel conditions. This cross-layer flow aggregation is fundamental for handling the

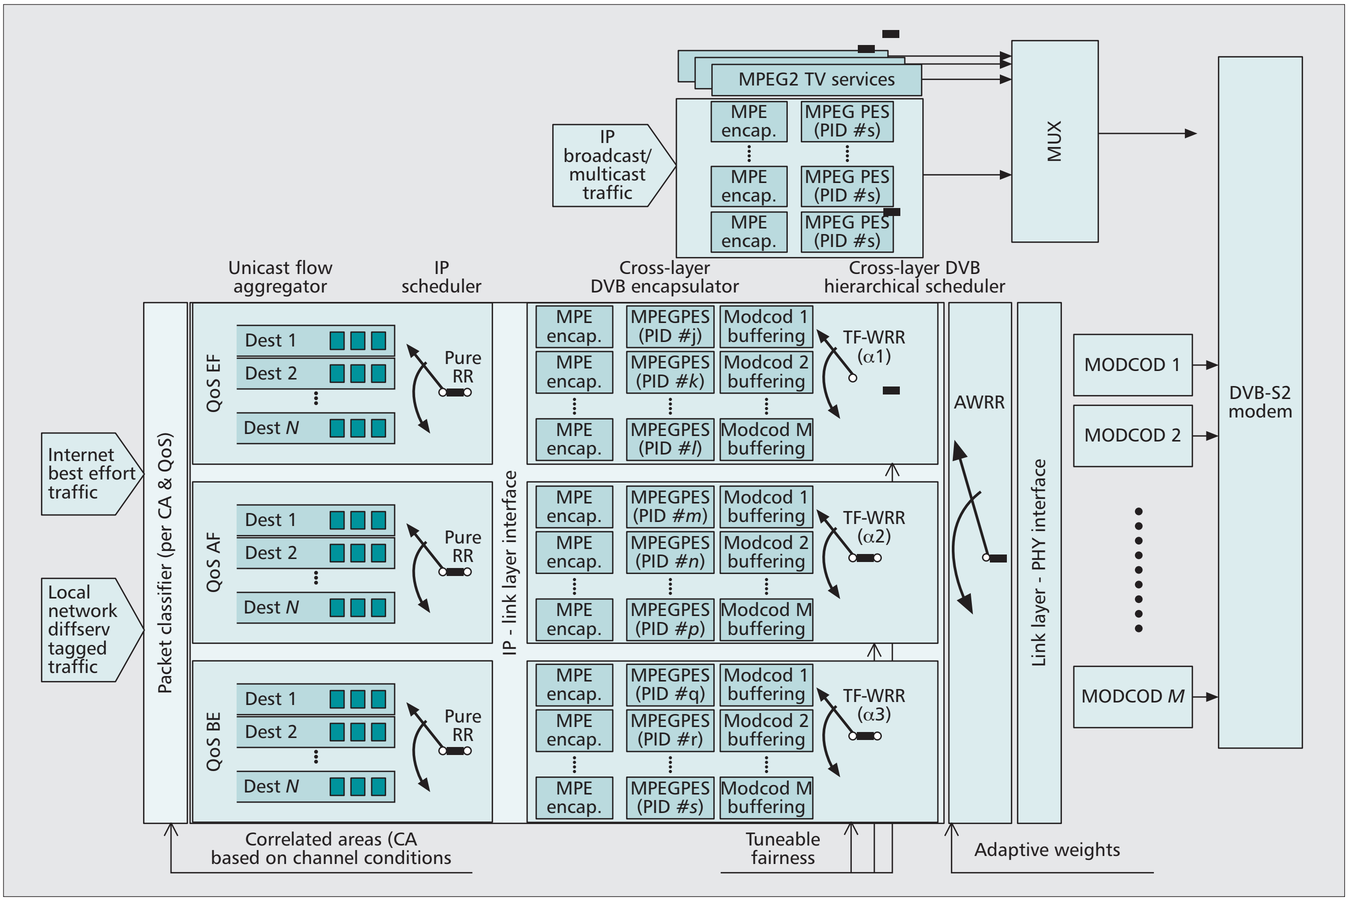

Figure 3. *QoS enabled DVB-S2/RCS cross-layer architecture.*

high data rates expected in a broadband satellite system, on the order of thousands of simultaneous active users and total throughput of GB/s.

IP Packet Scheduler — Once the IP packets are queued per CoS and CA, an IP packet scheduler sends the packets down to the DVB layers. This scheduler has been made intentionally simple and fully driven by the lower layers, where the cross-layer adaptation takes place. Specifically, the only role of this scheduler is to send packets in a pure Round Robin (RR) fashion to the DVB buffers in the cross-layer DVB encapsulator module at the pace dictated by back pressure of the cross-layer DVB scheduling.

Cross-Layer DVB Encapsulation — This module encapsulates the IP packets coming from the IP scheduler. This is a cross-layer configuration that allows QoS joint optimization at layers 2 and 3 by defining queues in terms of CoS and physical layer MODCODs. This process takes place at the DVB layer and consists of encapsulating the IP packets into MPEG-TS packets. The DVB-S2 standard provides up to 28 different MODCODs, although manufacturers may implement a subset only. IP packets are encapsulated into MPEG-TS packets using the Multiprotocol Protocol Encapsulation (MPE). MPEG-TS packets convey different Packetized Elementary Streams (PES), each having a different Packet ID (PID), and it is up to the designer to define one or more PES per MODCOD, for fine-grained control at a lower level. Herein below we describe how latest encapsulation schemes can be also adopted.

Cross-Layer DVB Hierarchical Scheduler — This scheduling module employs a hierarchical scheduling policy that exploits the cross-layer configuration of having IP packets queued in terms of both CoS and MODCOD. The hierarchical scheduling consists of the following policies.

Tuneable-Fairness Weighted Round Robin (TF-WRR) — This policy operates on a per-MODCOD basis within each CoS. It provides fairness among MODCODs as follows. Define the vector X = (x1, x2, …, xM) as the vector of weights assigned to each MODCOD within a given CoS. Weights X are defined as relative frequency the buffer is given access so that $\sum_{m=1}^{M} x_{m=1}$ We propose a tunable-fairness parameter α to derive the vector X such as fairness can be tuned across flows transmitted through the different MODCODs within a CoS. The derivation of the analytical expression for the weights per-CoS and per-MODCOD can be found in [7] and for the sake of completeness is presented here:

$$x_m(\alpha_{CoS}) = \frac{p(m,\alpha_{CoS})}{\sum_{i=1}^{M} p(i,\alpha_{CoS})}, m = 1,2,\ldots M, \quad CoS = BE, AF, EF \qquad (1)$$

where $p(m, \alpha_{CoS}) = 1/T_m, \eta_m^{\alpha\mathrm{CoS}}$, T_m and η_m are the transmission time and spectral efficiency respectively of MODCOD m. In [7] it is proved that by varying the α_{CoS} parameter, the fairness achieved among users under different channel conditions within each CoS is tunable, thus providing full control over the system throughput.

Layering in IP networks hides the complexity of adjacent layers by only exposing very simple interfaces. This provides a fundamental basis of Internet success. Therefore, we propose that cross-layer visibility must be provided as a service, with well-defined interfaces for populating external databases and querying the information.

Adaptive Weighted Round Robin (AWRR) — The first policy provides the preferred fairness, while this second scheduling policy provides the QoS guarantees. Let us consider the weights for the required Diffserv policy such as qBE + qAF + qEF = 1. Then in our case, the actual per-CoS and per-MODCOD weights are given by

$$w_{m,CoS}\,(\alpha_{CoS}) = q_{CoS}\,x_m\,(\alpha_{CoS}),$$
$$m = 1, 2, \ldots M,\ CoS = BE, AF, EF \qquad (2)$$

Note that Diffserv weights used on the Internet are usually fixed since wireline capacities are fixed. However, the adaptation of the qCoS weights is necessary in our scenario in order to follow the capacity dynamics due to physical layer adaptation. Note that the weights can be non-adaptive in order to reduce complexity, but at the cost of operating the system at a lower load.

Note that a hierarchical implementation of WRR and AWRR is assumed in Fig. 3. Naturally, a completely equivalent implementation is a single scheduler with a non-hierarchical structure and with weights given by Eq. 3.

Before deriving the AWRR adaptation algorithm, we first note that the steady-state value of the qCoS weights should be set to the value that corresponds to the percentage of each CoS traffic, i.e., if the percentage of the total traffic assigned to each CoS can be, for example, 10 percent (EF), 40 percent (AF) and 50 percent (BE), then the steady-state value of the Diffserv weights would be (0.1, 0.4, 0.5). It is well known that Longest-Queue-Highest-Possible-Rate (LQHPR) allocation strategy minimizes the average packet delay. An optimal policy needs to operate at the edge of stability, and knowledge of both arrival rates and current available physical-layer capacity allows it to verify whether that is the case. However, neither the arrival rate nor the current available capacity information can be obtained by any cross-layer signaling since only information on single-user physical layer capacity is available at the physical layer. Nevertheless, this information can be obtained from the queue lengths to compensate queue unbalancing produced by load and capacity variations among the three CoS.

We propose a scheme where as in [8] the basic idea is that an increase of first-class service's (EF) weight by first shifting the weight from best effort to EF and, if this is not enough, part of AF's weight will also be shifted to premium service. However, once the system load decreases, the nominal weights will be reinstated. Conversely to [8], we propose a logarithmic adaptation instead of linear. Note that this adaptive scheduling policy at layer 2 is clearly cross-layer since it requires knowledge of the IP queues occupation at layer 3.

CROSS-LAYER ARCHITECTURE FOR MOBILE GATEWAYS

In this section, we demonstrate the straightforward extension of our cross-layer architecture to the transmission to mobile gateways serving multiple end-users. Specifically, we focus on a potential scenario for DVB-S2: mobile gateways installed on trains.

In a railway scenario, short-term fading is superimposed to the channel described earlier. In particular, the presence of several metallic obstacles along the railroad such as electrical trellises (power arches) and catenaries typically result in an attenuation of 2–3 dB [9]. Additionally, electrical trellises (usually every 50 meters in Europe) result in deeper attenuations that, depending on the geometry and layout of the obstacles and on the orientation of the railway with respect to the position of the satellite, can even reach 15–20 dB. Therefore, this attenuation is very deep but deterministic and has a very short duration. Therefore, it can be modeled as an ON/OFF channel with a cycle that varies with the train speed. These effects were not contemplated in the DVB-S2 standard. In [9] it is shown that using Forward Error Correction (FEC) at the packet level (thus allowing a longer encoding period over several packets) is enough to provide the required performance against railway channel impairments. FEC was first introduced in the DVB family of standards within DVB-H, where a Reed-Solomon code is applied at the IP level, right before MPE encapsulation. However, DVB-H is a first-generation DVB standard and as such only pure broadcasting transmission is assumed. Current versions allow for different encapsulation schemes and require additional signaling. From an architectural perspective, it is straightforward to incorporate such additional signaling in our architecture as part of the Cross-layer DVB encapsulator, as shown in Fig. 4.

CROSS-LAYER IMPLEMENTATION

Layering in IP networks hides the complexity of adjacent layers by only exposing very simple interfaces. This provides a fundamental basis of Internet success. Therefore, we propose that cross-layer visibility must be provided as a service, with well-defined interfaces for populating external databases and querying the information.

In addition, this approach also allows many different means for providing the cross-layer information. Although the terminal and/or network elements themselves may generate the cross-layer parameters, information could also come from separate measurement devices or even human operators (as in our case, the preferred fairness policy). This approach preserves the inherent diversity across layers and the natural evolution of both layers and techniques for collecting the data.

Further, our proposed architecture is compliant with the ETSI BSM SI-SAP interface since satellite dependent and independent modules are kept separated [10]. The ETSI BSM architecture is characterized by the separation between common Satellite-Independent (SI) protocol layers and alternative lower Satellite Dependent (SD) layers.

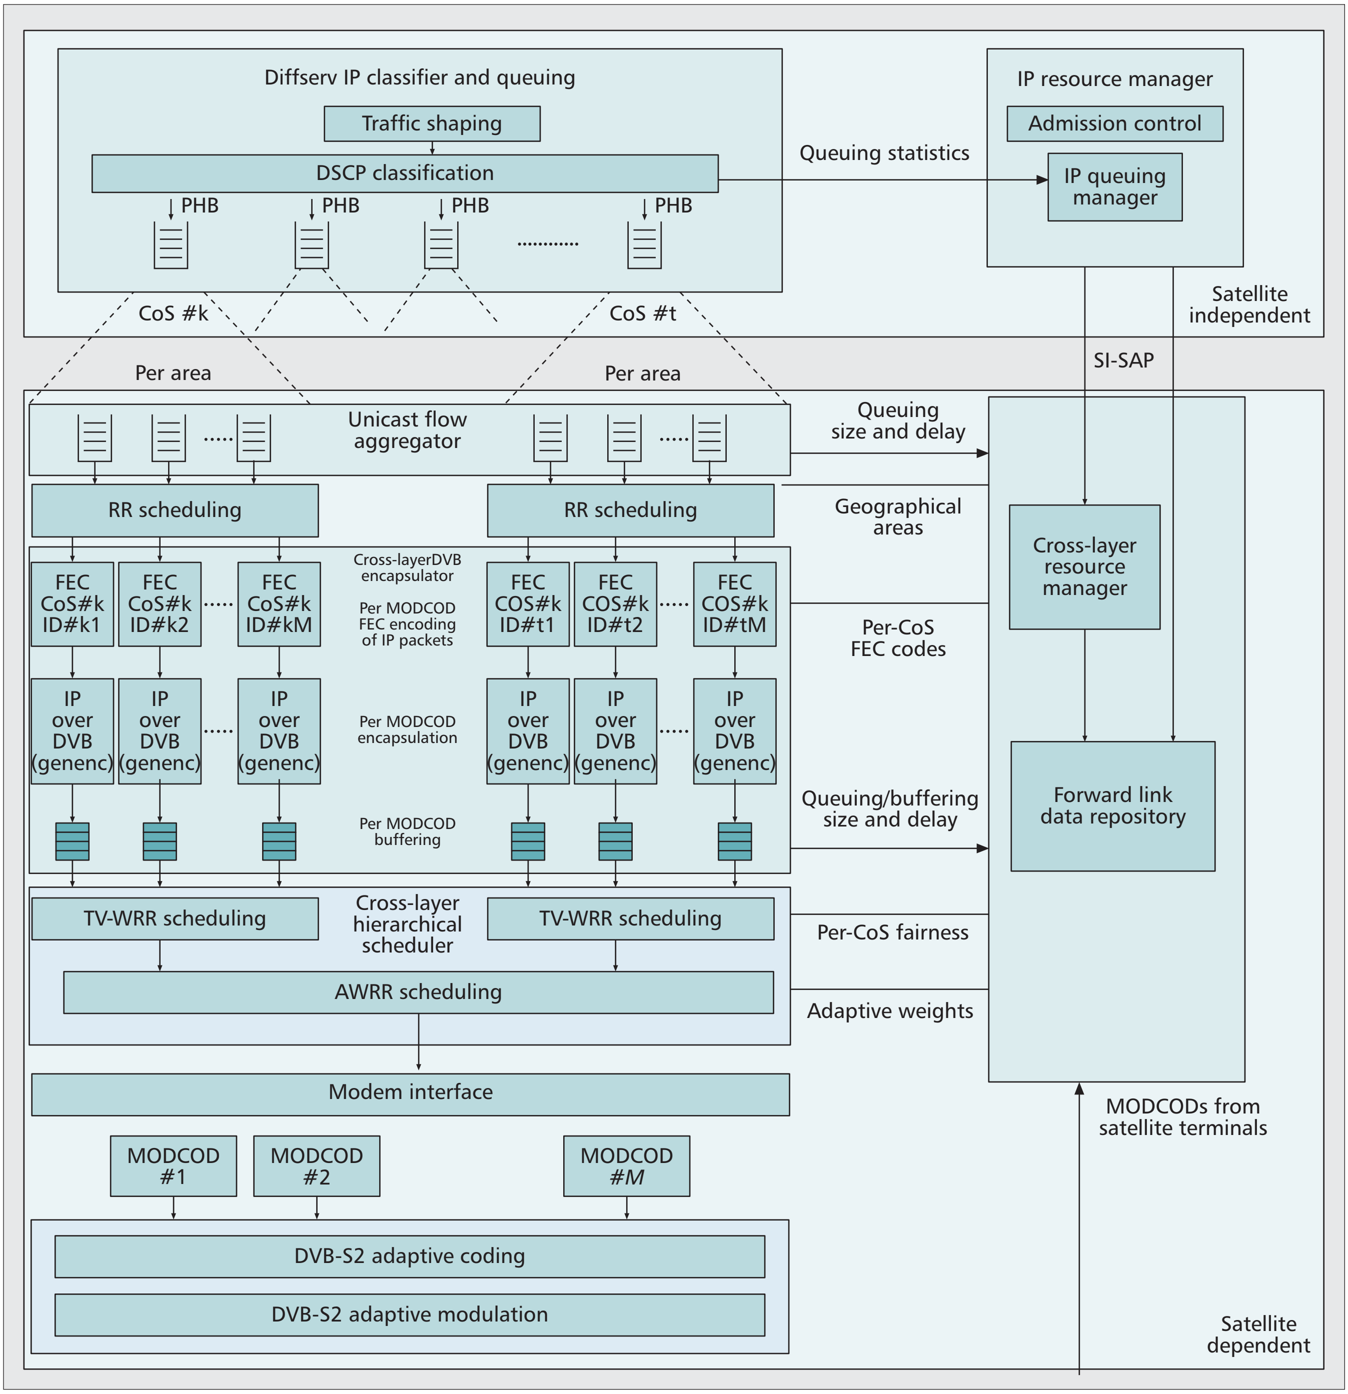

Figure 4. *DVB-S2/RCS architecture implemented in the ETSI BSM SI-SAP model.*

As shown in Fig. 4, our architecture does not make further assumptions on the satellite independent part of the protocol stack. Rather, we assume a fully complaint Diffserv implementation. For the satellite-dependent part we propose a cross-layer resource manager and a shared memory with interfaces to all lower layers via suitable primitives. This approach decouples information collection from information use and hence can adopt an implementation compatible with the protocol stack organization proposed by the ETSI BSM SI-SAP interface. Other aspects of our proposed implementation are:

- Circular buffers for the per-MODCOD packet encapsulation so packets can be continuously buffered between transmissions. This option is inline with the simplicity required to target high broadband capacity.
- Alternative encapsulation protocols such as Generic Transport Stream (GSE), which is compliant with DVB-S2 standard.
- The following IP queues statistics are needed: Average Queues size $\overline{Q}(t)$, Average Queues Delay $\overline{D}(t)$, and per-queue estimated throughput rat $\overline{Q}(t)/\overline{D}(t)$.
- The Forward Data Link Repository can be actually identified as the Dedicated Policy Repository of the policy management architec-

ture the IETF Policy Framework (POLICY) Working Group has developed, which is considered the best approach for policy management on the Internet. It can be described as the place to store and retrieve information.

- The Cross-layer Queuing Manager accesses the data repository to obtain queue statistics and takes decisions on hierarchical scheduler weights. Both fairness and CoS weights adaptation are key issues of packet scheduling, and a detailed analysis is presented in the following section. Fairness policies can also be modified by human operators.

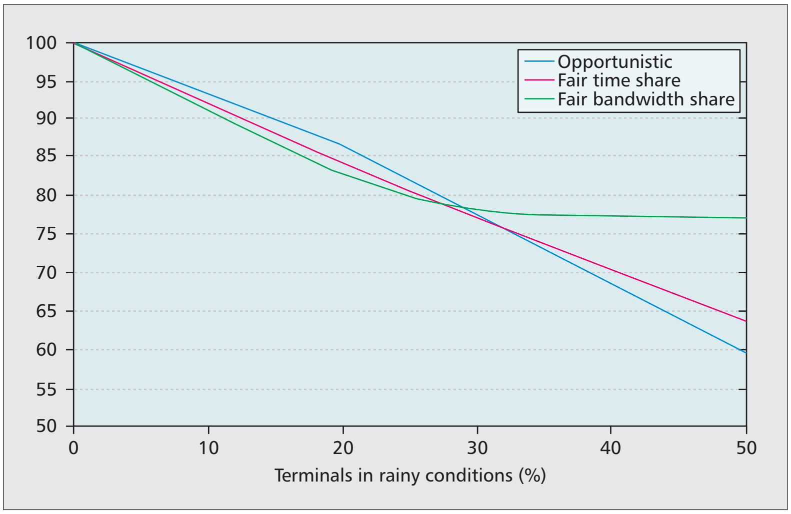

Figure 5. *Maximum stable system load with QoS guarantees for different fairness policies.*

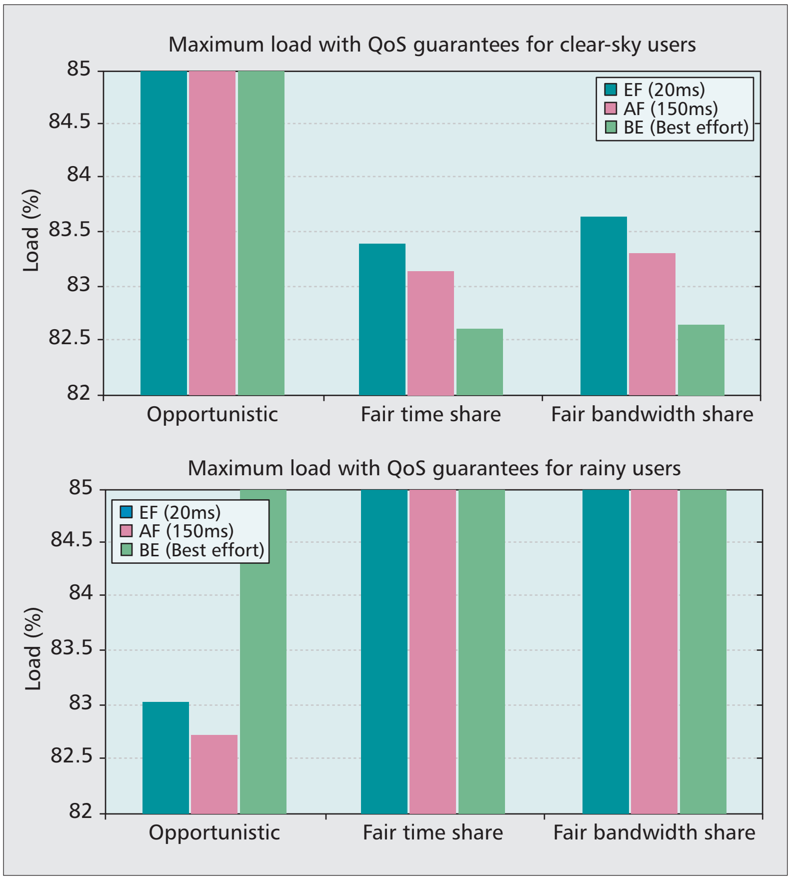

Figure 6. *Effect of fairness policies with QoS guarantees at the user-level (stable load is 85 percent).*

SIMULATION EVALUATIONS

A large number of event-driven simulations and realistic emulations have been carried out and have consistently validated the performance of the architecture for any channel conditions and population density within the coverage area. The study case analyzed considers a broadband system of 500 MHz bandwidth and 70 beams covering Europe. We show two illustrative results highlighting the effect of different fairness policies in terms of overall delay and throughput. Figure 5 shows the maximum stable load for different percentages of users under rainy conditions (maximum rain attenuation of 7 dB). It is observed how opportunistic policies can increase the overall throughput by granting more bandwidth to the highest spectral efficiencies. However, in order to guarantee the QoS constraints, this is only valid when the percentage of users under rainy conditions is small. The other two policies, equal bandwidth or equal time regardless of the spectral efficiencies, are more suited for very few occasions where there is a high percentage of users under rainy conditions. We show that a trade-off exists between system-wide and individual throughput and different objectives can be reached by tuning this trade-off appropriately.

Figure 6 presents the case of a non-uniform population density with 20 percent of users in rainy conditions corresponding to a theoretical maximum system load of 85 percent. Note that if the load increases above the stable limit, the Cross-layer Queuing manager will adapt the QoS weights and the excess load with be diverted to BE. Traffic may eventually be discarded in the BE queues since delay bounds are not guaranteed above the maximum system load.

CONCLUSIONS

We have presented a full cross-layer design of a DVB-S2 based satellite system for unicast services providing QoS guarantees. Both methodology and design of all functional elements are described. Our approach exploits the satellite channel characteristics of time-and-space correlation via a cross-layer queuing architecture and an adaptive scheduling policy that supports different fairness policies. Furthermore, we propose a practical implementation compliant with the ETSI BSM SI-SAP model. Our results show our tunable-fairness solution allows not only load system control but also system throughput adaptation. Moreover, our results prove that a trade-off exists between system-wide and individual throughput, which may be quite relevant when designing a suitable DVB-S2 Broadband Satellite Multimedia service.

ACKNOWLEDGEMENTS

The work reported in this article was supported by the European Space Agency contract "IP-friendly cross-layer optimization of adaptive satellite systems," http://telecom.esa.int/telecom/www/object/index.cfm?fobjectid=27706 .

REFERENCES

[1] 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," ETSI European Standard, v.1.1. March 2005.

[2] ETSI EN 301 790, "Digital Video Broadcasting (DVB); Interaction Channel for Satellite Distribution Systems," ETSI European Standard, 2000, v. 1.3.1, Mar. 2003.

[3] E. Chaput, A.-L. Beylot, and C. Baudoin, "Packet Scheduling Over DVB-S2 through GSE Encapsulation," Proc. IEEE Global Telecommun. Conf., GLOBECOM 2008, (New Orleans, US).

[4] G. Fairhurst *et al.*, "Multimedia Traffic Scheduling in DVB-S2 Networks," *IEEE Int'l. Wksp. Satellite and Space Comm. (IWSSC'08)*, Toulouse, France, Oct. 2008.

[5] B. Collini-Nocker and T. Soboll, "Class Based Scheduling for DVB-S2," *IEEE Int'l. Wksp. Satellite and Space Comm. (IWSSC'09),* Siena, Italy, Sept. 2009.

[6] M. A. Vázquez Castro and G. Seco Granados, "Cross-Layer Packet Scheduler Design of a Multibeam Broadband Satellite System with Adaptive Coding and Modulation," *IEEE Trans. Wireless Commun.*, vol. 6, no. 1, 2007, pp. 248–58.

[7] F. Vieira, M. A. Vázquez Castro, and G. Seco Granados, "A Tunable-Fairness Cross-Layer Scheduler for DVB-S2," *Int'l. J. Satellite Commun. and Networks*, vol. 24, no. 5, 2006, pp. 437–50.

[8] H. Wang, C. Shen, and K. G. Shin, "Adaptive-Weighted Packet Scheduling for Premium Service," *IEEE Int'l. Conf. Commun.*, Helsinki, Finland, June 2001.

[9] S. Cioni *et al.*, "Advanced Fade Countermeasures for DVB-S2 Systems in Railway Scenarios," *EURASIP J. Wireless Commun. and Networking*, vol. 2007, no. 2, pp. 15–25

[10] ETSI TS 102 357 "Satellite Earth Stations and Systems (SES); Broadband Satellite Multimedia; Common Air Interface Specification: Satellite Independent Service Access Point (SISAP)"

BIOGRAPHIES

MARIA ÁNGELES VAZQUEZ CASTRO [SM'08] received the Telecommunication Engineer degree (1994) and Ph.D. (cum laude, 1998) both from the Polytechnic University Vigo (Spain). She is an Associate professor at the Universitat Autònoma de Barcelona (Spain). She has been a Research Fellow at the European Space Agency (2002–2004) and at the University of Southern California (2000). Her research group belongs to the European Network of Excellence on Satellite Communications. She has co-authorized around 100 papers, holds one patent and contributes to the standardization bodies ITU and DVB. Her current research interests are cross-layer optimization and coding for interference management.

FAUSTO VIEIRA received his degree in Electrical and Computer Engineering from the University of Porto (2001). He joined the Telecom Section of the European Space Agency (ESA), in 2003. He became a researcher at the Universitat Autònoma de Barcelona (UAB), in 2004, contributing to different projects EU and ESA projects. In July 2008, he obtained his Ph.D. in Telematics Engineering from the Universitat Politècnica de Catalunya (UPC) and in September 2008, he became a post-doctoral researcher, and later on a Research Fellow at Instituto de Telecomunicações (IT) in Porto.

Our approach exploits the satellite channel characteristics of time-and-space correlation via a cross-layer queuing architecture and an adaptive scheduling policy that supports different fairness policies. Furthermore, we propose a practical implementation compliant with the ETSI BSM SI-SAP model.