

Research Article

Capacity Versus Bit Error Rate Trade-Off in the DVB-S2 Forward Link

Matteo Beriola, Christian Kissling, and Rémi Lapeyre

German Aerospace Center (DLR), Institute of Communications and Navigation, Oberpfaffenhofen, 82234 Wessling, Germany

Received 5 October 2006; Accepted 12 March 2007

Recommended by Ray E. Sheriff

The paper presents an approach to optimize the use of satellite capacity in DVB-S2 forward links. By reducing the so-called safety margins, in the adaptive coding and modulation technique, it is possible to increase the spectral efficiency at expenses of an increased BER on the transmission. The work shows how a system can be tuned to operate at different degrees of this trade-off, and also the performance which can be achieved in terms of BER/PER, spectral efficiency, and interarrival, duration, strength of the error bursts. The paper also describes how a Markov chain can be used to model the ModCod transitions in a DVB-S2 system, and it presents results for the calculation of the transition probabilities in two cases.

Copyright © 2007 Matteo Beriola et al. This is an open access article distributed under the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

1. INTRODUCTION

The original DVB-S standard dates back to 1995 and was intended for delivery of broadcasting services, the underlying transport stream of DVB-S was defined to be MPEG-2. DVB-S2 [1] is the second generation of the DVB-S standard and comprises a variety of new features. It can be used for provision of HDTV (high definition television) but it also allows for transportation of different multimedia streams such as, for example, internet traffic, audio and video streaming and file transfers with support of different input stream formats such as IP, ATM, single/multiple MPEG streams or generic bit streams, both for broadcast and unicast transmissions. For the support of interactive applications a return channel is necessary which can be provided by DVB-RCS [2]. DVB-S2 can achieve a capacity increase of up to 30% under the same transmission conditions compared to the older DVB-S standard what is achieved by applying higher order modulation schemes and by the use of low density parity check codes (LDPC) and Bose-Chaudhuri-Hochquenghem (BCH) codes.

The real novelty introduced by DVB-S2 was the possibility to use adaptive coding and modulation (ACM). In traditional nonadaptive systems the link dimensioning has to be made under considerations of service availability and worst case channel assumptions due to the deep fades caused by atmospheric effects; as a consequence the classical fade

mitigation techniques like power control result in an inefficient use of the system capacity since most of the time more transponder power than necessary is used. On the other hand in case of ACM, if a terminal is able to inform the gateway of its particular channel conditions (by means of a proper return link) the gateway can select an appropriate waveform, coding and modulation, to best exploit the spectrum and at the same time to overcome the channel impairments.

An efficient exploitation of the expensive satellite capacity has always been a key factor in the development of the satellite market, and the improvements brought by DVB-S2 give promising perspectives for the future of satellite communications. Nevertheless it is important to keep improving the exploitation of the satellite bandwidth, in order to guarantee reduced costs for all satellite services (broadcast, Internet, etc.). The aim of this work is to go one step further in this trend and to try to optimize the throughput and the spectrum efficiency in DVB-S2 forward links.

Today DVB-S2 links offer to the higher-layers protocols a terrestrial-like transmission medium, with recommended PERs around 10^{-7} . This is of course an excellent result, but not all services at higher layers require to reach such outstanding performance. This is in particular true for Internet and multimedia services [3].

Some audio codecs (e.g., AMR [4]) can typically accommodate packet losses with only a small impact in quality, and up to 15% failures before the speech is severely degraded.

Other modern media codecs (e.g., MPEG-4 [5]) have been designed to be highly resilient to residual errors in the input bit-stream, to detect and localize errors within the packet payload, and to employ concealment techniques, like for instance interframe interpolation, that hide errors from a human user. These codecs offer acceptable quality at a residual BER poorer than 10^{-3} , and some at poorer than 10^{-2} [6]. In order to support these error-tolerant codecs, the IETF has also standardized a new multimedia transport protocol, UDP-Lite [7], that allows to specify the required level of payload protection, while maintaining end-to-end delivery checks (verification of intended destination, IP header fields and overall length).

When these services are operating over the satellite connection, it is convenient to reduce the quality of the transmission in DVB-S2 forward links, by allowing higher BERs, in order to increase the precious capacity and the throughput. The first motivation for this is to make use of cross-layer mechanisms by voluntarily allowing higher bit error rates which can be compensated with error correction at higher layers. A second motivation to allow for higher BERs is that not all applications have the same stringent BER requirements. This represents a natural trade-off between errors and capacity. The present work analyzes this trade-off, proposes a way to tune the system parameters in order to work in optimal conditions, and investigates the performance of the system in this situation. The work is organized as follows: Section 2 presents the background and the scenario of the subject, Section 3 describes the main ideas of the paper and the original approach to the problem, Section 4 evaluates the performance of a system operating in the suggested conditions, and Section 5 drives the conclusions of the paper.

2. BACKGROUND AND SCENARIO

2.1. Overview of DVB-S2

The second generation of DVB-S provides a new way of fade mitigation by means of adaptation of the coding and modulation (ACM) to the different channel states. This of course implies the need for every terminal to signal its perceived channel state back to the gateway which can then make a frame-by-frame decision of the modulation and coding combination (ModCod) to be applied based on these measurements. DVB-S2 offers a broad range of modulations and codings for ACM. The supported modulation schemes comprise QPSK, 8-PSK, 16-APSK, 32-APSK and considered coding rates are 1/4, 3/4, 1/3, 2/5, 3/5, 4/5, 1/2, 5/6, 8/9, 9/10. The possibility to select the modulation and coding for an individual destination allows to make a more efficient use of the system capacity since transmission in a higher-order modulation in combination with a low coding rate (e.g., for clear sky conditions) allows to transmit more bits per symbol than a low-order modulation with high coding rate (e.g., for rainy channels). In this way it is possible to use individually for every ground terminal (or for every group of terminals in the same spot beam) the highest possible modulation scheme and the lowest coding rate which still allows to cope with

the channel impairments to provide low BER. A destination with a bad channel state can thus use a very robust modulation and coding pair (ModCod) while other terminals with a very good channel state can still transmit in highly efficient ModCods. The adaptive selection of the best suited ModCod results in an increased net data throughput while terminals in bad channel conditions are still able to receive their data since they can use ModCods with lower-order modulation and higher coding (but at the cost of lower spectral efficiency and thus lower throughput).

As can be seen in Figure 1 the system architecture of DVB-S2 is subdivided into six main components [1]. The mode adaptation subsystem provides an interface to the application specific data stream formats and also contains a CRC-8 error detection coding scheme. It is possible to merge different input streams together and to segment them into the so-called data fields which are the payload part of the so-called baseband-frames (BBFRAME) created at the output of the consecutive stream adaptation module. Buffers store data until they are processed by the merger/slicer and in case not enough data is available to fill a data field or if it is required to have only an integer number of packets in a frame (in general integer number of packets will not perfectly fit into a frame but their payload sum will always be slightly smaller or larger than the data field), the unused space can be padded, this operation is accomplished by the stream adaptation subsystem. In order to complete the baseband frame (BBFRAME) additional header information (BBHEADER) is added in front of the data field and scrambling of header and payload is applied. The final BBFRAME structure is illustrated in Figure 2.

The consecutive FEC encoding block performs outer and inner coding and bit interleaving. The coding scheme which is used is selected based on the channel measurements received from the terminals the data of which is contained in the frame. The outcome of this module, called forward error correction frame (FECFRAME), is shown in Figure 3. The FECFRAMEs can either have a length of 16200 bits for short frames or 64800 bits for normal frames. Since the length of the encoded frame is fixed, this means that the length of the payload in the underlying BBFRAME changes with the applied coding. For applying higher-order modulation schemes the subsequent mapping block performs a serial-to-parallel conversion. The mapper chooses the applied modulation schemes again based on the channel measurements for the destination(s) of the data contained in the frame. The outcome of the mapping of the data into symbols is called an XFECFRAME which is afterwards formed into a physical layer frame (PLFRAME) after pilots and PL signalling have been inserted and after final scrambling for optimization of energy dispersal. In case no XFECFRAMEs are provided by the preceding subsystems, the PLFRAMING module inserts the so-called DUMMY PLFRAMEs to provide a continuous TDM stream on the link. To allow every terminal independent of its channel state to receive the PLHEADER information (which also contains the used modulation and coding scheme for the underlying frame) this header is always modulated with BPSK.

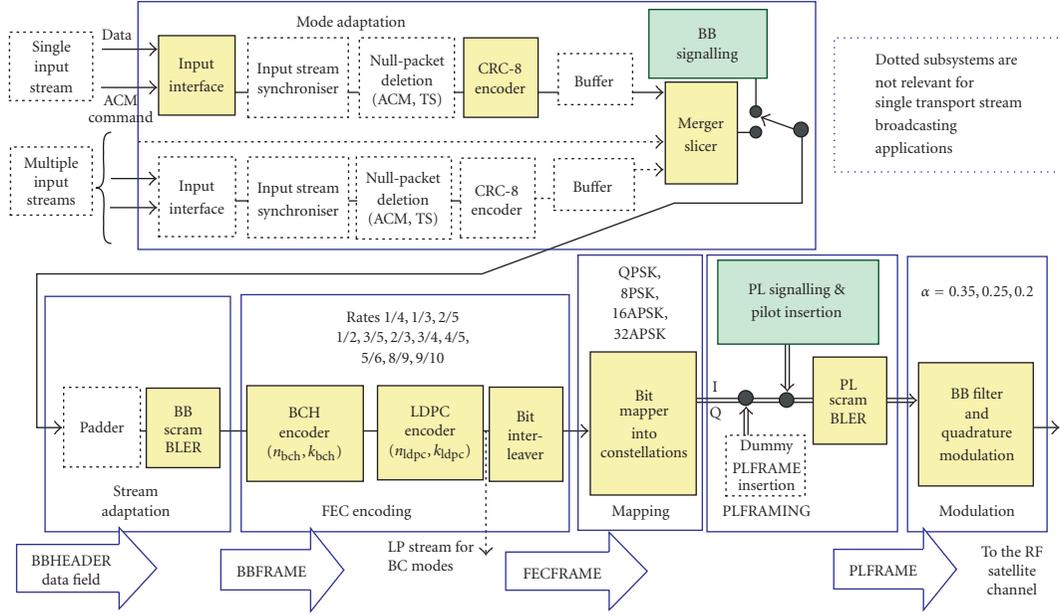


FIGURE 1: DVB-S2 system architecture [1].

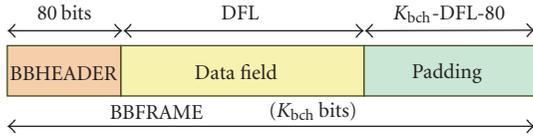


FIGURE 2: Structure of a BBFRAME [1].

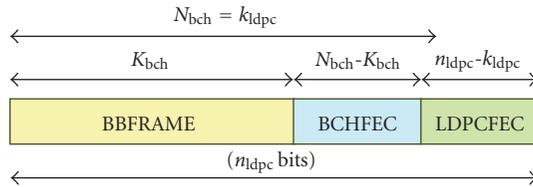


FIGURE 3: Structure of a FECFRAME [1].

For the selection of a ModCod that is adapting to the individual experienced channel states of the terminals, a return link must be provided to give feedback information about the measured channel states to the gateway. The gateway can then use this information to select a ModCod that suits transmissions in this channel state. This means the ModCod is selected to provide a quasi-error-free transmission as long as the critical SNR (signal-to-noise ratio) demodulation threshold for this ModCod ($\text{thr}_{\text{dem}}(\text{ModCod})$) is not crossed. If the signal drops below $\text{thr}_{\text{dem}}(\text{ModCod})$ then the BER will drastically increase due to the nature of the applied LDPC and BCH coding of having very steep BER-versus-SNR curves. In the GEO-stationary scenario investigated here, the propagation delay of the information feedback from the terminal to the gateway takes relatively long and it is in the order of sev-

eral hundreds milliseconds (250 milliseconds). This means that though the order of magnitude for the propagation delay allows for a compensation of very slow changing channel effects, like rain attenuation, it is too long to compensate fast, high-frequent changes in the SNR as those caused by scintillation, this will be explained in the next section.

2.2. Channel modelling

The selection of a ModCod scheme for transmission is very decisive for the performance of the system in terms of net data rate, bit errors and, respectively, packet errors. If the ModCods are selected too aggressively (meaning selection of ModCods with a too high modulation scheme and a too low coding) the transmission will result in a drastically higher PER. On the other hand, selection of safe ModCods (meaning a ModCod with a modulation lower than what would be necessary and a coding higher than necessary) will result in inefficiencies which reflects in a lower net data rate. In order to evaluate the influence of different parameters for the ModCod selection it is important to have a realistic channel model. The channels in satellite systems face mainly two sources of signal fading, rain attenuation and scintillation. The effect of rain attenuation is very significant for systems operating in K-band where the signal is attenuated by absorbing effects of the water. The second effect coming along with rain attenuation is scintillation which is basically a high frequent distortion of the signal amplitude and phase caused by small-scale irregularities in electron density in the ionosphere [8].

The scintillation in K-band can be considered to be a normal distributed random variable with a non linear spectrum (see [9, 10]) as shown in Figure 4. The standard deviation

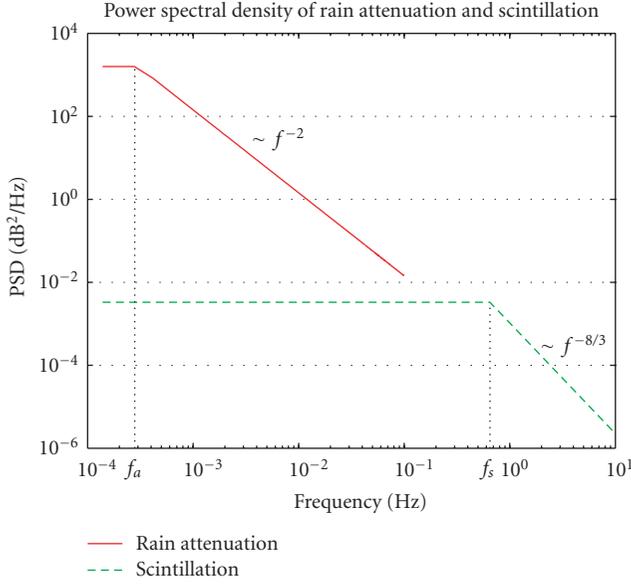


FIGURE 4: Attenuation and scintillation spectrum (typical values: $f_a \approx 10^{-4}$ Hz, $f_s \approx 0.1 - 0.65$ Hz).

of the scintillation process can be calculated according to (1) corresponding to the theory of Tatarskii [9] and the model of Matricciani [10],

$$\sigma = \sigma_0 \cdot A^{5/12}. \quad (1)$$

The value σ_0 is the standard deviation of the scintillation for a rain attenuation of A [dB]. [10] suggests a typical value of 0.039 for σ_0 in the frequency range of 19.77 GHz. According to (1) the resulting scintillation standard deviation σ is then in the order of tenths of a dB for rain attenuations smaller than 20 dB.

Within this work the main focus is on the scintillation effects since these cannot be compensated by signalling of the channel states via the return channel because of the long propagation delay of the GEO satellite. Nevertheless the channel simulations used in the rest of this work consider spatial correlated rain attenuation as well since the magnitude of the scintillation also depends on the intensity of the rain attenuation (see (1)). Similar to the generation of the scintillation, also the rain attenuation is created via a normal distributed random variable whereas its spectrum has a different corner frequency of f_a (see also Figure 4).

Figure 5 shows a channel example for the attenuation caused by scintillation and rain for a user located at longitude 8.6° E and latitude 52.7° N, in the area around Hamburg (Germany). It can be seen here that scintillation effects occur with a much higher frequency than regular rain attenuation events and how rain attenuation and scintillation are correlated.

2.3. ModCod switching strategies

While the rain attenuation occurs on a larger time scale scintillation effects occur very rapidly. For this reason rain fades

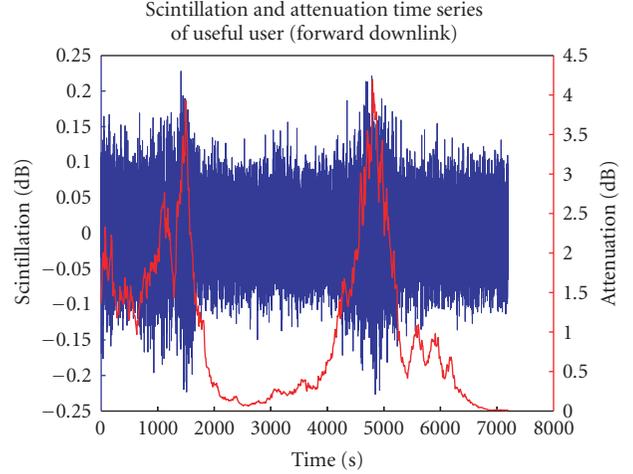


FIGURE 5: Example of scintillation and rain attenuation.

ing can be mitigated by mode adaptation whereas countermeasures for scintillation require a different compensation. For every combination of modulation and coding a threshold thr_{dem} (ModCod) exists which is needed to be able to decode the frame with a quasi-zero BER. The decision of the gateway on which ModCod will be used is thus driven by thresholds. For switching among ModCods these thresholds could theoretically be used directly for the decision about which ModCod will be used, but in practice this would result in frequent transmission errors since the high frequent variations of the channel (due to scintillation) would cause a frequent crossing of the threshold. On the one hand, this high frequent crossing cannot be compensated by signalling to the gateway, on the other hand such signalling would also mean a high frequent change of the ModCod which is as well undesirable.

To provide more reliability the minimal needed demodulation thresholds thr_{dem} (ModCod) can be replaced by thresholds which have a certain safety margin. This means that a lower ModCod is selected already before the critical threshold (the threshold below which a strong increase in bit error rate occurs) is reached. The size of the safety margin does thus determine the robustness against fast occurring scintillation fades. On the other hand, this size of the safety margin also influences the system performance since transmission in a higher ModCod would result in a higher net data rate. Since fast oscillations between neighboring ModCods are also possible when safety margins are used, an additional hysteresis margin is introduced. Figure 6 illustrates the different thresholds and margins.

Within Figure 6 the terms $\text{thr}_{\text{dem}}(N - 1)$ and $\text{thr}_{\text{dem}}(N)$ denote the minimum SNR values which are just enough to provide quasi error free decoding. If, for example, ModCod N is used and the signal strength falls below the $\text{thr}_{\text{dem}}(N)$ threshold, the BER will drastically increase. These thresholds have also been called critical in [11] for this reason. If on the other hand the signal strength increases, for example, while using ModCod $N - 1$, the next higher ModCod is not selected as soon as the demodulation threshold of the next higher

ModCod ($\text{thr}_{\text{dem}}(N)$ in this case) is crossed but after a higher threshold is exceeded ($\text{thr}_{\text{enable}}(N-1)$).

In case the signal strength decreases again, the ModCod is not switched when the enabling threshold $\text{thr}_{\text{enable}}$ is crossed, but just when an additional hysteresis margin is exceeded. The threshold for switching to a smaller ModCod is denoted as $\text{thr}_{\text{down}}(\text{ModCod})$ and the size of the hysteresis margin as ($\Delta\text{thr}_{\text{hyst}}(N)$). The distance between the critical demodulation threshold $\text{thr}_{\text{dem}}(\text{ModCod})$ and the threshold that triggers a downswitching $\text{thr}_{\text{down}}(\text{ModCod})$ is called the safety margin $\Delta\text{thr}_{\text{safety}}(\text{ModCod})$.

The safety margin ($\Delta\text{thr}_{\text{safety}}(\text{ModCod})$) can be seen as an additional security for high frequent oscillations which cannot be countervailed due to the long satellite propagation delay. If the signal strength oscillates within this area no increase in BER will occur since the $\text{thr}_{\text{dem}}(\text{ModCod})$ threshold is not crossed. The values for the safety margin and the hysteresis margin can be varied and they can also be different for every ModCod. Wörz et al. [11] presented a calculation method for all aforementioned parameters which provides a quasi error free system performance. The calculation of the parameters in [11] mainly depends on estimated values of the scintillation standard deviations and a numerically derived function which accounts for the fact that the standard deviation of the scintillation is also dependent on the intensity of the rain attenuation.

Within the remaining parts of this work the influence of the size of the safety margins with respect to gain or loss in channel net efficiency and increase/decrease of BER is investigated. The term “Wörz-Schweikert safety margins” denotes the safety margins calculated according to the algorithm presented in [11] while “zero-safety margin” denotes the fact that no safety margin is used.

2.4. Investigated environment

Within the examined scenarios, a set of user terminals has been located in a geographical region close to the city of Hamburg, Germany (longitude $9.5^\circ - 10.5^\circ\text{E}$, latitude $52.5^\circ - 54^\circ\text{N}$) within the aforementioned channel simulator. The set comprises 38 different terminal locations whereby the channel states are sampled with 10 Hz. The investigated duration is 7200 seconds per simulation run. In order to get statistical significant results the simulation duration of 7200 seconds per simulation have been extended to 60 hours. The 60 hours channel simulation results for the 38 terminals can be seen as 2280 hours of simulated channel states for a single terminal what in turn means that all results are based on channel information which corresponds to roughly a quarter of a year.

3. SYSTEM MODELLING

The main idea behind this study is that by reducing the safety margin in the ModCod switching strategy it is possible to gain in spectral efficiency, and thus to increase the net data throughput, at the expenses of an increased BER (and consequently a higher PER). In order to investigate this and to derive a detailed quantitative estimation of this trade-off, it

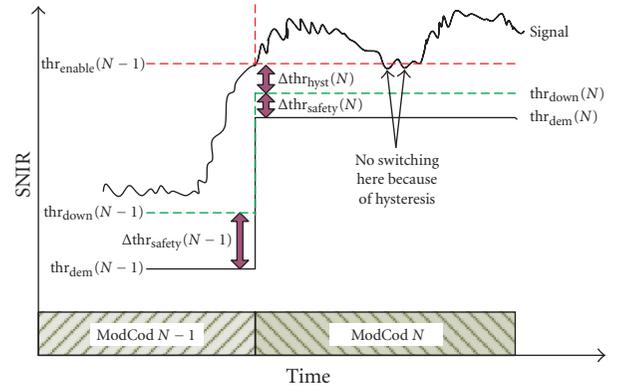


FIGURE 6: Illustration on the different thresholds.

is important to carefully describe the assumptions on which the analysis is based, this is what is presented in this section. Though the obtained results have a quantitative meaning only considering these assumptions, it is worth stating that their qualitative relevance has a general importance, as it will be later explained.

Existing systems compliant with the DVB-S2 standard can provide the higher-layers protocols with a quasi-error-free underlying physical layer ($\text{PER} = 10^{-7}$). For this purpose regular 8-bytes CRC (cyclic redundancy check) fields are used to identify errors in the BBFRAME, which were not corrected by the coding schemes (LDPC and BCH) at reception. In case an error is detected in a frame, it has to be considered that the wrong bit(s) cannot be singularly identified in the frame, so one of the two following choices can be made:

- (1) the whole frame is discarded (this is what is normally done);
- (2) the packets in the frame are passed to the higher protocols with uncorrected failures (this can be done in case the higher protocols are able to cope with errors).

These cases are very rare when high safety margins are adopted, and systems are normally dimensioned to avoid them, but they become more frequent if the system works closer to the demodulation thresholds (as we defined them in the previous section), for the reasons already explained.

If a system is dimensioned also to operate in these conditions, it is important to evaluate the statistical properties and characteristics of these situations, that is, how often they occur and what failures they bring in comparison to the capacity gain. In order to do that we performed three levels of analysis. They are theoretically described in this sections, whereas the results obtained for each of them are shown and discussed in the next one.

3.1. Markov model

The first analysis is a comparison of the new approach with a classical one existing in literature (the already mentioned Wörz et al. [11]), in terms of ModCod switching statistics.

The best way to show the difference between the two approaches is to model the system according to a Markov chain, where each state represents the system operating with one particular ModCod. The chain presents two states for each ModCod N : the good one, N_G , and the bad one, N_B ; so the overall number of states is twice the total amount of allowed ModCods (56). In the good state the SNR measured at the receiver is above the demodulating threshold for that ModCod, and so no failures are expected, in the bad state the system SNR is below the demodulating threshold for that ModCod, and so failures may occur with probabilities that are not negligible.

A similar Markov chain is an excellent model, because it summarizes very well the properties of a ModCod switching approach. So once the transition probabilities for one particular ModCod switching criterion have been calculated (normally by simulations), the Markov chain can be used as a basis for all types of analysis without the need of running again computationally heavy simulations, which might be very long in order to gather statistically meaningful data. In this sense the calculated Markov chain (i.e., the ModCod transition probabilities) can be considered independent from the simulated channel conditions, only if the simulation is long enough to represent general channel statistics. On the other hand, it should be mentioned that the same Markov chain depends on some parameters which might be characteristic of particular cases, for example, the link budget in clear sky, and consequently the system availability. So even if the resulting numbers are only meaningful bearing in mind these assumptions, the quantitative conclusions which can be derived have general relevance, and this will be clearer in the next section.

3.2. Error rate versus capacity trade-off

The second level of analysis describes the details of each state of the Markov chain. The good states present quasi-error-free conditions according to the DVB-S2 recommendations, so $PER = 10^{-7}$. Since the BER versus SNR characteristics for all ModCods are very steep, the BER values increase quite rapidly when the SNR level goes below the demodulating threshold. In particular they change of several orders of magnitude within a few tenths of dB, going from $BER \approx 10^{-10}$ when SNR is close or bigger than the demodulating thresholds, up to $BER \approx 10^{-2}$ when the SNR is just 0.3 dB below the threshold. Each bad state represents a set of different BERs, the proper BER is selected at each time step according to the received level of SNR with respect to the demodulating thresholds. The exact characteristics for the BER-versus-SNR functions, which were used in the simulations and to derive the Markov chain parameters, were taken from [12]. In that work, end-to-end performances of the BER versus the SNR are presented for the DVB-S2 system, the whole communication chain is modelled and simulated, including coding, modulation with predistortion techniques, satellite transponder impairments, downlink, demodulation with the synchronization, and the final LDPC and BCH decoders. In the Markov chain, each ModCod might have its own BER

versus SNR characteristic, but to use one single function for all ModCods already seems an excellent approximation to the real case, so this is how it was implemented in the simulator.

PERs are derived from these BERs under consideration of the payload length of each BBFRAME also regarding the applied ModCod. A BBFRAME is considered as erroneous if at least one of the payload bits is erroneous. For the rest of this paper the term PER denotes the BBFRAME packet error rate. Thanks to this definition of the states of Markov chain, this model allows to derive the properties of the communication in terms of PER and BER statistics, and by knowing the spectral efficiency associated to each ModCod it is easy to derive an average resulting capacity.

3.3. Error bursts analysis

The third and last level of analysis goes into the details of the failures introduced with this novel approach. In the previous section we explained how to derive a measure of the trade-off between average capacity and average BER (or PER). An average measure of the BER (or PER) does not seem a very precise information, since these failures come in bursts. The errors are mainly due to the ModCod switchings, and they are mostly introduced by reduced safety margins. So we want to investigate three main properties: (i) how often the error bursts arrive (interarrival times statistics), (ii) how long the bursts last (duration statistics), and (iii) how deep the fades are (i.e., how high are BER and PER during one error burst). These three properties can be estimated thanks to the Markov model, and this analysis produces interesting information, which will be presented in the next section.

4. RESULTS EVALUATION

4.1. Markov model

A software simulator was developed in order to derive the Markov model presented in the previous section. Once the ModCod switching criterion has been specified the software simulates the evolution over time of the system; from these simulations we can derive statistics about the permanence in the different ModCods for each ModCod switching criterion, this was done by computing transition matrices and solving them. In the following we present two full transition matrices for two different ModCod switching criteria.

Simulations equivalent to 3 months of SNR time series have been carried out, one using Wörz-Schweikert safety margins, the other one using zero-safety margin bounds with Wörz-Schweikert hysteresis bounds.

The matrices in Figures 7 and 8 represent the transition probabilities for those two approaches, where position (i, j) is the probability in each time step (0.1 second) to move from state i to state j ; the first line and the first column of each ModCod represent the bad state (i_B and j_B), the second one the good state (i_G and j_G). Figures 7 and 8 show the transition matrices for zero-safety margin and the Wörz-Schweikert safety margins. The cells marked black indicate that their content is unequal to zero. In Figure 8, we can see that the

	1	2	3	4	5	6	7	8	9	12	11	13	14	18	19	20	21	
1	1.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
2	0.00	1.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
3	0.00	0.00	1.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
4	0.00	0.00	0.00	1.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
5	0.00	0.00	0.00	0.00	1.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
6	0.00	0.00	0.00	0.00	0.00	1.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
7	0.00	0.00	0.00	0.00	0.00	0.00	1.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
8	0.00	0.00	0.00	0.00	0.00	0.00	0.00	1.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
9	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	1.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
12	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	1.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
11	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	1.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
13	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	1.00	0.00	0.00	0.00	0.00	0.00	0.00
14	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	1.00	0.00	0.00	0.00	0.00	0.00
18	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	1.00	0.00	0.00	0.00	0.00
19	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	1.00	0.00	0.00	0.00
20	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	1.00	0.00	0.00
21	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	1.00	0.00

FIGURE 7: Transition matrix for zero-safety margin.

	1	2	3	4	5	6	7	8	9	12	11	13	14	18	19	20	21	
1	1.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
2	0.00	1.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
3	0.00	0.00	1.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
4	0.00	0.00	0.00	1.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
5	0.00	0.00	0.00	0.00	1.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
6	0.00	0.00	0.00	0.00	0.00	1.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
7	0.00	0.00	0.00	0.00	0.00	0.00	1.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
8	0.00	0.00	0.00	0.00	0.00	0.00	0.00	1.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
9	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	1.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
12	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	1.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
11	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	1.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
13	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	1.00	0.00	0.00	0.00	0.00	0.00	0.00
14	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	1.00	0.00	0.00	0.00	0.00	0.00
18	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	1.00	0.00	0.00	0.00	0.00
19	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	1.00	0.00	0.00	0.00
20	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	1.00	0.00	0.00
21	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	1.00	0.00

FIGURE 8: Transition matrix for Würz-Schweikert safety margin.

more “stable” states are the good states. This makes sense if we look at the SNR time series, because switchings between ModCods are quite spaced in time compared to the time step of 0.1 second. Simulations estimate an average number of 4.5 ModCod switchings per hour.

Moreover on the diagonal, the probabilities of remaining in a bad state are not so high; this is also correct, since when we are in a bad state, we know that a down-switch should occur. The only bad state which has a higher stability is the bad state for ModCod 1_B. This results from the fact that when the SNR goes below the last demodulation threshold, the system cannot switch to a lower ModCod, so it remains in bad state until the SNR rises again. This is basically an outage where the DVB-S2 receiver is not available; the simulator was designed to give a system availability of 99.96% of the time, for both approaches.

With the *Würz-Schweikert* scheme, shown in Figure 8, the matrix is far more sparse, and no bad states are ever active, except for ModCod 1_B because of system unavailabil-

ity. Transitions only occur between good states and this confirms that this approach is designed to work only in good states.

For the novel approach, the zero-safety margin one, it may be interesting to derive the probability to be in each ModCod (bad or good state). Once the transition matrix for the Zero-Safety margin is solved [13], we end up with Figure 9 which shows a stacked probability density graph for good and bad states of each ModCod. This is the result of a simulation of an equivalent of 4.5 years of SNR evolution over time. What we can see is that the most used ModCods are those whose demodulation threshold is just below the SNIR in clear sky conditions. That makes sense because most of the time we are in clear sky conditions, so we use the highest ModCods. We can also notice the high value of the bad state in ModCod 1_B, because of system unavailability. Some ModCods are never used due to overlapping with other ones, some ModCods achieve a better spectral efficiency requiring less SNR.

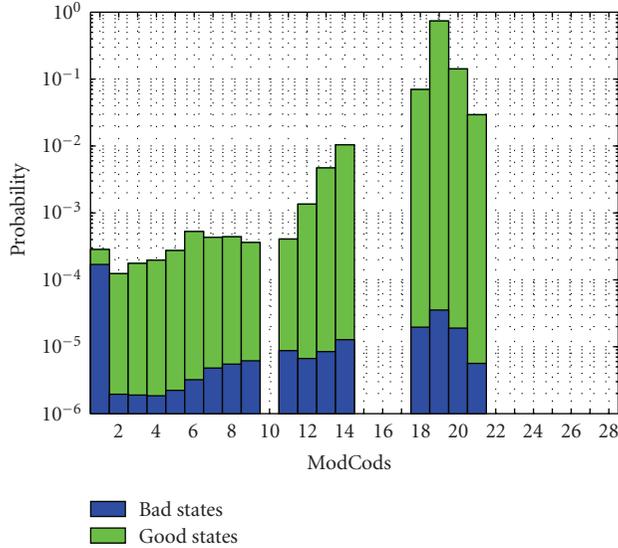


FIGURE 9: State probabilities.

4.2. Error rate versus capacity trade-off

This section presents the main results which are obtained when reducing the safety margins, in terms of increase spectral efficiency and increase errors. The starting point is the set of threshold selected by Wörz-Schweikert; this set guarantees a quasi-error-free system operation. We try to proportionally reduce those margins and even to have negative margins, to see how the system performs. The x axis in Figures 10 and 11 represents factors to be multiplied to the Wörz-Schweikert set to get the tested thresholds. This means that for multiplying factor 1 we have the Wörz-Schweikert set, for the factor 0, we have the zero-safety margin approach, and for negative values of the factors we are testing thresholds which are below those thresholds recommended by the DVB-S2 standard. This may seem strange, but it will appear clear how useful this is to show that there is a trade-off between errors and increase in capacity.

Figure 10 shows (as expected) that the PER objective of 10^{-7} is achieved already before the Wörz-Schweikert bounds. This is not surprising since the model has been designed to do so. As expected as well, PER and BER are fast-growing up to 1 when the safety margin becomes negative. A surprising fact here is that there are possibilities to achieve the goal PER even for margins which are 0.4 times the Wörz-Schweikert safety margins. That means that those Wörz-Schweikert margins may not be the optimum selection.

Figure 11 shows the core result of this work. A trivial thing is that the gross capacity (total amount of received bits with failures) is still increasing when we go for lower and lower bounds, because of course we are using less and less robust ModCods that provide better spectral efficiency. The very interesting point comes with the fact that the net capacity (throughput of correct bits) shows a maximum in the negative part of the scaling factor: -0.4 at the packet level

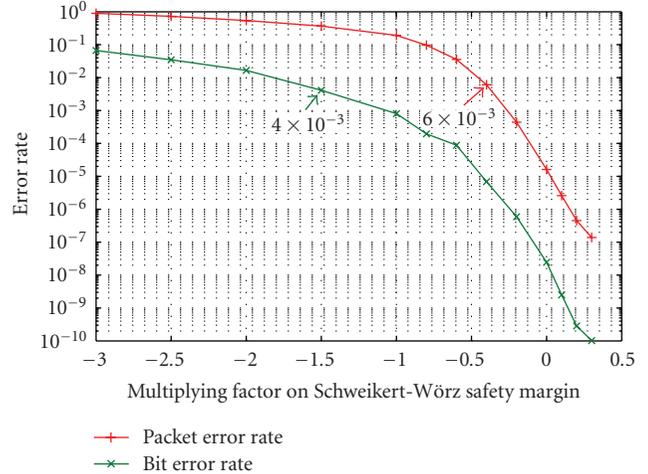


FIGURE 10: Packet error rate (PER) versus Safety margin.

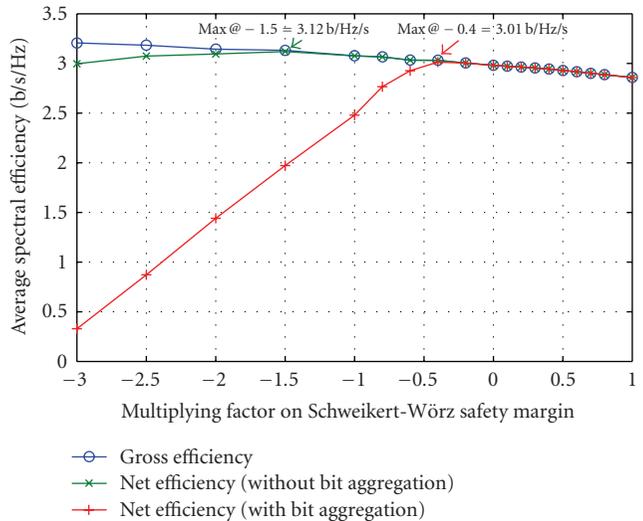


FIGURE 11: Average spectral efficiency versus safety margin.

and -1.5 at the bit level. Corresponding values of PER/BER at these maxima are $6 \cdot 10^{-3}$ and $4 \cdot 10^{-3}$. The two curves represent the two ways of operating described in Section 3: bit aggregation is when failures cause BBFRAME discard, no bit aggregation means when the frame is passed to the higher layers with failures. It should be noted that for bit error aggregation (see Figure 11) the PER (see Figure 10) is the relevant result since in case of a bit error the complete BBFRAME is discarded. Without consideration of bit error aggregation, the BER is the relevant result since erroneous bits within the BBFRAME are expected to be corrected by the higher layers. This means that a system which wants to have the indicated throughput with or without bit aggregation, is operating at those PER/BER.

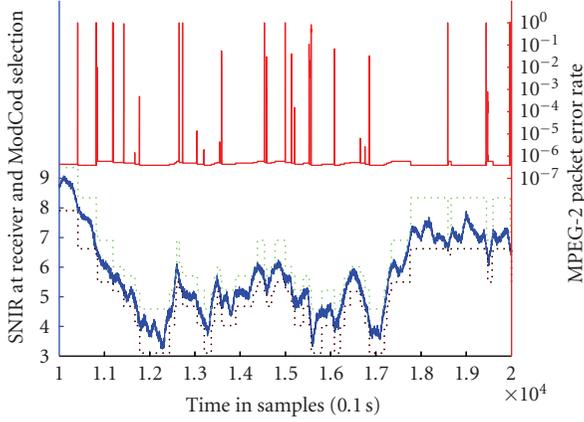


FIGURE 12: PER with SNR and ModCod selection for zero-safety margin.

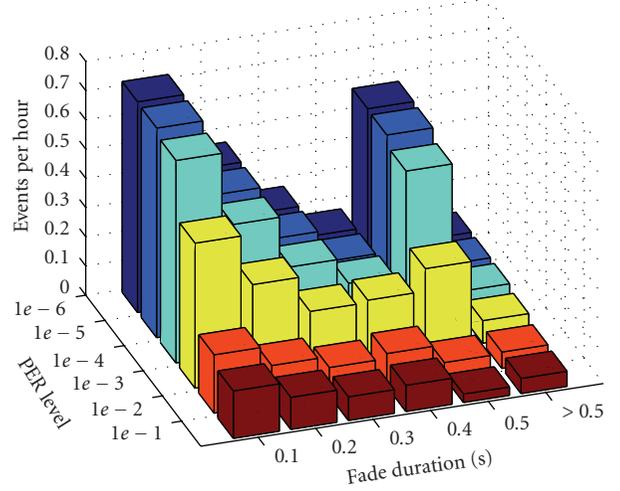


FIGURE 14: State probabilities.

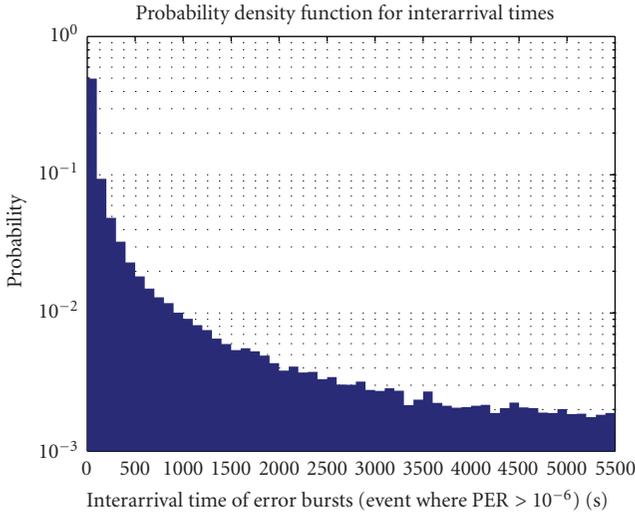


FIGURE 13: Interarrival times distribution.

If a system can cope with these error rates, then it may be interesting to design it with lower safety margins than those in the Wörz-Schweikert strategy, in order to gain throughput.

4.3. Error bursts analysis

To deeper investigate the quality of the transmission in case we reduce the safety margins, we have to look at the distribution in time of the error bursts. Figure 12 shows an example of simulated SNR time series with corresponding PER for zero safety margin. In contrast to Figure 10 which shows the averaged error PERs and BERs, here we investigate the distribution of the interarrival times between two PER peaks (without averaging), considering a detection threshold of $PER = 10^{-6}$. The simulation that has led to Figure 13 has been worked out on 4.5 years of simulated SNR, and it was conducted with the zero-safety margins approach.

We see that in 50% of cases, the time between two error bursts is in the range of 0–100 seconds. This distribution comes from the fact that during a rain fade, ModCods are switched down one by one, and as we saw on Figure 12, error peaks often occur at every down-switch. The question is now what is the duration/severity of these peaks?

Figure 14 shows the number of fade events per hour using zero-safety margins, sorted by their duration and PER strength. A sequence of samples is considered as one fade event if the associated PER is exceeding a given level. For each PER level, Figure 14 shows the number of fades per hour which exceed this PER level. For this graph, we have 6 different PER levels, and the fade events are distributed among their duration. We can see that the shorter fades are the ones that occur most of the time. This comes from the fact that for a normal process like scintillation, the probability of having a fade is decreasing exponentially with its duration. There are peaks for each PER level at 0.5 second, this is due to the fact that independently from how long the fade would be, in the worst case the system can switch to a lower ModCod within half a second (twice the GEO propagation delay), which is the time needed to signal to the gateway the fading situation and to receive a new transmission with a new ModCod. So in theory fade should not exceed 500 milliseconds, but the last bin of this bar plot shows that even if they are rare, fades exceeding 0.5 second do exist. There are two explanations for that. First, if we are in the highest ModCod of a couple of very close ModCods (in terms of demodulation threshold) and we enter a strong rain fade, with a steep decreasing SNR, it can happen that the SNR crosses the demodulation threshold of the lowest ModCod before the system has switched down. This results in a bad state to bad state transition, and we can see some of these cases in the transition matrix (12_B to 9_B or 13_B to 11_B , e.g.). A second explanation is the following, non-negligible contributions to this behavior are the outages due to system nonavailability, that is the fades that occur in the lowest ModCod.

5. CONCLUSIONS

The possibility to have a quasi-error-free transmission channel in DVB-S2 systems is not always an optimal solution in case the higher-layer protocols do not require such high performance. In this case the lower layers can provide a transmission with some resilient errors, and exploit more the spectrum to gain in throughput. The error-capacity trade-off can be tuned, according to the requirements of each particular system, with the adjustment of the ModCod safety margins. The paper presents the gain in spectral efficiency, which is obtained with this method, and the statistical characteristics of the “artificially” introduced error bursts, in terms of interarrival, duration and depth (PER). One additional interesting side-outcome of this work is the development of Markov chain to model the ModCod transitions and the failure occurrence in a DVB-S2 system.

ACKNOWLEDGMENTS

This work was partly supported by EC funds SatNEx under the FP6 IST Programme, Grant number: 507052. This work was supported by the European Satellite Network of Excellence (SatNEx).

REFERENCES

- [1] ETSI EN 302 307 V1.1.2, “Digital Video Broadcasting (DVB); second generation framing structure, channel coding and modulation systems for broadcasting, interactive services, news gathering and other broadband satellite applications,” June 2006.
- [2] ETSI EN 301 790 V1.4.1, “Digital Video Broadcasting (DVB); interaction channel for satellite distribution systems,” April 2005.
- [3] G. Fairhurst, M. Berioli, and G. Renker, “Cross-layer control of adaptive coding and modulation for satellite Internet multimedia,” *International Journal of Satellite Communications and Networking*, vol. 24, no. 6, pp. 471–491, 2006.
- [4] ETSI TS 126 102, “AMR Speech Codec,” 2001.
- [5] ISO/IEC 14496-2, “Coding of audio-visual objects (MPEG-4)—part 2: visual,” 2004.
- [6] ETSI TR 126 975, “Performance Characterisation of the Adaptive Multi-Rate (AMR) Speech Codec,” 2004.
- [7] L.-A. Larzon, M. Degermark, S. Pink, L.-E. Jonsson, and G. Fairhurst, “The Lightweight User Datagram Protocol (UDP-Lite),” IETF, RFC 3828, 2004.
- [8] S. Datta-Barua, P. H. Doherty, S. H. Delay, T. Dehel, and J. A. Klobuchar, “Ionospheric scintillation effects on single and dual frequency GPS positioning,” in *Proceedings of the 16th International Technical Meeting of the Satellite Division of the Institute of Navigation (ION GPS/GNSS '03)*, pp. 336–346, Portland, Ore, USA, September 2003.
- [9] V. I. Tatarskii, *Wave Propagation in a Turbulent Medium*, McGraw-Hill, New York, NY, USA, 1961.
- [10] E. Matricciani, M. Mauri, and C. Riva, “Relationship between scintillation and rain attenuation at 19.77 GHz,” *Radio Science*, vol. 31, no. 2, pp. 273–280, 1996.
- [11] T. Wörz, R. Schweikert, A. Jahn, and R. Rinaldo, “Physical layer efficiency of satellite DVB using fade mitigation techniques,” in *Proceedings of the International Communication Satellite Systems Conference (ICSSC '05)*, Rome, Italy, September 2005.
- [12] E. Casini, R. De Gaudenzi, and A. Ginesi, “DVB-S2 modem algorithms design and performance over typical satellite channels,” *International Journal of Satellite Communications and Networking*, vol. 22, no. 3, pp. 281–318, 2004.
- [13] M. F. Neuts, *Matrix-Geometric Solutions in Stochastic Models: An Algorithmic Approach*, Dover, Mineola, NY, USA, 1981.