TCP behavior in a DVB-RCS environment

Cesare Roseti
University of Rome "Tor Vergata" via del Politecnico 1, 00133, Rome, Italy

Erling Kristiansen
European Space Agency, Keplerlaan 1, NL-2200 AG Noordwijk, The Netherlands

In 2000, ETSI has approved and published a standard for interactive broadband satellite services (DVB-RCS) [1,2]. One of the most important features of this standard is the implementation of DAMA strategies aiming at a fair and efficient use of the shared resources in the satellite return channel. A DAMA scheme introduces a new control loop, based on a signaling exchange between STs and the NCC, installed in the GW. As a consequence, a new component is added to the overall end-to-end delay: The "access delay". TCP, on the other hand, the transport protocol used by most current Internet applications, is based on the RTT experienced in the network. Therefore, performance of TCP over DVB-RCS is effectively governed by two nested control loops with similar time constants. Depending on the circumstances, these two nested control loops may cooperate to give good performance, or may interact in a way that degrades performance. This paper analyses the interaction between the TCP and DVB-RCS control loops by exploring the performance of TCP over the different capacity allocation categories defined in the DVB-RCS standard. An access simulator, based on ns-2 [3], has been implemented and used to study the dynamics at both MAC and TCP layers, as well as their cross-interaction. Furthermore, a set of analysis approaches, including theoretical treatment, emulation and real measurements are presented as complementary tools to have an exhaustive knowledge about TCP performance in DVB-RCS environment.

Abbreviations

ACK	=	Acknowledgement.
AVBDC	=	Absolute volume based dynamic capacity.
awnd	=	TCP advertised window.
CA	=	Congestion Avoidance.
CRA	=	Constant rate allocation.
cwnd	=	TCP congestion window.
DAMA	=	Demand Assignment Multiple Access.
DVB-RCS	=	Digital video broadcast – return channel by satellite.
ETSI	=	European Telecommunication Standard Institute.
FR-FR	=	Fast-Retransmit and Fast Recovery.
GW	=	Gateway. DVB-RCS hub station.
NCC	=	Network Control Centre.
PDF	=	Probability Distribution Function
NEP	=	Network Engineering Platform.
Ns-2	=	Network Simulator 2.
RBDC	=	Rate based dynamic capacity.
RTO	=	Retransmission Time Out.
RTT	=	Round trip time. Time to travel from sender to recipient and back to sender.
SS	=	Slow Start.
ssthresh	=	TCP slow start threshold.
ST	=	Satellite terminal
TCP	=	The Internet Transfer Control Protocol.
TDM	=	Time division multiplex.
VBDC	=	Volume based dynamic capacity.

I. Introduction

As DVB-RCS (Digital Video Broadcasting – Return channel by Satellite) [1,2] is becoming established in the market of multimedia Internet via satellite, there is a growing interest in evaluating its performance under various conditions of traffic profile and network load, with the aim of understanding the underlying issues and optimizing its performance.

The subject of this paper is to study TCP performance over DVB-RCS networks. The scope is limited to the study of standard end-to-end TCP. The performance enhancement by means of protocol adaptations or introduction of new network elements in the path (e.g. Performance Enhancing Proxies) is outside the scope of the paper.

A simulation model, based on ns-2 [3], is proposed that reproduces the standard DVB-RCS access schemes and consequently allows a very flexible analysis of the behavior of a TCP flow as a function of a large number of parameters (i.e., RCS frame structure, TCP initial parameters, RTT, etc.).

The paper is organized as follows: section II introduces the main issues this work deals with; section III describes the reference system architecture focusing on the DAMA scheme; section IV provides some basic features of the TCP mechanisms; section V describes the simulation model; section VI proposes further approaches that can complement the simulation analysis; section VII shows the achieved results; and section VIII draws some conclusions.

II. Presentation of the problem

The Internet reliable transport protocol TCP [4,5] has been designed to cope with a very broad range of environments. Nevertheless, current TCP implementations tend to be optimized for typical terrestrial wired networks. When used over networks that do not fit this model, the protocol still works, but performance tends to suffer [6,7].

When using TCP over satellite networks, performance is influenced negatively by 3 factors:

1) Long link delay means long round trip time (RTT) for the mechanisms controlling sending rate and congestion management.
2) TCP acts defensively to variation in available bandwidth, potentially leading to under-use of communication resources.
3) TCP assumes all packet losses are due to congestion. A packet lost due to errors on the link induces TCP to reduce its sending rate unnecessarily, thus wasting resources and/or reducing performance.

In a DVB-RCS environment the use of DAMA mechanisms may introduce effective RTTs well above the propagation delay, and the available bandwidth may vary strongly and abruptly [8,9]. These two factors are essential in determining the achievable performance.

Packet loss due to link errors is less of an issue because DVB-RCS uses strong forward error correction at the link layer, providing to the network layer and above a quasi error free channel.

When TCP operates over a DAMA satellite link, this effectively leads to two nested control loops: TCP's rate control loop and DVB-RCS's bandwidth allocation control loop. The control bandwidth of both loops is of the same order of magnitude: Somewhere between half a second and a few seconds. This inevitably leads to the two control loops interacting with each other. The objective of this paper is to study the interaction, using a variety of methods.

III. Brief description of DVB-RCS

The reference model for a DVB-RCS network is shown in the Figure 1, where three main components can be identified:

- A Gateway (GW) or Hub station,
- A GEO-stationary satellite,
- Several Satellite Terminals (ST's).

The network topology is a star network with the GW station as hub.

The GW transmits a forward carrier towards all ST's. The ST's, in turn, share a channel towards the GW according to a Demand Assignment Multiple Access (DAMA) protocol.

The GW typically provides access to the Internet and/or to a corporate network.

ST's cannot communicate directly with each other, but some systems allow double hop connection via the GW.

DVB-RCS bandwidth management is asymmetric:

In the forward link (gateway (GW) to satellite terminal (ST)) a TDM scheme is used. This is managed by the GW, and the way the available bandwidth is shared between concurrent flows is an implementation issue. Typically, the GW maintains several FIFO queues with different priority.

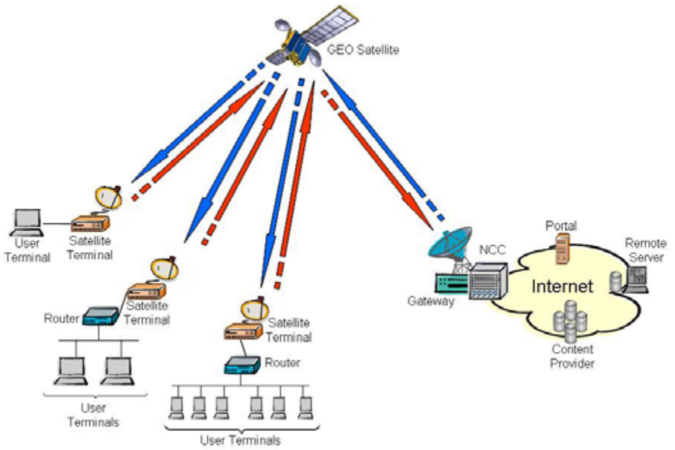

Figure 1: Reference model for the Satellite Interactive Network

In the return link, bandwidth is shared between ST's according to a DAMA scheme. This study considers 3 types of bandwidth allocation [1]:

- CRA (Constant Rate Allocation) is a fixed allocation, available permanently to an ST irrespective of whether it is used.
- VBDC (Volume Based Dynamic Capacity) is dynamically allocated on request of the ST. The request is for a given volume of data. Each request is a one-off request for a given volume, and requests are cumulative. The time between issuing the request and getting the allocated bandwidth is typically in the order of 1.5 seconds, but may vary from one system to another depending on the details of the frame structure and the precise allocation algorithm. The standard supports also an Absolute VBDC (AVBDC) scheme in which requests are absolute rather than cumulative, typically used to cope with potential loss of requests. In our analysis we will refer only to VBDC.
- RBDC (Rate Based Dynamic Capacity) is dynamically allocated on request of the ST. The request is for a given data rate and is typically valid for some seconds, being updated by subsequent requests as required. The extra delay of the request/allocation cycle is only felt on the initial request for RBDC, and on requests that modify the requested rate. The impact of dynamic allocation on RTT is therefore less severe than for VBDC.

Some systems support a fourth type of bandwidth allocation: FCA (Free Capacity Assignment). Any bandwidth not used for fulfilling requests is split between ST's in some way. This feature is not included in our study. It requires knowledge of, or assumptions about the total traffic in the channel, and its effect is therefore unpredictable from the viewpoint of a single or a few ST's. We felt that including it would bring little benefit, and would make interpretation of results more difficult, having one more parameter to cope with.

From a performance point of view, the 3 types of allocation can be described as follows:

- CRA has a constant RTT of around 500 ms for a geostationary satellite channel. Bandwidth is reserved whether or not it is used. Any unused bandwidth is wasted.
- VBDC has a much longer RTT, typically around 1.5 seconds, due to the request/allocation cycle. Normally, capacity will only be requested for data that is already in the buffer, so the bandwidth utilization is 100%.
- RBDC has an RTT similar to CRA, except some extra delay on first request and on adjustment of requested rate. The bandwidth provided will rarely match precisely the bandwidth needed, so some bandwidth will be wasted most of the time, but less than for CRA.

Most implementations offer combinations of two, or all three of the allocation schemes.

The precise algorithm for requesting and granting VBDC and RBDC is not part of the DVB-RCS standard, and it is difficult to get information about the details of implementations. So we had to make some assumptions.

VBDC is rather straightforward, and there are not many choices available.

RBDC, on the other hand, is much more open to different algorithms. We chose a relatively simple model based on a sliding average of recent traffic demand.

It is our understanding that VBDC is intended for best effort traffic like TCP, while RBDC is intended for stream traffic that has more or less constant bandwidth requirements over extended time periods. We will discuss in the conclusion whether these expectations are met.

IV. TCP basics

TCP is a reliable connection-oriented transport protocol in which the in-sequence delivery of each data unit, called "segment", is notified to the sender by a cumulative acknowledgement (ACK) [4,5]. Based on this, TCP implements three control mechanisms: flow control, congestion control and error recovery.

The flow control mechanism uses the "sliding window" algorithm [10] to continuously control the amount of data "in-flight", while congestion control mechanism gradually probes the network status by increasing the window until a segment loss, interpreted as a congestion signal, occurs and an error recovery algorithm is triggered. Two status variables are used to perform the TCP control mechanisms: the "congestion window" (cwnd), representing the size of the sliding window, and the "advertised window" (awnd), indicating the free space in the receiver buffer. The minimum between cwnd and awnd determines the transmission window.

In this work, we consider the standard TCP NewReno version [11], which is based on 5 algorithms: Slow Start (SS), Congestion Avoidance (CA), Fast Retransmit and Fast Recovery (FR-FR), Retransmission Time-Out (RTO) [5]. SS runs either at the start of the connection or after a timeout expiry (RTO), performing an exponential increase of the cwnd every time ACKs are received. Subsequently, TCP switches to the CA algorithm when cwnd reaches a threshold value (ssthresh) or after the FR-FR phase, allowing a gentler probing of the available capacity by a linear increase of the cwnd.

As far as the error recovery mechanisms are concerned, the RTO is the default mechanism used by TCP. Basically, TCP sets a timer for each sent segment and, if the corresponding ACK is not received within the timer expiration, it retransmits the segment. The optional FR-FR algorithms allow a faster recovery of the lost segments by exploiting the fact that TCP sends a copy of the last ACK (duplicate ACK) every time a segment is received out of sequence. In practice, the reception of a small number of duplicate ACKs (usually 3) is considered evidence that a segment has been lost and the Fast Retransmit performs an immediate retransmission (without waiting for the time out expiration) Fast Recovery prevents possible congestion events by both halving cwnd and switching TCP to CA.

V. A simulation platform

A simulation platform, modeling the DBV-RCS access schemes, has been built on the satellite network extensions of the network simulator Ns-2 [3]. In particular, we used these extensions to model a geostationary "bent-pipe" satellite node connected to STs and a GW through asymmetric links, a population of STs and a satellite network interface stack for each ST. In addition, modifications on the internal C++ code and the creation of *tcl* procedures have allowed to model a DVB-RCS-like MAC layer, supporting all relevant DAMA disciplines.

The structure of the implemented model is represented in the figure 2. First, to configure a particular allocation discipline, several configuration parameters must be set for each ST:

- Maximum capacity allowed (in slots);
- Number of slots assigned by CRA (N_{CRA});
- Maximum number of slots managed by RBDC (N_{RBDC});
- Sampling weight γ for the RBDC algorithm.

CRA is implemented by a module acting as a "setup filter". It is run at the beginning of the simulation and communicates to the "DAMA controller" the number of slots per frame (N_{CRA}) permanently granted to a given ST. As a consequence, the DAMA controller decreases the number of slots it will dynamic assign "on demand".

The RBDC module computes capacity requests based on the principle of making input and output rate in the ST queue equal. To this purpose, it acts as a filter that, every superframe, gets as input the amount of data in the queue not served by CRA, computes the RBDC request on the basis of both the N_{RBDC} value and the following algorithm:

$$R_{RBDC}(k) = (1-\gamma)\cdot R_{RBDC}(k-1) + \gamma \cdot Q_k \qquad (1)$$

(where Q_k is the amount of data queued at the k-th superframe and γ is a weighting factor that determines the amount of "memory"of past history).

The VBDC discipline is managed through a "DAMA Entity" module. Every superframe, a capacity request is computed by using the following algorithm [12]:

$$R_{VBDC}(k) = \left\lceil Q'(k) - A(k) - \sum_{j=1}^{L-1} R(k-L+j) \right\rceil^{+} \qquad (2)$$

The inputs to the VBDC algorithm are

- the amount of data (Q'_k) not processed by others access schemes (i.e., CRA, RBDC);
- the capacity allocated for the considered ST in the current superframe ($A(k)$);
- the requests sent in the previous superframes but not yet taken into account by the DAMA controller. In particular, the parameter L indicates the "resource allocation periods" between the transmission of the capacity request and the activation of the corresponding capacity (advertised by the BTP), named "System Response Time" or "access delay". It includes the propagation delay and the processing delays at both STs and GW sides. A dedicated module, called "Pending Requests", stores such requests.

A "DAMA Controller", logically located in the GW, manages the request scheduling and, every superframe, assigns the available slots on the basis of the received requests. Then, depending on the assigned capacity, the DAMA Controller regulates the transmission time of the queued data (managed by the "MAC Object").

Furthermore, the implemented allocation scheme has been designed to respect typical values concerning both the frame structure and processing delays. In the default setting, the frame supports a capacity of 2048 kb/s and is divided into 32 units/slots. The frame duration is 24 ms, while a superframe is constituted by 4 frames (the duration is 96 ms). The "System Response Time", defined as the time between the request sending and the corresponding capacity assignment, is set equal to 11 superframes (1056 ms).

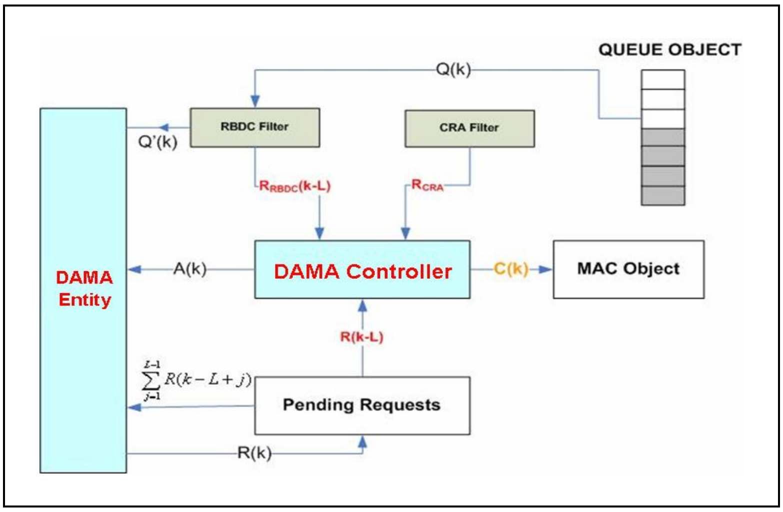

Figure 2: Allocation scheme

A key goal of this simulation platform is to provide a flexible tool to evaluate any kind of combination among the DVB-RCS standard allocation schemes.. For example, an ST could transmit data through 2 slots permanently assigned (CRA), while 3 additional slots could be requested in RBDC and 4 further slots in VBDC.

A vast gamut of statistics (graphs or data files) can be obtained as simulation outputs. In particular, such outcomes can be divided into thee categories:

- *Lower layer statistics*. These concern the allocation dynamics at the MAC layer, as:
  - The "System Response Time" measurement;
  - Trace of the slots allocated in each superframe;
  - Calculation of the "Saturation Time", the time at which the maximum sustainable transmission rate has been achieved.

- *End-to-end statistics*. Statistics gathered in the end-systems:
  - RTT;
  - Throughput.
- *TCP statistics*. Monitoring of the evolutions of the TCP internal parameters:
  - Cwnd and ssthresh trend;
  - Trace of the retransmissions.

VI. Further analysis approaches

A complete evaluation of the dynamics characterizing the communication protocols requires the reproduction of the "real life" of the network. In fact, delay, bandwidth constraints, loss distributions, queue disciplines, impact of competing traffic, hardware limitations are just some of the factors affecting the protocol dynamics and the performance. A simulation aims at isolating the most critical features of the reference communication system in order to allow a simplified and cost-effective analysis. This tendency to move from the general to the particular can lead to the loss of some secondary, but relevant, effects impacting on the system behavior, especially when dynamic processes are performed. For example, in the simulation platform illustrated in the previous section, we have modeled the layers below IP just as a delay and a variable bandwidth due to the DAMA algorithm, so all aspects concerning ATM/MPEG2 functions or the processing tasks at the physical layer have been ignored. On the other hand, simulation allows to easily experiment with a very wide range of parameter values, and allows to measure any desired parameters, including some that are not easily accessible in real life.

Since simulation is based entirely on a simplified mathematical model, there is a need to validate simulation results against other data.

In this section, we briefly describe three further analysis approaches that are useful in order to verify the consistency of the simulation results:

- A theoretical performance model;
- Emulation platforms;
- Measurements on real systems.

VI-A. A theoretical model

A theoretical approach allows to isolate all the intrinsic characteristics of the involved protocols from the overall behavior that can be observed in the real network. Then, a reformulation of the main formulas, governing TCP window dynamic, can emphasize the effect of the two aforementioned nested loops.

In particular, the TCP bandwidth control loop is based on the RTT perceived by the sender, while DAMA schemes affect the RTT by an additional delay contribution named "access delay" due to the exchange of capacity request/allocation messages. Therefore, by assuming that the NCC is installed in the GW, DAMA and TCP loops present similar time constants, proportional to the physical RTT.

Then, if a volume-based allocation discipline (VBDC) is considered, the RTT perceived by TCP can be expressed as:

$$RTT(t) = RTT_0 + L \cdot n_s + \beta(t) \cdot n_s \qquad (3)$$

where:

- RTT_0 is the physical round-trip delay;
- n_s is the number of TDMA frames constituting the superframe and representing the "resource allocation period";
- $L \cdot n_s$ is the system response time;
- $\beta(t)$ is a component related to the congestion state of the network.

But, DVB-RCS supports also CRA and RBDC allocation disciplines either as stand-alone access schemes or combined with VBDC. In the latter case, an ST could have, superframe per superframe, slots allocated by different allocation disciplines, and (3) can be re-formulated as follows:

$$RTT(t) = RTT_0 + L \cdot n_s + \beta(t) \cdot n_s - \alpha(t) \cdot n_s \qquad (4)$$

The α parameter is the CRA/RBDC gain, and depends on both the number of slots assigned in CRA/RBDC discipline and the amount of data stored in the ST queue. In general, CRA and RBDC allow to decrease the average

RTT (α >0) as a consequence of the fact that CRA discipline allocates statically slots without dynamic negotiations, while RBDC computes request on the basis of the rate the data feed the ST queue and then, once that slots are allocated, they are booked also for a certain number of successive superframes.

Therefore, by both assuming an un-congested network state ($\beta(t) = 0$) and considering the average CRA/RBDC gain ($\bar{\alpha}$), the time needed to transmit a file consisting of d_T TCP segments can be approximated as follows [13]:

$$T_{TRANSFER} = \underbrace{\underbrace{\left(RTT_0 + L \cdot n_s - \bar{\alpha}\right)}_{DVB-RCS\_control\_loop} \cdot f\left(TCP_{cong\_phase}\right)}_{TCP\_control\_loop} \quad (5)$$

where $f(TCP_{cong\_phase})$ depends on both the performed TCP congestion algorithm (i.e., SS, CA) and the initial internal parameters (i.e., ssthresh, cwnd, RTO value).

VI-B. Emulation

The use of real satellite equipments may be costly, and often the presence of proprietary features in the protocol stack makes the outcomes very difficult to interpret. A further approach to test protocols/applications in real environments is emulation. In fact, emulation platforms allow to include “real network entities”, while several aspects of the communication (e.g. the propagation delay) are emulated by software/hardware tricks. Of course, to reproduce the dynamic behavior of the DVB-RCS DAMA scheme, an “active” emulation platform is needed, providing the ability to dynamically change network parameters (e.g. bandwidth) on the basis of the instantaneous traffic.

The Network Engineering Platform (NEP) is a network emulation tool developed by Alcatel Alenia Space, initially for internal use by the company, but later upgraded for more general use. The European Space Agency has procured a copy of the NEP and offers it for use by ESA contractors.

The NEP aims at emulating DVB-S/DVB-RCS satellite access systems as far as network, access and resource management layers are concerned. In particular, it provides a real-time environment to assess and benchmark the performance of both popular internet applications and satellite dedicated applications.

NEP is composed of 6 interconnected Linux PCs covering all the functionalities of a real satellite access network DVB-RCS like:

- 3 PCs acting as STs;
- 1 PC implementing GW and NCC functions;
- 1 PC providing the Terrestrial Interconnection Sub-System (TISS);
- 1 configuration station implementing the “Man Machine Interface” (MMI) and coordinating NEP operation.

Working on real equipments and protocol implementations, a certain number of practical issues have to be addressed in order to have a complete control of the tests. In particular, parameters concerning application programs and operating systems had to be appropriately set to allow a TCP window growth to fill the channel pipe value. To this purpose, we performed some tests by running Iperf [14] to create a TCP flow from a user terminal to the GW. Adjusting the maximum TCP buffer size set in operating system kernel and the Iperf window size, we determined the optimal values allowing TCP to fill the channel pipe under all the supported access schemes. With default buffer sizes, the maximum achievable throughput will be limited by buffer size, and will be lower than the available bandwidth.

Figure 3 compares the RTT measurements obtained by considering the main aforementioned access schemes. In the CRA and VBDC cases, a perfect match with the expected values can be observed (physical delay in the CRA case, physical delay + access delay in the VBDC case), with the exception of some higher values detected during the start-up phase. These higher values are likely to be artifacts of the detailed implementation of the NEP.

As far as RBDC is concerned, some aspects in the RTT pattern need to be commented:

1- At the start up, the RTT values are very similar to those of VBDC. The rationale is that, due to the bursty nature of the TCP flow, RBDC needs some time to adjust to the real source activity. In fact, when the new rate samples are zero for several superframes, the RBDC gradually reduces the resources granted, and when new data feed the queue it behaves exactly like VBDC.

2- Due to the bursty nature of the TCP traffic, RBDC is slow in establishing a stable estimate of the transmission rate. Consequentially, the oscillations in the rate estimate lead to oscillations in the perceived RTT.

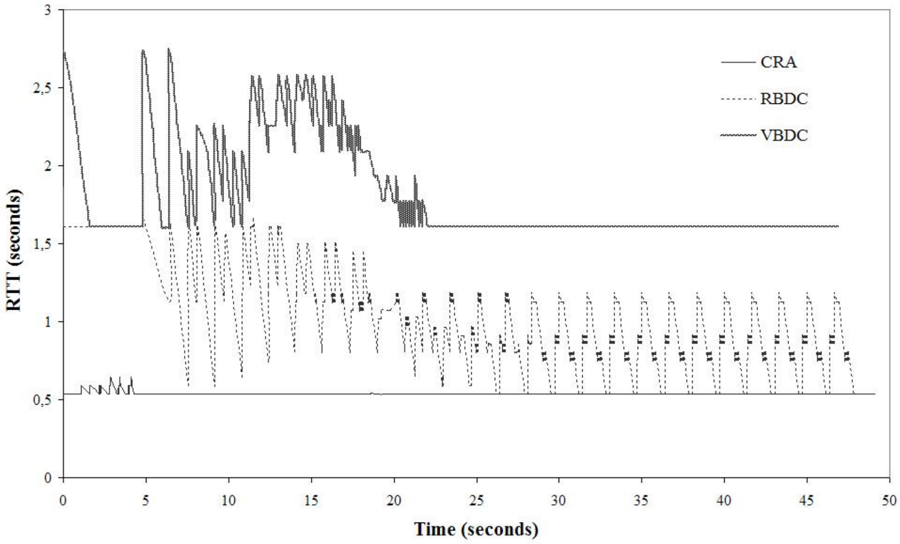

Figure 3: NEP tests: RTT measurements

VI-C. Measurements over real systems

A full analysis about the interaction between DVB-RCS access schemes and TCP performance should be completed by trials over real satellite systems. Unfortunately, this turns out to be extremely difficult due to a number of factors. Basically, commercial systems use proprietary solutions as far as both request/allocation algorithms and TCP acceleration are concerned. The satellite segment often appears as a “black box”, of which the precise function is not known, and the interpretation of performance, in terms of DVB-RCS characteristics, is a very hard task.

Nevertheless, it is possible to exploit tests over a real DVB-RCS link to gather useful information concerning statistics of some physical parameters as, for example, the probability density function of the satellite delay. To this purpose, we performed some tests over a DVB-RCS network operated commercially by Belgacom. In particular, we obtained the probability density function of the observed RTT between an ST and the GW as shown in figure 4.

Then, even if simulation and emulation appear as the best approach to analysis details, real trials can be used as a support tool aiming to provide most reliable inputs.

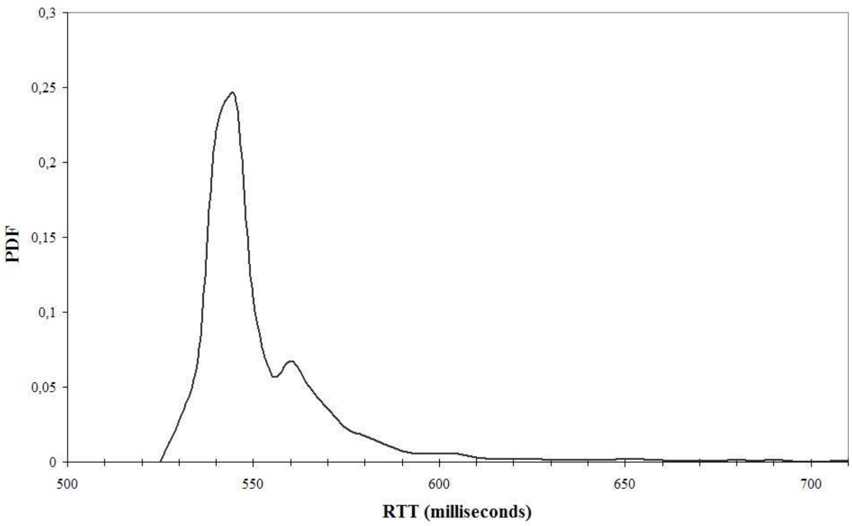

Figure 4: RTT distribution - CRA

VII. Simulation results

Our analysis aims at showing TCP behavior over the different access schemes provided by the DVB-RCS standard. In particular, we consider the case where a TCP source is on an ST while the TCP destination is supposed to be installed on the GW. Such assumptions are not restrictive since usually PEP mechanisms are implemented at the two ends of the satellite link isolating the satellite segment from the others along the whole communication path. However, the eventual terrestrial sub-links do not contribute significantly to the overall delay and the satellite segment can be considered, without loss of generality, as the bottleneck.

Furthermore, we focused on the possibility to combine the different standard access schemes in order to study the correspondence between the access policy and the achieved performance over a large set of alternative schemes. Our main goal is to highlight the dynamics of TCP over a DVB-RCS network and then identify the best access scheme for a targeted QoS.

It is clear that by allocating statically the capacity (i.e., CRA) there is a benefit to the end-to-end performance at the expense of the overall resource usage efficiency, while a “demand-based” access scheme (e.g., VBDC) assures an optimal resource usage efficiency at the expense of a deterioration of the end-to-end performance. Our aim is to identify the trade-off “solution” that allows to achieve the required QoS by maximizing the overall efficiency as well.

Therefore, the following analysis can be divided into two parts:

- Emphasis on the capability of the created simulation platform in handling layer 2 parameters/statistics. This is the main limitation when real equipments are used.
- Evaluation on the performance perceived at the end-systems (higher layer statistics) corresponding to the provided QoS. In this frame, we investigate the effects produced by all the standard access scheme and combinations of them.

Figure 3 shows four output graphs concerning different access schemes. In particular:

3a) Dynamics of the slots allocated per superframe, in the case a hybrid access scheme (“1 CRA slot” + “31 VBDC slots”) is run.

3b) Overlaying of the graphs of respectively the data queued in the ST buffer and capacity (in bytes per superframe) allocated, when RBDC (γ=0.1) is considered as access scheme.

3c) Comparison between PDF of the RTT with different access schemes.

3d) Trace of the RTT in the case an RBDC (γ=0.1) access scheme is considered.

Then, due to the flexibility of the simulation approach, it appears evident how we could use the simulation to obtain the most accurate statistics and to have complete control of all the parameters at both the MAC and the transport layers.

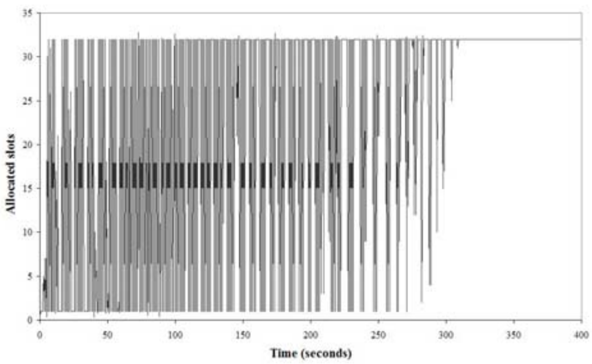

(a)

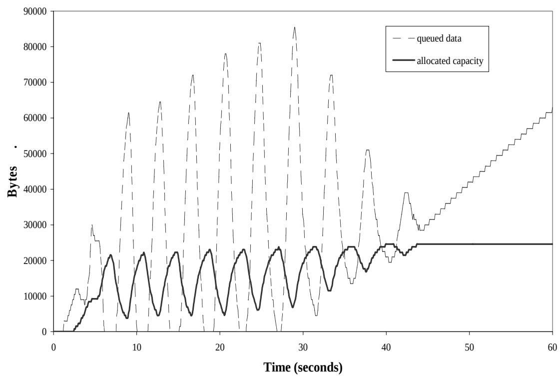

(b)

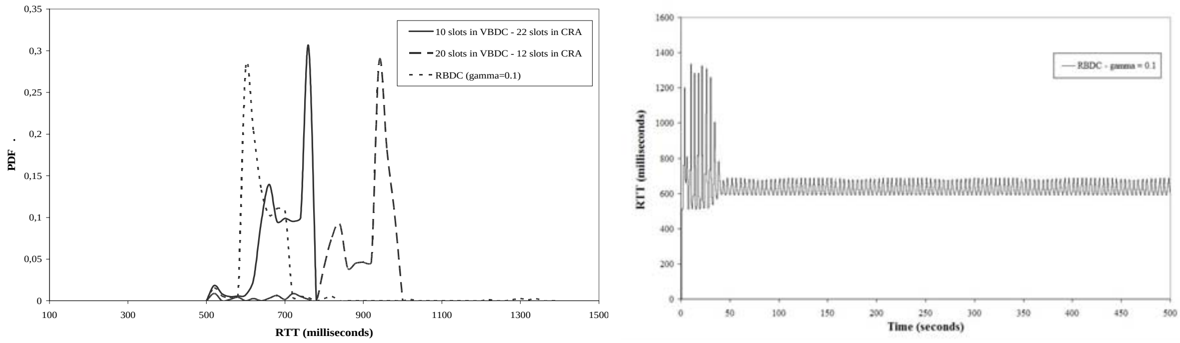

(c) (d)

Figure 5: Examples of simulation outcomes

As seen in the theoretical analysis (section VI-A), the key parameter for TCP performance is the RTT. For this reason, by associating to each access scheme an estimated RTT value, it is possible to deduce the corresponding theoretical performance perceived by the end-systems. In the following graphs we show the RTT values for various access schemes, representative of the main allocation classes treated in this paper, and the corresponding throughputs achieved by the TCP source.

In particular, Figure 6a shows the RTT perceived by TCP over 4 different access schemes:

- CRA;
- VBDC;
- RBDC (γ=0.1);
- Combination of VBDC and CRA (5 slots statically assigned in CRA, and the remaining 27 slots negotiable by VBDC).

In the simulations, we set the receiver buffer size equal to about 1 MB in order to avoid that it contributes to the perceived RTT due to ST “internal” congestions, occurring when the amount of data feeding the buffer exceed the bandwidth-delay product. By observing the trend of the graphs, a certain number of considerations can be made:

- VBDC and CRA are characterized by constant RTT values, respectively equal to the physical delay and the physical delay + the access delay (as defined and discussed in section III);
- RBDC leads to clear oscillations in the RTT at the beginning due to the bursty nature of the TCP traffic, whereas when saturation occurs (the maximum TCP window is achieved) the RTT becomes quite stable. In particular, the oscillations are caused by the alternation of two phases:
  1. In the previous superframes R_k=0. Then, when new data feed the queue, the time spent in the request-allocation process is very close to that one needed in the VBDC.
  2. In the previous superframes $R_k \neq 0$. Then, with reference to (1), the new request will be not null even if there is no data in the queue. Consequently, when further data will feed the queue, a given amount of capacity could be already assigned.
- In the “hybrid” VBDC/CRA access schemes, since 5 slots are statically assigned, the dynamic allocation process concerns just a part of the stored data. Therefore, when the amount of data is less or equal to the assigned capacity, the RTT involves just the physical delay; in contrast, when the amount of stored data grows, a part of capacity is requested by the VBDC algorithm (2), and the RTT is increased as well, until it approaches a saturation value that, on average, is less than that one perceived over the VBDC scheme.

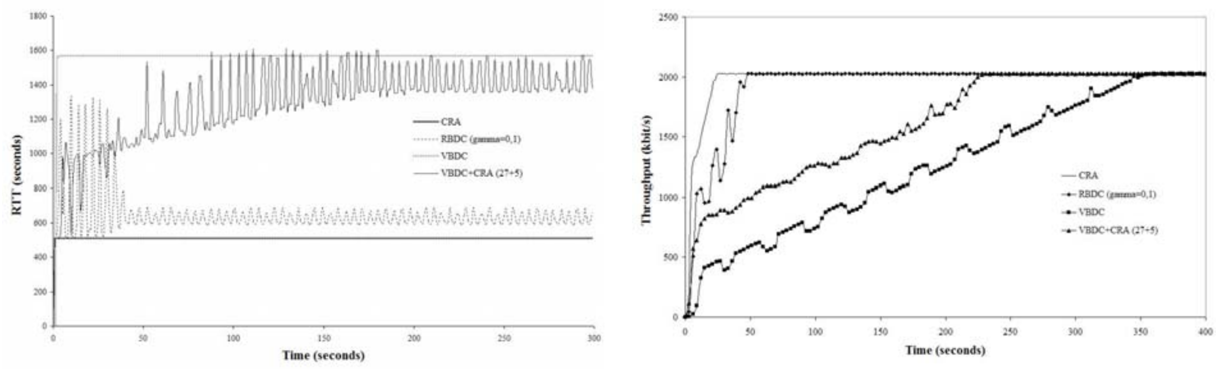

(a) **(b)**

Figure 6: End-to-end TCP statistics

Figure 6b shows the throughput variations by considering the 4 access schemes already mentioned in the figure 6a. These graphs summarize all the conclusions achieved up to here, highlighting the time needed to exploit the total capacity (throughput = 2048 Kbit/s), called saturation time, and the dynamics of the allocation algorithm at the start-up (oscillations in the throughput).

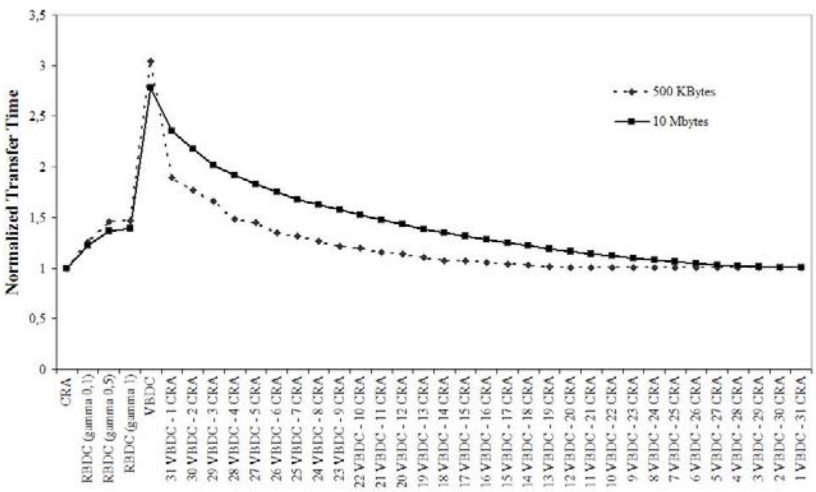

Figure 7: Normalized file transfer time

The last set of results, illustrated in the figure 7, concerns the simulation of file transfers by FTP from an ST to the GW. All the standard DVB-RCS access schemes, all the possible combinations between VBDC and CRA and different weights γ (0,1-0,5-1) in the RBDC algorithm (1) have been considered. Moreover, the tests have been carried out with a small file of 500 Kbytes and a large file of 10 Mbytes in order to distinguish short and long term behaviors. The transfer time has been normalized to the transfer time detected in the CRA case. By analyzing the figure, a certain number of observations can be made:

- The normalized transfer time over the VBDC access scheme is equal to 3. This validates the conclusions achieved by theoretical model expressed in (5), for which the transfer time is proportional to the RTT. In fact, as aforementioned, the RTT for VBDC is about 3 times the RTT of CRA.
- RBDC schemes present better performance for large files than for small files. This confirms the fact that RBDC suffers from the bursty nature of the TCP traffic during the start up phase.
- Lower γ values makes the RBDC algorithm more stable (at the expense of the network efficiency), leading to better performance.

- The hybrid CRA-VBDC schemes provide higher gains in the transfer time in the case of short files. The rationale is that the percentage of the data served by CRA is higher.

VIII. Conclusion and future work

The main goal of this paper has been to analyze the effects of the interaction between the TCP and DAMA control loops on the end-to-end performance by different approaches: theoretical analysis, simulation, emulation, real trials. Then, we used the results coming from different sources to achieve reliable conclusions. Furthermore, such a cross-analysis allows to investigate TCP behavior over a DVB-RCS compliant network by highlighting aspects belonging to different layers of the OSI protocol stack: i.e. lower layer statistics coming from real trials or simulation or the performance perceived by the end-systems reproduced by emulation.

As far as the comparison among the standard access schemes is concerned, three basic comments can be made:

- All the proposed access schemes can be ranked in-between the VBDC scheme, that optimize the network utilization (at the expense of the end-to-end performance concerning a single TCP flow), and the CRA scheme, that aims at avoiding of any “access delay” and then at optimizing TCP performance (at the expense of the overall network utilization);
- RBDC and the hybrid CRA-VBDC access schemes represent trade-off solutions that, depending on the traffic characteristics and the required QoS, provide more or less improvements on the performance.
- Our expectation was that VBDC was designed to serve best-effort traffic like TCP. We must, however, conclude that, though VBDC has the highest potential for high resource utilization, it actually offers low throughput performance for TCP traffic, in particular for the transfer of relatively short files.

Therefore, the most important conclusion is that it is impossible to identify a best access scheme in general, but on the basis of the nature of the traffic, the network load, the QoS required by each terminal, a scheme can become more suitable than another. In this frame, the proposed work provides the guidelines for seeking the access scheme suitable for a particular communication scenario, by showing the main characteristics of each.

In the future, it could be useful to join all the parameters characterizing the communication scenario under a qualitative index to use as “key” in seeking the best access scheme for a given environment. In other words, the next step is to make tables where the “optimal” access scheme can be identified as function of parameters depending on the ST activity and on the economics of network resource allocation. Such a framework would also need to consider non-TCP traffic, a topic that we were unable to include in our study, being constrained by the available time and resources.

References

[1]ETSI, *Digital Video Broadcasting (DVB); Interaction Channel for Satellite Distribution Systems*, DVB-RCS standard, EN 301 790.

[2]ETSI, *Digital Video Broadcasting (DVB); Interaction Channel for Satellite Distribution Systems; Guidelines for the use of EN 301 790*. TR 101 790, V. 1.2.1, 2003.

[3]http://www.isi.edu/nsnam/ns/ns-documentation.html.

[4]Stevens, W., *TCP/IP Illustrated vol. 1,* Addison Wesley, MA, USA, 1994.

[5]Stevens, W., *TCP Slow Start, Congestion Avoidance, Fast Retransmit, and Fast Recovery Algorithms*, RFC 2001, Jan.1997.

[6]Partridge, C., Shepard, T. J., *TCP/IP Performance over Satellite Links*, IEEE Network, Sep. 1997, pp. 44-49.

[7]Balakrishnan, H., Padmanabhan, V. N., Fairhurst, G., Sooriyabandara, M., *TCP Performance Implications of Network Path Asymmetry*, RFC 3449, Dec. 2002.

[8]Sooriyabandara, M., Fairhurst, G., *Dynamics of TCP over BoD satellite Networks*, International Journal of Satellite Communications and Networking, Vol. 21, No. 4-5, Jul. 2005, pp. 427-449.

[9]Neale, J., Mohen, A., *Impact of CF-DAMA on TCP via satellite performance*, Proceedings of IEEE Globecom ’01, Nov. 2001.

[10]Luglio, M., Roseti, C., Gerla, M., *The impact of Efficient Flow Control and OS Features on TCO performance over Satellite Links*, ASSI Satellite Communication Letter, 9th edition, Vol. 3, No. 1, 2004.

[11]Floyd, S., and Henderson, T., *The NewReno Modification to TCP’s Fast Recovery Algorithm*, Internet RFC 2582, 1999.

[12]Karaliopoulos, M., Tafazzoli, R., Evans, B. G., *Providing Differentiated Service to TCP Flows Over Bandwidth on Demand Geostationary Satellite Networks*, IEEE Journal on Selected Areas in Communications, Vol. 22, No. 2, Feb. 2004.

[13]Cardwell, N., Savage, S., Anderson, T., *Modelling TCP Latency*, IEEE INFOCOM, Mar. 2000.

[14]http://dast.nlanr.net/Projects/Iperf/