

Modeling the Behavior of a DVB-RCS Satellite Network: an Empirical Validation

D. Adami, S. Giordano, M. Pagano, R. Secchi

Dipartimento di Ingegneria dell'Informazione - Università di Pisa, Via Caruso, I-56122 Pisa
E-mail: {d.adami, s.giordano, m.pagano, raffaello.secchi}@iet.unipi.it

Abstract

Since satellite networks provide a valuable bandwidth access to vast regions at relatively low costs, recently the interest in Internet access through satellite links has considerably increased. At the same time, the TCP/IP protocols suite is widely used in all the communication environments, encompassing the most important network applications. Thus, the actual exploitation of satellite technology into the global Internet is possible only by moulding satellite access control scheme to the TCP/IP network features. Starting from a set of measurements carried out in a DVB-RCS network, in this paper we propose an analytical model of the RBDC medium access control scheme. The model highlights how bandwidth-on-demand allocation mechanisms may lead to large packet delivery delays when sudden variations in the incoming traffic rate, typical of multimedia traffic flows, occur. Due to the effectiveness of the model, it is possible to use it for improving resources allocation strategies.

1 Introduction

In contrast with traditional satellite broadcasting systems, the new generation satellite networks allow the convergence of low-cost broadband connectivity with the technology required for user interaction. Since the satellite links provide Internet access at high bitrate to vast regions, thus allowing a broad portfolio of services with different characteristics and QoS requirements, the broadband satellite networks have been advertised as a significant part of the future global Internet. As a consequence, understanding the behavior of TCP/IP applications in a satellite network is a compelled step for the integration of such networks into the global communication infrastructure.

Generally, depending on the service class, the access to a satellite network can be provided by means of *static* or *dynamic* capacity allocation schemes. When the static approach is used, during the connection set-up phase, the satellite terminal that submitted the query negotiates the bandwidth with Network Control Center (NCC). This technique, which guarantees a fixed amount of bandwidth to a given traffic terminal (TT) for the whole connection lifetime, is not widespread in satellite networks, because bandwidth is a precious resource and must be cost-effectively used.

Therefore, a number of dynamic bandwidth allocation schemes have been proposed and developed [1][2]. In these schemes, the NCC dynamically assigns a certain amount of bandwidth to each TT on the basis of explicit capacity requests according to a specific Demand Assignment Medium Access (DAMA) scheme. This way, the NCC is able to keep track of the actual traffic generated by the TTs and to modify the bandwidth assignment in order to better exploit the satellite link capacity. On the other hand, since the TTs have to wait for an explicit capacity assignment from the NCC, the connections traversing the satellite link may experience long delays before gaining access to the satellite network. This phenomenon may reduce the throughput mainly for short-lived connections.

Unlike traditional approaches, which aim to describe the satellite link only in terms of delay and losses, in this paper we propose a model for a satellite connection that explicitly accounts for the rate-based bandwidth-on-demand (BoD) allocation scheme. Many reasons motivated our choice to develop the theoretical model. First, an analytical approach is useful to have a better understanding of satellite system behavior at IP layer and to identify the most critical aspects of BoD schemes in TCP/IP traffic management. Moreover, we believe

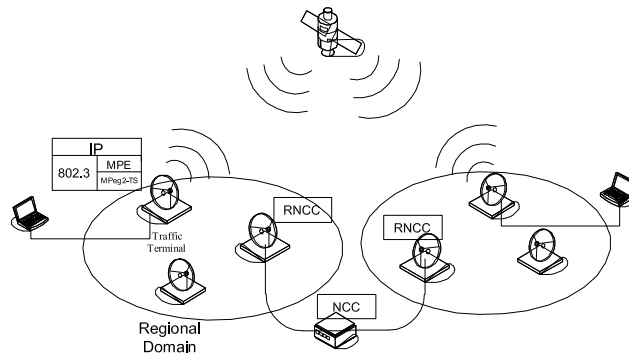


Figure 1: DVB-RCS Network Architecture

that a model of the system can be helpful in the design of the resource management and allocation strategies because it discloses the *tuning knobs* to enhance the satellite link utilization.

Since the adaptation of networking features to the satellite MAC is a relatively new research topic, a little amount of literature can be found on TCP/IP performance modeling over BoD access networks. The interaction of TCP with BoD schemes was discussed first in [3][4] with some simplistic approximations for the resource allocation procedures. Instead, a more rigorous approach has been adopted in [5], in which an analytical model (based on Fixed-Point approximation), describing the interaction between BoD and an aggregate of TCP connections, was developed. However, the authors of [5] mainly focus on volume-based allocation methods, while in this paper we are treating rate-based schemes.

The rest of the paper is organized as follows. Section 2 describes the characteristics of the satellite technology and introduces an empirical model of rate-based allocation schemes. In section 3, the satellite field trial setup is presented and some experimental results are discussed. Section 4 concludes the paper with final remarks.

2 DVB-RCS System

Digital Video Broadcasting with Return Channel via Satellite (DVB-RCS) [6] has emerged as the platform to integrate the satellite broadcasting capabilities with the Internet infrastructure. The IP/DVB architecture defines the basic set of functionalities required for users interactivity across a satellite DVB-S distribution system. Basically, IP packets are firstly encapsulated into an MPE (MultiProtocol Encapsulation) structure containing the source/destination MAC addresses of the traffic stations involved in the data transfer. Then, MPE packets are delivered using a transport stream of 188 bytes *MPEG-2* cells.

Fig. 1 shows the main elements of our DVB-RCS network. The network is divided into regions, corresponding to the geographic areas covered by each satellite spot beam. The group of TTs that belong to a specific region is monitored and controlled by a centralized authority, called Regional Network Control Center (RNCC), which is connected through a leased line with the NCC, an higher level network management system. The Multi-Frequency Time Division Multiple Access (MF-TDMA) scheme [6] is used to share the satellite uplink capacity. In this way, the time/frequency domain is organized in a hierarchical structure, which consists in superframes (the largest containers), multiframes and frames.

In order to establish the time-slot assignment for the traffic terminals, the RNCC broadcasts periodically a burst time plan (BTP) table message to its group of TTs. The DVB-RCS standard allows several methods of bandwidth allocation, which can be combined to guarantee better QoS requirements:

1. **Constant Rate Assignment (CRA)**. The link resources are negotiated at the beginning of the transmission and are maintained for all the duration of the connection. Thus, this scheme does not require bandwidth consuming dynamic signalling from TTs to RNCC. However, since the resources are allocated even when not used, this method may lead to waste the yet scarce satellite bandwidth.




Figure 2: Example of message exchange between the BoD controller and a traffic terminal.

2. **Rate Based Dynamic Capacity (RBDC)**. In this DAMA scheme, a TT periodically submits to the RNCC a capacity request message based on measurement of the local incoming traffic. Every explicit request overrides the previously submitted one and new requests are submitted only if needed. RBDC has two parameters: the ceiling rate ($RBDC_{max}$) and the instantaneous rate. The ceiling rate is established during the call set-up using a method similar to CRA and, taking this threshold into account, the scheme tries to guarantee the requested rate. Moreover, a request remains effective as long as it has not timed out after a certain number of frames.

3. **Volume Based Dynamic Capacity (VBDC)**. In this scheme, the TT dynamically signals the data volume required to empty its buffer. New requests are sent any time more traffic is queued. The scheduler assigns the capacity according to these requests, while taking the constraints remaining after CRA and RBDC into account. Since the BoD controller cannot guarantee the requested bandwidth to the TTs, this access scheme is more suitable for best effort traffic flows.

A VBDC variant, known as *guaranteed VBDC*, has been proposed to guarantee a minimum amount of data per frame to each TT.

4. **Free Capacity Assignment (FCA)**. In this scheme, the capacity is assigned from the BoD controller to the TTs. Since capacity allocation does not involve any explicit demand from the TTs to the controller, FCA is designed to satisfy utilization and fairness criteria. Hence, this scheme improves the MAC scheduling performance especially in low load conditions and reduces the jitter in traffic flows.

2.1 Satellite Testbed Overview

Our system uses a combination of the CRA and RBDC resource allocation algorithms: for each TT connected to the network an amount of bandwidth equal to one time-slot per MF-TDMA multiframe is reserved.

Such slot can be used both to transmit signalling (as the terminal configuration information and allocation requests messages) as well as data messages. However, signalling packets have always higher priority than data packets.

In our system, the multiframe is made up of $N_f = 8$ frames, which contain $N_b = 6$ time-slots. Then, since every time-slot is composed by $N_s = 8$ cells (MPEG2 packets), a frame carries $N_s N_b = 48$ cells. Since the total duration of a multiframe is $T_f = 273.3\text{ms}$ and the size of each cell is 188 bytes, it is straightforward to compute the channel available capacity C_T and the time-slot bandwidth granularity C_S as

$$C_T = 1504 \frac{(N_s N_b) \cdot N_f}{T_f} = 2.112 \text{ Mbps} \quad (1)$$

$$C_S = \frac{C_T}{(N_f N_b)} = 44 \text{ Kbps} \quad (2)$$

Actually, the system can be configured also with a different number of time-slots per frame, provided that the number of cells per frame ($N_s N_b$) is 48. The BTP is delivered once per superframe (every 820 ms), which includes 3 multiframe. It is noteworthy that, in our system, all the four regions have a common RNCC traffic terminal, which is located in the area where the four spot beams overlap. Then, this TT brings together the functionalities of NCC and RNCC.

2.2 The BoD-controlled link model

In the following discussion, we assume a simplified version of the RBDC allocation algorithm as described in [7]. Periodically, as long as new packets are buffered into the TTs, a BoD agent in each terminal submits

a request to a BoD controller, specifying the amount of data per frame $r(k)$. Fig. 2 illustrates the dynamic of the BoD procedure. Each active TT sends a request every T seconds (the resource allocation period) and receives the response after L_s seconds (the system response time). The system response time includes the request/response propagation delay, the processing delay at the BoD controller and the *safe frame period* at the TT, which is required to ensure that all the TTs have received the BTP on time (fig. 2).

The BoD agent monitors the local incoming traffic and estimates the amount of data requested $r(k)$ using a smoothing filter

$$r(k) = \alpha x(k) + (1 - \alpha) r(k - 1) \quad (3)$$

where $x(k)$ is the amount of data seen by the terminal during the k -th allocation period. As long as the requested slots are less than the frame available space RBDC_{max} , the BoD controller is able to assign the bandwidth according to the following formula:

$$a(k) = \min\{\max\{r(k - L), \text{CIR}\}, \text{RBDC}_{max}\} + \varphi_F(k) \quad (4)$$

where the term $\varphi_F(k) \geq 0$ is a *noisy* term accounting for the discrete nature of slot allocation. Notice that the allocated bandwidth cannot be smaller than the threshold CIR, since our system assigns at least one slot to each active station.

In order to obtain a coarse model of the BoD algorithm, a continuous time version of (3) and (4) has to be adopted. Let us indicate with $r(t)$ the data rate requested and with $a(t)$ the rate assignment at time t . When the input rate is between CIR and the available bandwidth, we may assume that, after L_s seconds, the BoD system is able to provide the requested bandwidth,

$$a(t) = r(t - L_s) + \varphi_F(t). \quad (5)$$

On the other hand, the queue size evolution is governed by the differential version of Lyndley's recursion

$$\dot{q}(t) = x(t) - 1_{q(t)>0} \cdot a(t) \simeq x(t) - a(t) \quad (6)$$

where $x(t)$ is the instantaneous rate offered to the queue at time t and $1_{q(t)>0} \simeq 1$ corresponds to heavy load condition. Typically, in order to run after the rapid rate variations in the incoming traffic, the smoothing parameter α is kept near unity. Here, to simplify the analysis, we assume $\alpha = 1$, or equivalently $r(t) = x(t)$ from (3). Thus, substituting (5) in (6) and integrating, we have

$$q(t) = \int_{-\infty}^t [x(\tau) - x(\tau - L_s)] d\tau - \int_{-\infty}^t \varphi_F(\tau) d\tau \quad (7)$$

This expression allows to evaluate the RTT under the heuristic that packets are buffered at the TT (single bottleneck assumption) as follows

$$RTT(t) = T_0 + \frac{q(t-R)}{a(t-R)} \quad (8)$$

where T_0 is the sum of the propagation and transmission delay. Notice that we take into account the buffering at the TT by introducing a temporal shift equal to the average RTT R on the right hand side of the previous equation.

2.3 Single TCP Connection Modeling

The window scale option [8] has been introduced to cope with the low throughput in LFN (Long Fat Networks). Nevertheless, most of TCP implementations limits the advertised window size to the original value of 64 KB, without exploiting this additional capability. As a consequence, a single TCP connection is unable to *keep the pipe full*.

Thus, the cwnd evolution is well characterized by an initial slow start phase, with an exponential increase of cwnd up to the maximum value w_{max} , followed by a period where $\text{cwnd} = w_{max}$. In other words, a D seconds TCP connection can be described as

$$x(t) = \min\left\{\frac{2^{t/R+1}}{R}, \frac{w_{max}}{R}\right\} \cdot 1_{0 \leq t \leq D} \quad (9)$$

Note that in deriving this simple TCP rate profile, we have implicitly assumed that TCP connections suffer no losses and initial cwnd is equal to two packets. As it will be more clear in later sections, the absence of losses is reasonable since TCP is unable to overload the network link.




Figure 3: Comparison between experimental and theoretical UDP throughput rates.

3 Experimental Validation

Measurements were carried out over a bidirectional satellite link between two stations, referred to as A and B in the following, located in different sites in Pisa. The satellite link is accessed through a ground terminal (SKYPLEX¹) connected to a 90cm satellite dish antenna directed towards a transponder of the geostationary satellite HotBird 6TM operating in Ka band.

The system was configured with a resource allocation period of 400ms and the value of the system response time was about 3sec. During our measurements, the number N_t of active TTs was set to 18, but only two of them were involved in the monitored unicast connection. Thus, since one time-slot was allocated to each active TT, the maximum available bandwidth was

$$RBDC_{max} = C_t - (N_t - 1)C_s = 1364 \text{ Kbps} \quad (10)$$

Our measurements aimed to collect end-to-end IP traffic statistics to characterize the satellite link behavior with UDP and TCP flows. To this purpose (see Fig. 1), each site in Pisa was equipped with a Linux PC, connected to the TT through an Ethernet 10/100 LAN.

3.1 UDP Traffic Measurements

In the first set of measurements, using BRUTE [9], we generated a CBR UDP traffic flow. The receiver, by means of the Linux Netfilter tool, swapped the IP source and destination fields and sent back the packet. Moreover, in order to measure the RTT, the drop rate and the link capacity, the transmitter stored the relevant information of each packet through a robust and low-invasive Linux kernel analysis tool [10]. Several tests were performed considering different rates and packet sizes. Table 1 shows the statistics of the RTT experienced by the packets, when the traffic generator was located either at site A or at site B.

As expected, the satellite link is highly symmetrical in delay and the RTT is only slightly affected by the packet size. Moreover, the RTT variations are rather small as compared to the mean value (about 2%) and the estimated packet drop rate is lower than 10^{-4} . These results confirm our previous assumptions that the system introduces a negligible delay jitter and may be considered loss free.

To evaluate the behavior of the resource allocation scheme when the offered traffic is progressively increased, the input rate of the UDP flow was raised in steps of 200Kbps every 50 seconds. Fig. 3 reports a comparison between the experimental and theoretical throughput when the packet size is 1024 bytes. Similar results were observed in both directions also for 128, 256 and 512 bytes packets; this suggests that the behavior weakly depends on the packet size. The theoretical throughput was calculated by solving the system of equations (5), (6) and (7) for $a(t)$ with a step rate as input $x(t)$ and assigning one half of bandwidth granularity to φ_F term, which is equal to $C_s/2$ from (2).

¹SKYPLEX is a commercial device mainly used for digital video streaming and interactive multimedia services.




Figure 4: Comparison between theoretical and experimental RTT.

As expected, the satellite connection was able to satisfy the bandwidth requests up to $RBDC_{max}$. However, we notice an overshoot in the received throughput just after each rate increment. This behavior indicates that, just after a sudden change in the offered rate, the system provides to the flow more bandwidth than necessary. Equation (5) confirms this experimental evidence: indeed, due to the φ_F factor, the allocated bandwidth can be higher than the requested one.

The same concept is highlighted by the temporal evolution of the RTT (fig. 4): when the input rate changes, the RTT drastically increases, reaching values several times greater than the steady-state RTT. In particular, the height of the RTT peaks is inversely proportional to the rate. Once again this behavior is confirmed by the proposed model: since the rate increment in (7) is the same at each step, the queue fills up to the same level and so the lower is the rate, the higher is the delay.

When the input rate is larger than $RBDC_{max}$, the outgoing buffer of the transmitting TT fills up and the in excess packets are dropped. Therefore, in this conditions, we can evaluate the buffer size as the ratio between the maximum delay and the service time. From our measurements we deduced a value of about 1500 packets; a so large buffer is useful in reducing the transient drop rate, but may introduce very large delay for real-time traffic.

3.2 TCP Traffic Measurements

This section reports the results of the experimental tests performed by using netperf [11] when TCP SACK [12] is used. Other TCP versions (i.e. TCP Westwood [13]) were also considered, but only marginal difference in the results were highlighted. In the experiments, each packet header was dumped at the sender side by means of tcpdump in order to keep track of the TCP parameters dynamics. To evaluate the temporal evolution of the average cwnd and RTT (see fig. 5 and 6 respectively) for a single TCP connection, we generated a periodic traffic flow, scheduling every 60 seconds a new TCP connection carrying 592KB.

As shown in figure 5, after the initial exponential rate increase due to the slow start phase, the mean cwnd gets stuck at the size advertised by receiver (32KB in this test). Since no loss occurs, the TCP never exit from the initial slow start phase and the slow start threshold remains unset. Note that the measured TCP rate fits well the profile (9) described by our model, as long as the term R is substituted by the mean measured RTT. Thus, as already pointed out, the small throughput (250Kbps) is due to the limited value of the cwnd.

Table 1: Mean and std. dev. of the RTT between the two hosts A and B

Packet Size (bytes)	direction	mean RTT (ms)	std. dev. (ms)
1024	A → B → A	607.71	13.28
	B → A → B	596.50	17.69
512	A → B → A	601.35	12.40
	B → A → B	590.07	15.19




Figure 5: Mean cwnd evolution of 34 TCP connections in a 2000 seconds overall duration test.

The instantaneous RTT was computed as the difference between the time an ACK was received and the corresponding packet transmission time. Figure 6 shows the RTT evolution of all the 34 TCP connections and the results obtained with our RTT model when the real TCP traffic traces are used (bottom plot). The graph highlights that BoD reservation scheme affects the behavior of a single TCP connection by delaying the initial RTT. This occurrence, confirmed by the simulations, slows down the TCP startup phase, strongly penalizing short-lived connections.

4 Conclusions

The paper aimed at evaluating the behavior of a DVB-RCS satellite network with respect to TCP/IP traffic. Several traffic measurements sessions have been performed between two traffic terminals located in Pisa, and an analytical model for the BoD access scheme has been proposed. In a first set of tests, we used CBR UDP traffic flows to infer the characteristics of the satellite link and validate our model. Then, we used several TCP traffic flows, periodically scheduled, to estimate the dynamic behavior of the cwnd and the RTT for a single TCP connection. In both cases, our analytical model matches with the experimental behavior in terms of RTT and throughput.

Our analysis clearly highlights some issues related to the use of satellite connections. First, the large delay variations at the beginning of a connection or during a period where a sharp rate transition occurs determine long delays and performance degradations. In other words, the BoD controlled links are unable of quickly reacting to sudden rate variations and result in poor performance particularly at low load levels. For instance, in our measurements, we observed that, at the connection startup, TCP packets may experiment nearly two times the round trip time (RTT) observed during the steady-state phase. Such longer initial queuing delays must be attributed, as explained by the model, to the deferred reservation of bandwidth induced by the demand/assignment mechanisms.

A further problem to face with TCP connection is the limitation of the cwnd enforced by the advertised window, used to avoid receiver buffer overflows, which allows a 64KB congestion window at most. Such maximum window size is insufficient for the high bandwidth-delay product of long high-speed connections, such as satellite connections. Indeed, according to our measurements, a single TCP connection achieves just about one-half of available bandwidth. A trivial solution to the advertised-window problem, would be to enable and properly configure the window scale option in TCP header. However, TCP does not know a priori the characteristic of a communication path and the correct selection of window scale both at transmitting and receiving side can be difficult. As an alternative, the use of multiple parallel TCP connection and split connection techniques could be the right direction. While for the former technique the behavior of multiple socket congestion control has to be investigated, for the latter the problems inherent to breaking the paradigm of end-to-end connection should be overcome.




Figure 6: Experimental evolution of experimental and RTT

References

- [1] E.W. Fulp and D.S Reeves. On-line Dynamic Bandwidth Allocation. In *in proceedings of INFOCOM*, 1996.
- [2] Y. Afek, M. Cohen, E. Haalman and Y. Mansour. Dynamic Bandwidth Allocation Policies. In *in proceedings of INFOCOM*, 1996.
- [3] J.Neale and A. Mohsen. Impact of CF-DAMA on TCP via satellite performance. In *proc. IEEE GLOBE-COM 2001*, November 2001.
- [4] M.Sooriyabandara and G.Fairhurst. Dynamics of tcp via satellite networks. *Int. Jour. of Satellite Communications and Networking*, July 2003.
- [5] M. Karaliopoulos, R. Tafazolli, and B. G. Evans. Providing Differentiated Service to TCP Flows Over Bandwidth on Demand Geostationary Satellite Networks. *Int. Jour. On Selected Areas In Communications*, Feb 2004.
- [6] ETSI EN 301 790 v1.2.2. Digital Video Broadcasting (DVB): Interaction channel for satellite distribution system. December 2000.
- [7] Eros Feltrin PhD. thesis. Design optimization and test of an on-bord processing-based satellite network. University of Padova, 2004.
- [8] R. Fox. RFC 1106: TCP Big Window and NACK options.
- [9] N.Bonelli, S.Giordano, G.Procissi and R.Secchi . BRUTE: A High Performance and Extensible Traffic Generator. In *proceedings of SPECTS'05*, July 2005.
- [10] Wu-Chun Feng, Mark K. Gardner, and Jeffrey R. Hay. The magnet toolkit: Design, implementation and evaluation. *J. Supercomput.*, 23(1):67–79, 2002.
- [11] Helvet-Packard company. Netperf. <http://www.netperf.org/netperf/NetperfPage.html>.
- [12] M. Mathis, J. Mahdavi, S. Floyd and A. Romanow. RFC 2018: TCP Selective Acknowledgment Options, October 1996.
- [13] S. Mascolo, C. Casetti, M. Gerla, M. Sanadidi and R. Wang. TCP Westwood: Bandwidth Estimation for Enhanced Transport over Wireless Links. In *proceedings of the 7th annual international conference on Mobile computing and networking*, pages 305–317. ACM Press, 2001.