Flexible QoS Support in DVB-RCS2

Ziaul Hossain, Arjuna Sathiaseelan, Raffaello Secchi, and Gorry Fairhurst

Electronics Research Group, School of Engineering, University of Aberdeen,
Fraser Noble Building, King's College, Aberdeen AB24 3UE, UK
{ziaul.hossain,a.sathiaseelan,r.secchi,g.fairhurst}@abdn.ac.uk

Abstract. Satellite is considered a vital technology for enabling ubiquitous access to broadband services in many countries. This paper explores provision of IP-based broadband satellite access with Quality of Service (QoS). It analyses a set of example scenarios, based on the recently published DVB-RCS2 standard involving web and voice traffic. It specifically explores the interaction between Bandwidth-on-Demand (BoD) and QoS, showing that this interaction offers the flexibility required for satellite Internet service operators to manage the bandwidth of broadband users in a multi-service access network.

Keywords: QoS, BoD, DVB-RCS2, VoIP.

1   Introduction

Satellite systems can provide broadband access that is not dependent on the cable infrastructure. The advances in satellite technology in recent years have significantly increased the capacity available to customers, enabling satellite to become a strong contender for providing fast broadband to places that cannot be covered by other technologies due to physical, technical or economical limitations.

The second generation Digital Video Broadcast Return Channel Satellite (DVB-RCS2) [1] is the only open IP-based next-generation multivendor standard for broadband satellite access. This specifies IP QoS functions at layer 3 that can ensure appropriate traffic classification and prioritisation, essential for a satellite Internet Service Provider (ISP) to guarantee application performance. Bandwidth-on-Demand (BoD) methods, defined at Layer 2, are vital components to maximise the capacity utilization of satellite resources. The Quality of Service (QoS) architecture in DVB-RCS2 is sufficiently flexible to implement policies for specific classes of traffic/application and can be associated with different BoD methods.

QoS support in satellite networks is an important topic that has been explored extensively in the past. Examples of frameworks to support QoS and bandwidth allocation are provided in [2] [3] considering the first generation standard DVB-RCS. [4] and [5] analysed the behaviour of applications: web, VoIP, bulk file transfer etc. with BoD methods. Applications differ in their QoS requirements and a possible mapping of applications to BoD methods is presented in [6].

This paper examines the interaction between the Layer 2 and Layer 3 and implications of BoD methods on application performance through simulations. Our results illustrate the flexibility offered in DVB-RCS2 by tuning the BoD methods to support multiple network services, each of which can be customised to meet the

requirements of a particular traffic class. The analysis considers both VoIP and web traffic separately and simultaneously.

2 BoD in DVB-RCS2

BoD allows multiple Return Channel Satellite Terminals (RCSTs) to share Return Link (RL) satellite capacity. RCSTs transmit explicit Capacity Requests (CRs) to the Network Control Centre (NCC) in order to notify their instantaneous capacity requirements. The NCC uses the CRs to generate a burst-time plan (BTP) that contains RCSTs allocations. The BTP is broadcast periodically on the Forward Link (FL).

DVB-RCS2 permits an RCST to use a range of BoD mechanisms to request capacity, including Continuous Rate Assignment (CRA), Rate Based Dynamic Capacity (RBDC), Volume Based Dynamic Capacity (VBDC), Absolute VBDC (AVBDC) or Free Capacity Assignment (FCA). Some of these depend on the CRs made by the terminals while some of them are based on the Service-Level-Agreement (SLA) between the RCST and NCC.

Table 1. BoD methods

Method	Request Type	Set by
CRA	Rate	SLA or Control plane
RBDC	Rate	Explicit Capacity Request
VBDC	Volume	Explicit Capacity Request
AVBDC	Volume	Explicit Capacity Request
FCA	Volume	SLA or Control plane

These methods may be used alone or in combinations to form the basis of a bandwidth allocation strategy, known as a Request Class (RC). In DVB-RCS2, RCs can be mapped to particular traffic classes to satisfy the capacity requirement of a specific type of traffic. When multiple mechanisms are combined in an RC, DVB-RCS2 mandates not to request more than 110% of the resources needed [1].

3 Simulation Setup and Traffic Description

Simulations were performed to study the interaction between BoD and QoS for a set of services. Our DVB-RCS2 simulator [7] implemented IP-based QoS with priority queuing. An RCS Terminal (RCST) was configured with traffic classes to differentiate VoIP and Web flows using the IP Diffserv code-point assigned by the application.

Our experiments considered two key request methods: RBDC and VBDC. An RBDC Capacity Request (CR) reflected the rate of incoming traffic, averaged over the past request period. The following formula was used in the simulator to calculate the RBDC requested rate (R_{now}) in the terminal:

$$R_{now} = \alpha \times R_{prev} + (1 - \alpha) \times X \tag{1}$$

where α is the smoothing factor between 0 and 1, R_{prev} is the capacity requested in the previous request interval, X is the incoming traffic rate. This method (with a high α value) prevents the system against the high fluctuations of the RBDC burstiness which was described in [14]. A VBDC CR reflected the estimated volume of incoming traffic in an allocation period. This included all the capacity requested and not allocated at the time the CR was issued.

We considered two types of service with different traffic patterns: web and VoIP. Our web model [8] simulated a web page download comprising several Hyper-Text-Transfer-Protocol (HTTP) request/response exchanges. Each web page consisted of several objects distributed on multiple servers. The HTTP client opens multiple parallel connections to the servers and reuses a connection for multiple object requests. VoIP traffic consists of fixed-size packets of 76 B (with TCP/IP headers) and constant inter-transmission time of 40 ms. The voice encoder generates data with a media rate of 8 kb/s resulting in a constant flow of 15.2 kb/s including TCP/IP overhead.

Table 2. Key Simulation Parameters

Network	
Internet delay (sat. gateway and servers)	50ms
FL capacity	2 Mb/s
FL encapsulation	GSE
Max RL capacity	512 kb/s
RL encapsulation	RLE
Request interval	100 ms
Superframe period (one frame/superframe)	26.5ms
Min RBDC request	16 kb/s, 32 kb/s
Web page model	
Number of servers	4
Number of objects	40
Obj. avg. size	7 KB
Obj. size distribution	Pareto; shape=1.2
Max no. of persistent connections	6
HTTP request size	320 B
VoIP Model	
Packet size (IP/RTP/media)	76B
Inter-transmission time (without VAD)	40ms
Traffic rate	15.2 kb/s
VoIP/VAD Model	
Packet size (IP/RTP/media)	76B
Inter-transmission time (without VAD)	40ms
Peak Traffic rate	15.2 kb/s
Average codec rate	8 kb/s
Avg. burst duration (exp. distributed)	0.46 s
Avg. idle duration (exp. distributed)	0.54 s

When VAD was enabled, voice activity was modeled as a two state ON/OFF Markov chain [9]. The generation of comfort noise packets was ignored, since these packets are small and infrequent and produce negligible effects on the return link. The packet inter-arrival time and VAD design follows guidance in G.729 Annex B [10].

The performance of HTTP traffic was evaluated by performing 500 simulation runs and plotting the 90 percentile of the transfer duration (the time between issuing an HTTP request and download all the requested objects). The performance of VoIP traffic was evaluated by calculating a Mean Opinion Score (MOS) derived from the R-score [11] using an empirical mapping. This provided an objective perceived-quality measurement with a range from 1 to 5, where 1 is the poorest perceived quality and 5 the best. The allocation efficiency provided a metric for the effectiveness of the requestor/allocator in predicting the capacity required for the traffic. This was the ratio between the used slots and the total allocated slots.

Table 2 identifies the key simulation parameters. There are several additional delays with the request-allocation mechanism besides the satellite Round-Trip-Time (RTT), such as CR submission delay, NCC processing delay, waiting time for the allocated superframe, etc. For simplicity, only the effect of RTT was introduced between a request and the corresponding allocation and it was assumed that the NCC was always able to grant the requested amount. There was no L2 congestion at the satellite Network Control Centre, NCC, (i.e., capacity was always available).

4   Experiments and Results

This section investigates the suitability of specific RCs for web and VoIP traffic. The experiments demonstrate that the performance of different traffic classes is significantly impacted by the choice of RC parameters.

4.1   Web Performance Tuning with Different RCs

Web flow performance was evaluated across a range of RCs. We define RC1 using RBDC and VBDC, where the estimated input rate and queue size in the RBDC and VBDC CR are multiplied by p and $(1-p)$ respectively. Thus, the total amount of capacity requested per request period (in bytes) C_{REQ} is:

$$C_{REQ} = \quad p \times C_{RBDC} \times RP + (1-p) \times C_{VBDC} \qquad (2)$$

where C_{RBDC} (B/s) is the RBDC request calculated according to (1), RP (sec.) is the request period, and C_{VBDC} (bytes) is the VBDC request as described in Section 3. This way, C_{REQ} remains below the 110% threshold described in Section 2. Simulations varied p in the range [0,1]. When $p=1$ and $p=0$, RC1 respectively corresponds to RBDC and VBDC. Fig. 1 shows the page download time and efficiency versus the number of objects on the webpage. The results confirm the findings in [3]: An increase of p increases the allocation efficiency, but also increases the completion time. Indeed, as p varies between 0 and 1, the completion time curve moves upwards maintaining a quasi linear behaviour. The case $p=0.2$ can be regarded as a good compromise between performance and efficiency for web.
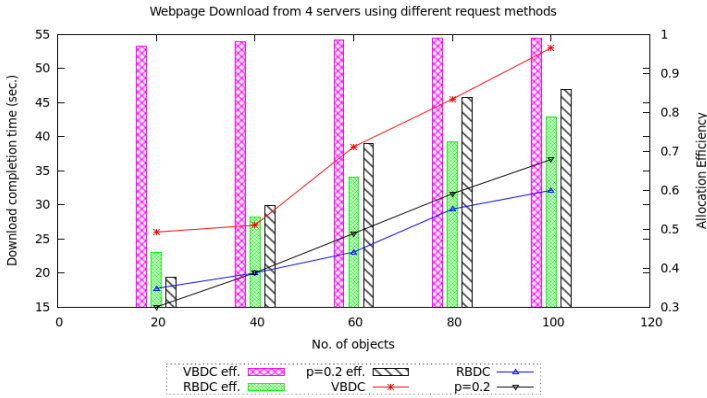
Fig. 1. Webpage download time as a function of the no. of objects in the webpage; download time and allocation efficiency increases with the no. of objects and they depend on the 'p' value

This analysis suggests that a satellite ISP can select a suitable RC for any traffic by varying the parameter to weight the different CR components. This allows an operator to trade efficiency and performance.

4.2   RC Tuning for VoIP Traffic

VoIP (not using VAD) flows result in constant bitrate traffic (Table 2).   Hence, an RC using the simple RBDC method (defined by Equation 1) can support this kind of traffic and it was named RC2 in our simulations. In contrast, VoIP flows that use VAD, result in a traffic flow that continuously varies its rate. RC2 requests reflect the average traffic rate. These were unable to make appropriate CRs for this variable traffic and resulted in reduced performance.

One way to mitigate this problem would be to use a request/allocation method that predicts the capacity using an algorithm tailored to the VoIP/VAD traffic characteristics. We defined RC3 so that it allocates extra capacity over the average rate to accommodate the burstiness of the VoIP/VAD traffic. RC3 is implemented by measuring the input rate each second and using a high smoothing factor ($\alpha = 0.9$ in equation (1)) for RBDC. The NCC allocates capacity to the RCSTs by multiplying RC3 CRs with a constant factor.

Fig. 2 plots the MOS against the allocated capacity divided by the average rate for combinations of VoIP flows. This shows that good performance (around MOS 3) can be achieved for VoIP (with VAD enabled) traffic when the allocation is substantially higher than the average rate. It can be noted that multiple flows of VoIP with VAD requires less allocation on average than a single flow to achieve a good performance level. Multiplexing of several flows reduces the rate variation in the traffic [12], reducing the need for extra capacity to compensate for rate fluctuations. RC3 outperformed RC2 in the presence of multiple flows. It can also be noted that increasing the ratio of allocated capacity to average rate beyond a certain point does not offer any better performance than a certain level. This defines the upper-bound of the performance that can be observed in satellite network. For example, 1.5 for 4 VoIP flows - allocating more than this not bring any significant performance improvement.
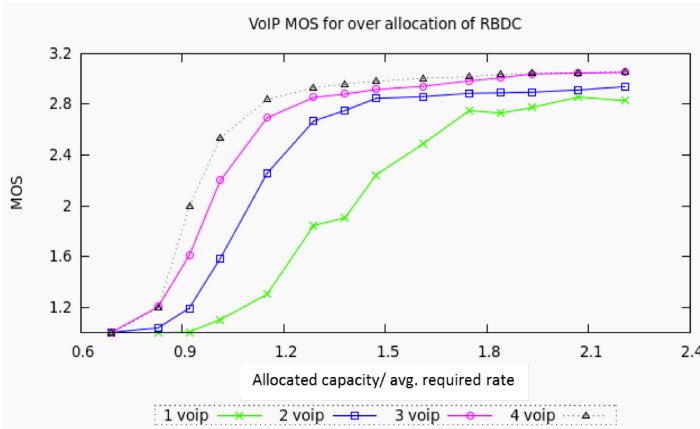
Fig. 2. MOS for VoIP traffic as a function of the allocated capacity per flow normalised to the flow peak rate assigned using RC3

4.3   Simultaneous Multiple Web and VoIP Flows

The previous sections illustrated how the performance of an application depends on the strategy implemented by the RC. In this section, we focus on the performance of a traffic mix in a multi-service scenario. In particular, we show that performance and efficiency gains are possible when different traffic share the allocated capacity.

Simulations considered traffic with a varying number of parallel running web and VAD-enabled VoIP flows. Web and VoIP flows were respectively assigned to RC1 and RC3. Moreover, the VoIP traffic was given priority over web, as recommended in [13]. Fig. 3 shows the webpage download time in this mixed traffic scenario as a function of the number of competing web and VoIP flows. The webpage download
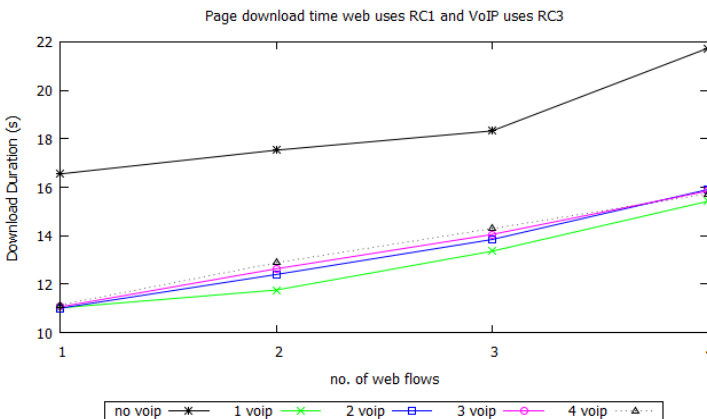

Fig. 3. Excess capacity from VoIP flows benefits web flows when downloading web pages

time in absence of VoIP (labelled 'no voip') ranged from 16 to 22 seconds as the number of web flows varied between 1 and 4. The download time increased almost linearly as the number of web flows increased. Indeed, the burstiness of web traffic caused a transient backlog of ACKs at the RCST buffer, which increased as the number of web flows increased.

Fig. 3 shows a significant reduction in the latency of webpage download time - independent of the number of web flows. The introduction of competing VoIP traffic had a beneficial effect on web performance. It takes less time to download a page when at least 1 VoIP is present. Since each VoIP/VAD flow had an ON/OFF behaviour, the VoIP aggregate left some of the allocated capacity unused. Web traffic could use this capacity to clear any backlog at the RCST, avoiding the need for additional CRs. This reduced the time for the download. The addition of more VoIP flows had only a small impact of web performance.

RC3 was suited to support multiple VoIP flows also in a mixed traffic scenario. Fig. 4 shows that the MOS was fair (around 3) and it did not deteriorate significantly when the number of VoIP flows increased.
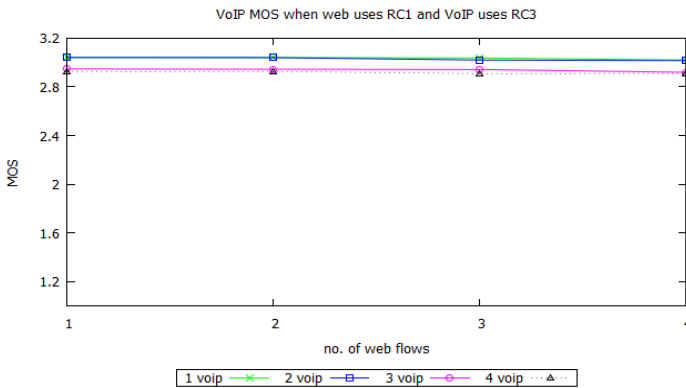

Fig. 4. VoIP MOS for multiple flows with multiple Web flows; the web flows have negligible effect on the VoIP performance level

We also report the efficiency of a broadband resource sharing scheme. Resource sharing is an important metric for commercial operation. This metric (Fig. 5) can be used to determine whether the resource distribution can be improved. Any unused capacity may be made available to support additional users or new services. The efficiency for one or two VoIP (with VAD enabled) flows was low (around 25% and 50% respectively). This poor efficiency was dominated by predictive over-allocation, required to ensure VoIP flows did not experience jitter, but exceeded the average capacity required for a few VoIP flows. Efficiency increased with the number of flows (VoIP or web).
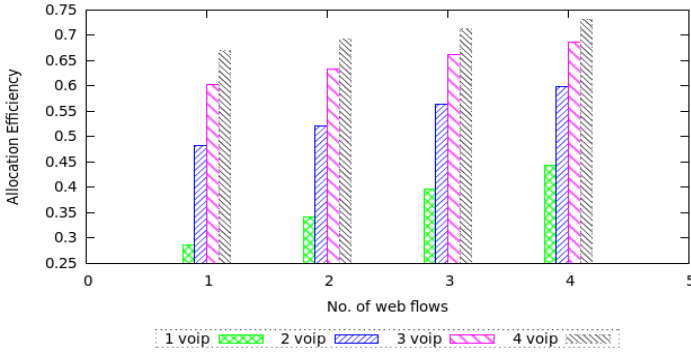
Fig. 5. Allocation efficiency as a function of number of simultaneous web and VoIP flows; higher number of flows contribute to better efficiency

5   Discussion

Resource sharing is important in all types of commercial access networks. In cabled services, an operator typically needs to consider contention in the backhaul since all downstream users share the resource of this backhaul. To avoid degrading the user performance, the decision to add a new user must consider whether the backhaul resources will become saturated. Contention ratios of 10:1, 20:1 or 50:1 are not uncommon, with smaller ratios for premium services.

A satellite access network has a similar need to share the resource of the return link capacity. However, the resource pool is typically much larger than that of cabled systems, where a satellite operator may share available resource among a large number (~1000s) of active users. The larger pool size allows an operator more flexibility in the way that the resources are assigned. Appropriate assignment can then be used as a tool to ensure economic system operation – assigning more capacity than needed incurs unnecessary cost to an ISP, allocating less than needed impacts the user performance.

The first two sets of experiments provided examples of ways to tune the performance and efficiency by choosing an RC with appropriate BoD mechanisms. RCs can be further customized to satisfy the service needs of a specific application.

The third set of experiments presented the performance and system efficiency when used to support multiple concurrent web and VoIP sessions, each supported by the RC best suited to the corresponding application. These results illustrate the advantage of selecting appropriate BoD methods and QoS parameters. We did not find a single optimum RC that could support the entire range of traffic in a multi-service network. Different applications have different traffic profile and delay requirements (Table 3).

It was also shown that a combination of RCs can provide the flexibility needed for bandwidth management by allowing an operator to tune the BoD to achieve a particular set of QoS classes (i.e. to meet a customer service level agreement –SLA). RCs can be designed combining RBDC and VBDC requests with different weights to meet these varying requirements.

Table 3. Traffic profile and delay requirements for different applications

| Application	Traffic Profile	Delay Requirements
Web	Variable rate (Bursty)	Moderately delay sensitive
VoIP	Fixed rate	Real-time
VoIP/VAD	ON/OFF between 0 and fixed rate	Real-time
Video-conferencing	Fixed rate (based on encoder)	Real-time
Web Video	Variable rate/ Bursty	Low delay and jitter
Instant messaging	Very short interactive transfer	Medium delay, low jitter
File Transfer	Bulk transfer	Medium delay
Email	Short transaction	Medium delay
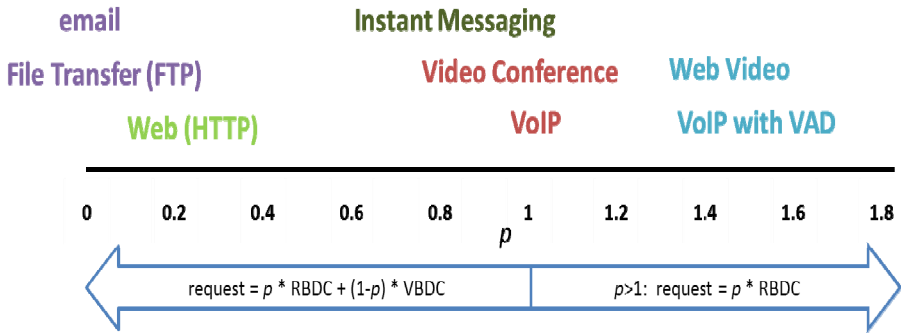
Fig. 6. A possible mapping of different applications to different RCs varying the value '*p*'

Fig. 6 shows a possible mapping of different applications to RCs. In our experiments, we considered two different request strategies: one combines RBDC and VBDC requests with a weighted value '*p*' (*p*<1), while the other multiplies the RBDC requested amount with '*p*' for allocation (*p*>1). Our results showed that application performance (for web and VoIP) varied greatly on the values of '*p*' and an RC can be designed with a particular value of '*p*' matching the application requirement. Other applications can also use the same allocation strategy, but certainly requires a suitable '*p*' value. Based on the traffic profile and delay requirements, an RCST running different applications may possibly follow the mapping in Fig. 6. However, the satellite operator has the full flexibility in selecting the appropriate value of '*p*' for different applications based on the factors like SLA, number of users etc.

6   Conclusion

This paper explored the interaction between BoD and QoS for multi-service support in an IP-based DVB-RCS2 system. Simulation results evaluated BoD methods for tuning the QoS. A combination of RCs can provide the flexibility needed for

bandwidth management by allowing an operator to tune the BoD to achieve the set of QoS classes. This offers functions for the satellite that are similar to those used to provide a cabled broadband service. As a result, satellite broadband can be an effective complimentary access technology to enable universal broadband roll-out.

References

1. ETSI EN 101 545, Second Generation DVB Interactive Satellite System; Part 2: Lower Layers for Satellite Broadband (March 2011)
2. Iuoras, N., et al.: An IP-based satellite communications system architecture for interactive multimedia services. International Journal of Satellite Communications and Networking 21(4-5), 401–426 (2003)
3. Skinnemoen, H., et al.: VoIP over DVB-RCS with QoS and Bandwidth on Demand. IEEE Wireless Communications 12(5), 46–53 (2005)
4. Lexow, H., Navekvien, T., Paxal, V.: Satellite resource management combining distributed cooperative methods and central control using BoD variants of DVB-RCS. In: 11th Ka-Band Communications Conference, Rome (September 2005)
5. Kalama, M., et al.: VoIP over DVB-RCS satellite systems: a novel capacity request mechanism for improved voice activity detection. In: IEEE Vehicular Technology Conference, Singapore, pp. 2957–2961 (May 2008)
6. Chini, P., Giambene, G., Hadzic, S.: Broadband satellite multimedia networks. In: Cranley, N., Murphy, L. (eds.) Wireless Multimedia: Quality of Service and Solutions, pp. 377–398. IGI Global (2009)
7. Secchi, R.: DVB-RCS(2) for ns2 (September 11, 2012), `http://homepages.abdn.ac.uk/r.secchi/pages/dvbrcs_ns2.htm`
8. Secchi, R., Sathiaseelan, A., Fairhurst, G.: Evaluating web traffic performance over DVB-RCS2. In: Pillai, P., Shorey, R., Ferro, E. (eds.) PSATS 2012. LNICST, vol. 52, pp. 148–155. Springer, Heidelberg (2013)
9. Adas, A.: Traffic Models in Broadband Networks. IEEE Communications Magazine, 82–89 (July 1997)
10. ITU-T: G.729 Annex B: A silence compression scheme for G.729 optimized for terminals conforming to Recommendation (October 1996)
11. Cole, R., Rosenbluth, J.: Voice over IP performance monitoring. ACM Computer Comm. Review 31(2), 9–24 (2001)
12. Castro, V., et al.: Quality of Service of VoIP over DVB-RCS. In: 6th Workshop on Signal Processing in Communications, Baiona (September 2003)
13. Sathiaseelan, A., Fairhurst, G.: Multimedia congestion control for broadband wireless networks. IST Mobile Summit, Budapest (2007)
14. Adami, D., Giordano, S., Pagano, M., Secchi, R.: Modeling the Behavior of a DVB-RCS Satellite Network: an Empirical Validation. In: Proceedings of Performance Modelling and Evaluation of Heterogeneous Networks (HET-NET's 2005), Ilkley, UK (July 2005)