Improving TCP Friendly Rate Control Behaviour in DVB Satellite Networks[1]

Raffaello.Secchi[(*)], Arjuna Sathiaseelan[(o)], Francesco Potortì[(*)], Alberto Gotta[(*)], Gorry Fairhurst[(o)]

[(*)]CNR-ISTI, Via Moruzzi 1, 56124 Pisa (Italy)
<Raffaello.Secchi,Potorti, Alberto.Gotta>@isti.cnr.it

[(o)]Electronics Research Group,
University of Aberdeen, United Kingdom
<arjuna,gorry>@erg.abdn.ac.uk

Abstract

In the not so distant future we envisage an Internet where the biggest share of capacity is used by streaming applications, with no QoS reservations in the network core. Such traffic calls for a robust and ubiquitous end-to-end multimedia congestion control mechanism such as TCP Friendly Rate Control (TFRC), that provides fair sharing with the other Internet traffic. We also suggest satellite links will be more common than today, utilising Demand Allocation Multiple Access (DAMA) to maximise satellite transponder utilisation when providing triple-play services. The interaction of TFRC with DAMA therefore needs to be investigated. Simulations show that DAMA slows down the start-up phase of TFRC, causing non-negligible delays. To mitigate this problem, we propose judicious use of the experimental Quick-Start (QS) mechanism.

Keywords: multimedia; TFRC; satellites; DAMA; Quick Start

1 Introduction

The advent of popular services like YouTube and Google Video, are evidence of a strong push towards Internet television, and the momentum that streaming p2p is gaining [1] makes us envision a future where the majority of Internet traffic is produced by streaming applications. Currently, these streaming applications have limited impact on the Internet infrastructure. They do not need a Quality of Service (QoS) reservation infrastructure, and do not yet pose a congestion problem, whereas the emerging commercial TV over IP (TVoIP) have been typically provided using strong QoS within a single Internet service provider network. We expect video streaming to become the most bandwidth demanding Internet application, with high-quality video streaming gaining a significant place in the next generation Internet. In this future network we suggest that hard QoS reservation will continue to be absent in the network core. In such a scenario, network over-provisioning will be a requirement – as streaming without reservations can only work in the absence of congestion – but nonetheless congestion control will be mandatory for granting overall network stability. Consequently, it is likely that congestion control will become important for streaming applications. One appropriate congestion control method for use in the general Internet is TFRC (TCP Friendly Rate Control), an equation-based algorithm that is designed to be fair to streaming applications running over it and to concurrent TCP flows traversing the same routers [15]. TFRC is also supported as congestion control in the newly standardised Datagram Congestion Control Protocol) (DCCP) [14].

At the same time, we envisage that the current success and the continued development of the DVB (Digital Video Broadcasting) standard for packet communications over satellite [16,18] will make satellite equipment interoperability a reality, and the deployment of small and large satellite networks cheaper and consequently more common than it is now. DVB-RCS (DVB with Return Channel by Satellite) [17] or other bidirectional satellite communication standards will make both private and public satellite networks affordable for generic Internet access, through the use of dynamic bandwidth allocation schemes on the satellite segments. DAMA (Demand Allocation Multiple Access) will be necessary both for better exploitation of the expensive shared satellite bandwidth and as a means of implementing BoD (Bandwidth on Demand) schemes on satellite networks. This mix of technologies is the basis for the provision of triple-play services on satellite networks [21].

Prediction is difficult, especially about the future, for this reason we do not mean to make forecasts, but only to analyse a possible future scenario where unreserved streaming traffic on the Internet is common. In such a scenario, it is important to assess the performance of TFRC-controlled streaming applications traversing a path with links whose bandwidth is varying under the control of a DAMA dynamic bandwidth allocation scheme. The start-up phase of TCP is significantly slowed when the TCP connection uses a satellite link controlled by DAMA, with respect to a satellite link with a fixed bandwidth allocation [3]. The problem is caused by the interaction of DAMA with the slow-start algorithm. Since this algorithm is also used by TFRC, we expect a degradation in the performance of TFRC streaming applications during the start-up phase on DAMA satellite links. The QS (Quick-Start) mechanism [7,10] significantly

[1] Work supported by the European Commission under the European Satellite Communications NoE (SatNEx II, IST-27393) within the 6th Research Framework Programme.

mitigates the problem when running TFRC over a fixed allocation satellite link [4].

In this paper we couple the QS mechanism with a satellite DAMA system and show that careful integration of QS with DAMA can offset the performance degradation caused by the interaction of TFRC with DAMA. We first consider the performance of a TFRC streaming application over a satellite link that does not deploy DAMA, with and without use of QS. This is what a TFRC-enabled streaming flow would experience today when traversing a satellite link. We then show that managing the satellite link with DAMA significantly worsens the streaming application performance, but a careful implementation of QS can almost completely offset this performance loss. We consider the usage of TFRC as implemented by the sending application directly, that is, without the use of DCCP. Extension of this work to consider the implications of using DCCP may be the subject of further research.

Performance of streaming applications is measured in terms of median and maximum one-way delay [9], measured from the sending host to the receiving host. While variable one-way delay is absorbed at the receiver using a de-jitter buffer, completely absorbing this delay implies a total delay at the decoder that is greater than the maximum one-way delay. It is possible to trade total delay at the decoder with packet loss, thus reducing the delay in exchange for a small packet loss rate. However, TFRC-enabled streaming applications running on DAMA experience the maximum delay only during the first few seconds of transmission, exactly the most important ones for content zapping, so reducing delay by sacrificing video quality in terms of packet loss may not be an option: our opinion is that minimising the maximum delay during start-up is highly desirable for the application.

In Section 2 we describe and comment on the technologies in our scenario: DVB, DVB-RCS, DAMA, Skyplex Data, TFRC and QS. In Section 3 we compare by simulation how TFRC over a fixed-allocation satellite network fares with respect to TFRC on a DAMA system. This comparison considers cases with and without QS, for a single TFRC streaming flow. Section 4 presents the scenario, where several TFRC flows of different types traverse a DAMA system. We analyse different methods of interaction between DAMA and QS and compare their relative merits. The conclusions describe the implications of our work and identify possible future developments of this analysis.

2 Technology overview

This section presents a brief overview of the main technologies involved in the scenario described in the introduction.

2.1 DVB Architecture overview

DVB (Digital Video Broadcasting) is a mature technology used to send packet data over wireless media, which is particularly popular in the DVB-S (DVB for satellite) format. IP packets are first encapsulated by MPE (Multi Protocol Encapsulation) [12] or using ULE (Unidirectional Lightweight Encapsulation) [13]. This encapsulation also adds the MAC address of the destination satellite earth terminal. The encapsulated data is then delivered using a transport stream of 188-bytes MPEG-2 cells.

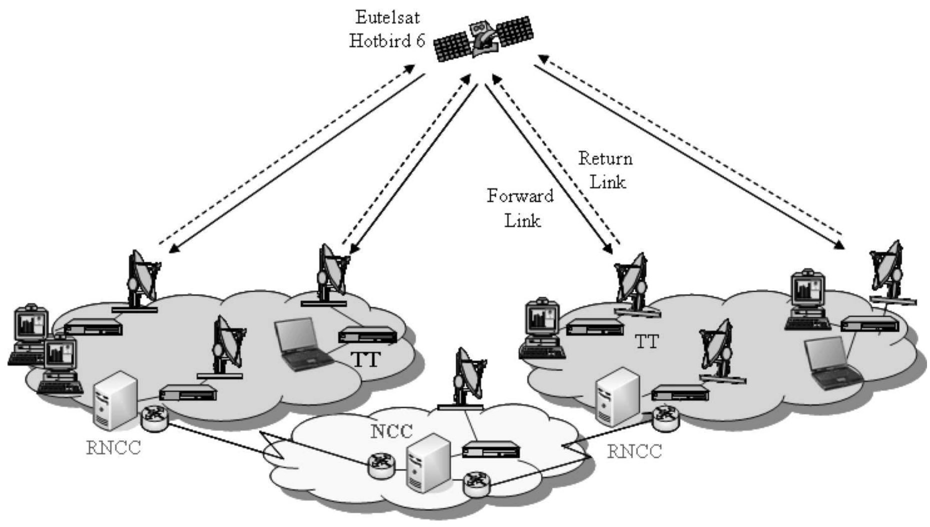

Fig. 1. Network Architecture

Figure 1 shows the main elements of a DVB-RCS multi-spot mesh network based on the Skyplex system. The network is divided into regions, corresponding to the geographic areas covered by a satellite spot beam. The group of traffic terminals (TT) that belong to a specific region are monitored and controlled by a centralised authority, called the Regional Network Control Centre (RNCC), which is connected through a terrestrial line with the Network Control Centre, a higher-level network management system. A Time Division Multiple Access (TDMA) scheme [17] is used to share the satellite uplink capacity, possibly on multiple frequencies (MF-TDMA). In a TDMA system, a *time slot* is the time granularity defined for access to the shared medium. Each *terminal* gets the right to transmit on one or more slots (an assignment) from a *master terminal* or Network Control Center (NCC). The NCC terminal periodically broadcasts a BTP (*burst time plan*) detailing the assignments for a *frame*, that is, a periodically repeating group of slots for which a BTP is valid. This design resembles, but is subtly different to most other deployed DVB-RCS networks.

In Figure 1, the RNCC assumes the role of master terminal for the terminals served by a spot. Several methods exist for deciding how to allocate assignments to the terminals. The simplest uses fixed assignments for each terminal, decided by a human operator. More complex, dynamic assignments are described under the umbrella of DAMA (demand assignment multiple access) methods. DVB-RCS defines the following terms:

- *Constant Rate Assignment (CRA)*: fixed allocation scheme.

- *Rate Based Dynamic Capacity (RBDC)*: capacity requests from the terminals to the NCC equal to the incoming traffic rate seen by the terminal from the terrestrial network. Every explicit request overrides the previous one and new requests are submitted only when needed.

- *Volume Based Dynamic Capacity (VBDC)*: capacity requests from the terminals to the NCC equal the length of the local uplink queue from the terrestrial network. This method is more stable than RBDC, because it inherently compensates for any assignment mismatch, but is slower to react to traffic changes.

- *Free Capacity Assignment (FCA)*: allows the unassigned capacity to be redistributed among the terminals without any explicit request.

Any combination of the above is allowed. For example, when combining all four methods, the request is computed as:

$$R_{req(t)} = R(t) + \frac{q(t)}{Ts} + C + F \,, \tag{1}$$

where $R(t)$, which accounts for RBDC, is the mean traffic rate entering the terminal from the terrestrial network over a time interval $[t - Tr, t]$ and where Tr is equal to the super-frame duration; $q(t)$, which accounts for VBDC, is the uplink queue length at time t, Ts is a time, C is a constant term accounting for CRA.F is the FCA volume allocated to the terminal. The logic behind Ts is that once the requested bandwidth is allocated, the terminal is able to empty its uplink queue in a time equal to Ts, while serving the traffic arriving at the terminal..

2.2 The Skyplex Data platform

The Skyplex Data system [5] was the first commercial system to implement a satellite mesh IP network for point-to-point or multicast transmission. It operates in the Ka band, and implements the DVB standard on the direct channel (downlink) and some features of DVB-RCS on the return channel (uplink). An On-Board Processor (OBP), performs stream multiplexing, allowing a fully meshed network without the need for a double-hop satellite network.

The HotBird6 satellite carries four Ka band transponders, each equipped with a Skyplex unit. Each manages several channels configurable in either a Low Rate (LR, 2.112 Mbps) or High Rate (HR, 6.226 Mbps) mode. SCPC (Single Channel Per Carrier) is possible, while TDMA mode is used for DAMA access. The total number of channels is between 6 (all HR channels) and 18 (all LR channels).

The time domain on the return channel is organised in a hierarchical structure, which consists of super-frames (the largest containers), multi-frames and frames. The BTP is sent once each super-frame, and is valid for each multi-frame contained in a super-frame. Each frame is composed of 6 time slots that are filled by bursts of 8 MPEG packets, for a total of 48 packets per frame. Since a slot can be assigned to a single terminal, the TDMA mode allows a maximum number of 6 users, with a bandwidth assignment granularity of 352 kb/s for both LR and HR. The number of users can be extended by grouping together 8 LR frames or 24 HR frames in a multi-frame. In this case, the number of users per channel becomes 6×8 in LR and 6×24 in HR, with a bandwidth assignment granularity of 44 kb/s for both LR and HR. Finally, 3 multi-frames are grouped into an 820 ms long super-frame.

Bandwidth assignment can be either static or dynamic. In the static case a fixed number of time slots are assigned to each terminal (CRA, [17]). In the dynamic case (RBDC, [17]), each terminal sends a request to the NCC, which builds the BTP, thus assigning a time slot to a terminal. The assignment does not guarantee any priority between terminals nor any service quality. The request and assignment is computed as follows:

1. Each terminal computes a bandwidth request (*BWreq*) based on the traffic arriving at the terminal from the terrestrial network, computed in a *Tr* period equal to the super-frame duration, and sends it to the RNCC. As far as we were able to measure with our equipment [3], Skyplex Data terminals use RBDC mode, so the request is computed on the base of the input data rate only.

2. The RNCC computes the bandwidth to be assigned based on the information received with the *BWreq*, by assigning bandwidth proportional to each request, with a minimum equal to the allocation granularity. If the sum of the requests exceeds the availability, it is normalized to the maximum capacity.

When a request arrives that is less than the previous one, the allocation is reduced exponentially. If $N_B(t_n)$ is the number of bursts allocated with the n-th BTP at the instant t_n and the local traffic detected in the *Tr* interval (t_n, t_n+1) correspond to a number $N_L(t_n)$ of traffic bursts, then at the (n+1)-th BTP the bursts allocated are:

$$N_B(t_{n+1}) = aN_B(t_n) + b\left[N_L(t_n) - N_B(t_n)\right] .$$

This behaviour is described in [17], but no detail is given on the values of a and b. We assumed $a = 0.5$, $b = 0.5$, thus obtaining an exponential decay with a time constant of 2 super-frames.

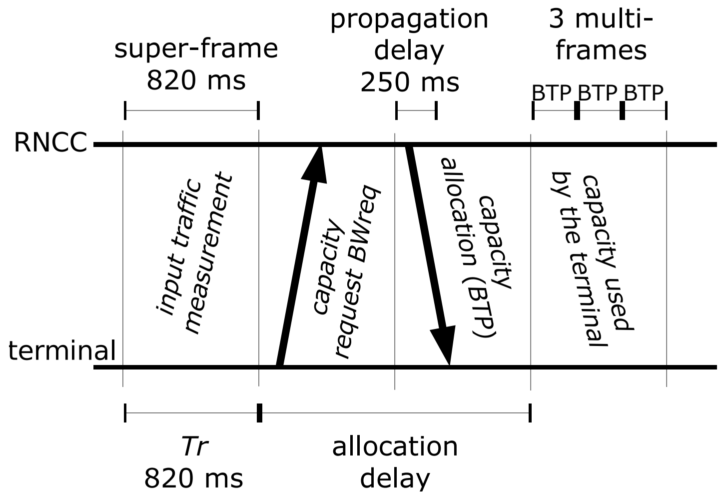

Fig. 2. Messages for bandwidth request and assignment

Figure 2 shows the sequence of messages exchanged between a terminal and the RNCC for the bandwidth assignment.

a. The terminal evaluates the rate of the incoming traffic for a time span of *Tr*, equal to the super-frame period.
b. In the next multi-frame after receiving the request the terminal sends a capacity request *BWreq*.
c. The RNCC receives the request messages from all terminals, computes the value of the bandwidth to be allocated and multicasts the BTP.
d. The terminal uses the received BTP for the next 3 multi-frames.

The *allocation delay* is the delay between the moment a request is ready, at the end of the *Tr* period, and the availability of the desired bandwidth in the BTP, which for Skyplex Data is equal to two super-frames. During the allocation delay, the terminal uses its uplink queue to absorb exceeding incoming traffic waiting for delivery.

2.3 TCP Friendly Rate Control algorithm

TFRC operates above the IP layer and was designed to provide multimedia applications with congestion control capabilities when using UDP across the Internet. TFRC is a rate-paced equation-based congestion control mechanism that requires both the sender and receiver to participate in determining the allowed sending rate [2]. In TFRC, the receiver periodically sends a feedback report informing the recent loss event rate that has been witnessed bya multimedia flow. The sender uses an equation that models the equivalent throughput that would have been obtained by a TCP flow to calculate the allowed sending rate for the next path RTT period. This includes sum of all link RTTs along the path between the sender and the receiver. The equation calculates the throughput that a TCP connection would receive under steady-state conditions given the loss event rate and the RTT of the connection:

$$M = \frac{s}{RTT\sqrt{\frac{2bp}{3}} + 3T_{RTO}\sqrt{\frac{3bp}{8}}\,p(1+32p^2)} ,$$

where M is the sending rate in byte/s, s is the packet size in bytes, RTT is the round-trip time, p is the steady state loss event rate, T_{RTO} is the TCP retransmission timeout, and b is the number of packets acknowledged by a single TCP acknowledgement, which in TFRC is set to 1. TFRC can be used by the application itself as a standalone protocol, or it can be one of several allowed mechanisms used by DCCP [15], a new IETF transport-level protocol. The simulations in this paper consider usage of TFRC as a standalone protocol.

TFRC exhibits a much lower variation of throughput over time compared to TCP, making it suitable for multimedia applications such as (TVoIP) and Voice over IP (VoIP) since it allows the sending rate to vary more smoothly by decreasing and increasing the sending rate gradually, while ensuring that it competes fairly with TCP. However, this makes TFRC respond slower to changes in available bandwidth compared to TCP. TFRC can be used for both unicast and multicast traffic.

TFRC was designed to be reasonably fair when competing for bandwidth with TCP connections using the same packet size, and does not perform well when a low-bandwidth TFRC flow using small packets shares a bottleneck with high bandwidth TCP flows using large packets, because the TFRC flow is unfairly slowed down. To solve this problem, TFRC-SP [6], a Small-Packet (SP) variant of TFRC was designed for applications such as VoIP that send small packets. TFRC-SP seeks to achieve the same bandwidth as a TCP flow using packets of up to 1500 bytes. TFRC-SP also enforces a minimum interval of 10 ms between packets, to prevent a single flow from sending small packets arbitrarily frequently.

Streaming applications use *transmit buffers*, possibly provided by the underlying network stack, to queue the encoded packets while the transport protocol dequeues the packets from the transmit buffer based on its sending rate. Streaming audio (e.g. Radio over IP) requires large transmit buffers, whereas for VoIP applications (e.g. telephony) the size of transmit buffer is usually very small. A mismatch of the encoding rate and the sending rate could lead to packets being queued in the transmit buffer and eventual buffer overflow. Similarly, the receiver uses a *dejitter buffer* to remove the effects of jitter before play out.

When traversing a satellite link, the path RTT is an important factor in determining the performance of a TFRC flow, because the sending rate is inversely proportional to RTT. Satellite networks have a high RTT: a minimum of about 500 ms for geostationary satellites, plus processing and framing delays. Some problems of using TFRC or TFRC-SP on satellite networks are:

- *Initial slow-start*: the time duration of slow-start phase is proportional to the RTT, and can be very long for applications with large encoding rates. While TFRC sends 4 packets at the beginning of the slow-start phase, this may not be enough.

- *Slow-start after an idle period*: during an idle period the sending rate may be reduced to a minimum of 4 packets per RTT, after which slow-start is again needed to return to full rate, taking longer than the initial slow-start.

- *Sending rate limit*: TFRC's sending rate can be at most twice the current receiver rate. This growth rate may not be sufficient to keep up with the encoding rate when an application oscillates between silence and talk periods.

The first two problems can be overcome with the help of Quick-Start.

2.4 The Quick-Start protocol

Quick-Start (QS) is an IETF experimental protocol designed to provide lightweight signalling of the level of congestion (specifically available capacity) between routers and a pair of communicating end hosts [7, 10]. QS was originally conceived to improve the performance of TCP bulk transfers over lightly-loaded network paths, but it may be useful in the case of multimedia flows, where it can mitigate the effect of slow starting to the encoding rate, both initially and after an idle period [11].

The sending end host sends a QS-Request for its desired sending rate (measured in bytes per second) using a Quick-Start option placed in the IP header. Each router along the path, either approves the requested rate, reduces the requested rate, or indicates that the QS-Request is not approved. Quick-Start is not a reservation protocol: routers do not reserve bandwidth, nor do they specially treat subsequent packets from the same flow; when a router approves a request, it is only asserting that at the moment there is sufficient free capacity for the requester to send at the requested rate without risking congestion. If the QS-Request is approved by all routers along the path, then the sender receives a transport-level QS-Response, after which it can send at up to the approved rate for an entire path RTT. In fact, QS only temporarily affects the flow, speeding it up, then releases control to the default congestion mechanism. If the request is not accepted, QS does nothing, and the start-up algorithm of the default congestion mechanism is used.

QS is not an end to end mechanism. To approve a QS-Request therefore requires that all routers along the path are QS enabled. While this is not necessarily an issue with equipment directly connected to a satellite terminal or connected via a (private) network that implements QS, there is less incentive to implement QS within the general Internet.

In this paper we explore the use of QS with TFRC to mitigate the effects of slow-start over a DAMA-enabled satellite link. The receiver generally sends a feedback report as soon as the first packet is received during the start of a connection [2]. Hence the QS sender ignores the first feedback report from the receiver, thus streaming at the accepted QS rate until the sender exits the QS validation phase [11]. The satellite terminal takes the role of a QS-enabled router and uses QS as a mechanism to probe the satellite sub-network about available capacity and possibly to allocate resources. This may involve some cross-layer communication, as discussed later in section 3.4.

3 Performance of TFRC over DAMA

To highlight the main characteristics of a TFRC flow we consider a simple case, where a single TFRC connection is analysed over a geostationary link managed either with constant allocation or DAMA.

3.1 Simulation scenario

Simulations were performed using the ns-2.28 TDMA-DAMA module [8]. We consider a geostationary satellite network with a HotBird6 satellite seen from Pisa (IT), with a propagation delay Tp = 250 ms. We define two configurations, both compliant with the Skyplex OBP, which we will refer to respectively as the *small-frame configuration* (24 slots per frame) and *large-frame configuration* (144 slots per frame); only the latter is compliant with Skyplex data modems. Each slot is 1504 bytes and contains 8 MPEG cells. Once per super-frame each traffic terminal generates a bandwidth request and the RNCC broadcasts the BTP containing the bandwidth allocation for all satellite terminals. The capacity of the shared satellite channel is 6.336 Mb/s, so the allocation period, i.e., the multi-frame, is 45.6 ms and 273 ms in the small- and large-frame configurations respectively, while the BTP period, i.e., the super-frame, is 3 times longer. The allocation delay is computed differently in the two configurations, because in one case the propagation delay is greater than the BTP period, while it is smaller in the other case.

TABLE I. SYMBOL DEFINITIONS

parameter	**symbol**	**small-frame**	**large-frame**
allocation period	Tf	45.6 ms	273 ms
BTP period	$Tb = 3Tf$	137 ms	820 ms
measurement intv.	$Tr = Tb$	137 ms	820 ms
propagation delay	Tp	252 ms	252 ms
allocation delay	$Ta = 6Tb, 2Tb$	820 ms	1641 ms
mean RTT	$RTTm = Tf+2Tp$	550 ms	777 ms
max QSD delay	$Tqmax = (Tr/2+Ta)-RTTm$	339 ms	1273 ms

The satellite network consists of 5 traffic terminals (including the master terminal) with directly connected end hosts. The allocation scheme is the one used in the Skyplex network, that is, it uses a rate-based allocation scheme (RBDA), each terminal is assigned at least one slot per multi-frame, the allocated rate reduces exponentially following a reduction in the requested rate, and no FCA is used.

The slow-start behaviour of TFRC leads to significant increase in delay during the start-up phase. To illustrate this phenomenon, we consider the *one way delay* (OWD) and the *end-to-end delay* (E2ED) of packets of a single TFRC flow initiating a connection over an idle satellite link (Figure 3). The OWD [9] is the delay between the wire-time transmission of the packet and its wire-time reception, which accounts for the queuing delays in router buffers and the propagation delays. The E2ED is the difference between the time the sending application generates the packet and the time the packet is delivered to receiving application's dejitter buffer. The difference between the two delays is the time spent in the transmit buffer of the sending host.

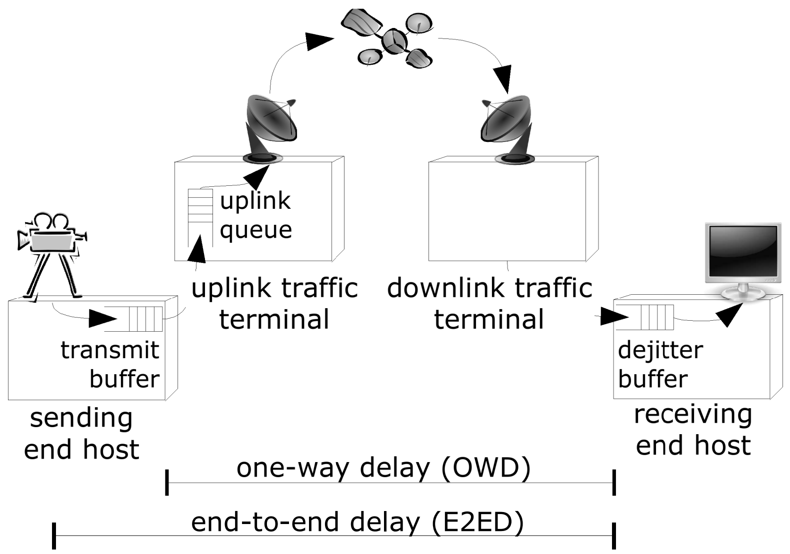

Fig. 3. One-way delay (OWD) and end-to-end delay (E2ED)

The size of the uplink queue is set to 1000 packets, as in the Skyplex platform, while the transmit buffer is unlimited in the simulations described in this section, because we are interested in the dynamics of interaction of a single connection with DAMA. The uplink buffer is set to realistic sizes in the next section, where a complex scenario with multiple connections is considered.

3.2 The TFRC start-up problem

In this section we consider a deterministic simulation scenario, where a single terminal transmits a single flow at 256 kb/s with a packet length of 1316 B. To have a complete picture of the dynamics of delay during the start-up period, we repeated the simulation starting the flow with all possible offsets with respect to the super-frame, in 1.9 ms steps. Where vertical bars are depicted, the plots show the minimum, maximum and median of the OWD and E2ED versus the time elapsed distance from the beginning of connection; other plots show only the median value.

Figure 4 shows the simulated OWD and E2ED in both the small- and large-frame configurations with TFRC and DAMA.

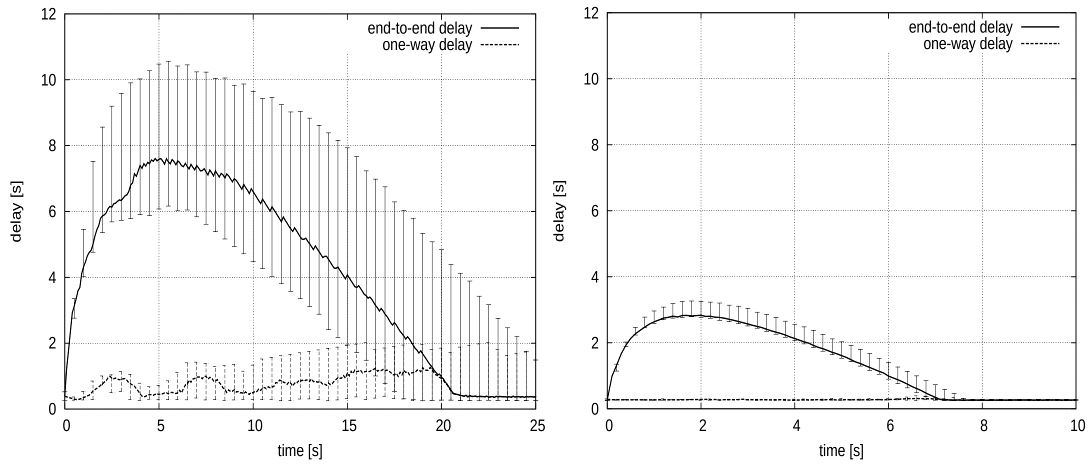

Fig. 4. Performance of a single 256 kb/s flow with DAMA.
(a) large-frame configuration, (b) small-frame configuration.

In Figure. 4 the minimum delay is equal to the propagation delay *Tp*. The difference between the top curve (E2ED) and the bottom curve (OWD) is the time packets are queued at the transmit buffer used by TFRC. The difference between OWD and the minimum delay, which is visible in the large-frame configuration, is due to queuing at the uplink terminal waiting for DAMA to assign the requested allocation. The OWD oscillates during start-up because TFRC increases its throughput exponentially using a slow-start algorithm. Once the sending rate exceeds the allocated rate, a new DAMA request is issued, and packets are queued at the uplink queue until the request is honoured one allocation delay later (1.6 s for the long-frame configuration and 0.8 s for the short-term one). During this time, the queue keeps growing and the rate keeps growing slowly following the slow-start algorithm. The effect is that slow-start is much slower on DAMA than on a channel with preallocated bandwidth having the same round-trip time [3]. Since during slow-start the sending application keeps encoding at a fixed rate, there is a mismatch between the encoding rate and the transmit rate governed by TFRC, which translates into queuing at the transmit buffer and the observed E2ED. This is the main effect that significantly slows TFRC on a DAMA link.

E2ED exhibits a large variability; for example, in the large-frame configuration, after 15 s from the beginning of transmission, packet delay can have any value in the range between 5 and 9 s, approximately. Specifically, small delay variations at the beginning of the connection cause large variations in E2ED. This is due to the way the receiver estimates the rate: since the flow rate is estimated as the number of packets seen by the TFRC receiver in a RTT, the estimator suffers from a quantization effect when the number of packets transmitted is small. Both the E2ED and its variability are significantly reduced in the small-frame configuration. This is an important observation for designers of DAMA systems.

This, and subsequent sections, assume a constant media rate for the data sent using TFRC. Actual application dynamics depend upon the codec design, application behavior and encoded media content. Future applications may employ variable rate codecs to tune performance to available capacity. Such codecs result in traffic bursts. We note that although TFRC is a rate paced protocol, the packets that arrive at the terminal uplink are not necessarily evenly paced due to implementation constraints and buffering in the network (these effects are not considered in this work). This depends on the way the time slots are allocated and could result in Quick-Start packets to be sent as bursts which could aggravate the level of congestion.

3.3 QS improves TFRC start-up performance

Section 3.2 identified the start-up delay problem as the interaction of slow-start in TFRC with DAMA, this section seeks to mitigate this problem. We analyse how much DAMA makes things worse with respect to a constant rate allocation (CRA) on the satellite link, and we analyse how well QS helps at start-up. We assume that the uplink traffic terminal implements a QS agent that receives QS requests, evaluates how much free capacity is available on the shared satellite channel and forwards the QS-Request accordingly. More details on the evaluation process are given in Appendix A.

Figure 5 summarises these comparisons. The simulation scenarios are the same as described above, with and without DAMA, with and without QS. Only the median of E2ED is shown, for clarity.

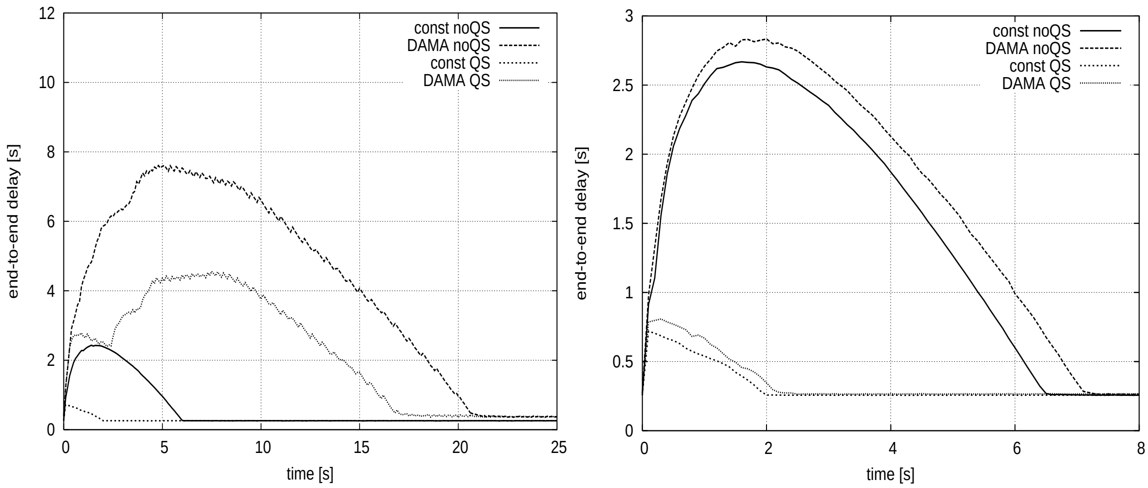

Fig. 5. Single 256 kb/s flow with and without DAMA, with and without QS.
(a) large-frame configuration, (b) small-frame configuration.

In the CRA case the preassigned bandwidth is 1.2 Mb/s, more than 4 times than that needed by the flow. As expected, performance with CRA is better than with DAMA, because the DAMA effect of slowing slow-start disappears [3]. As expected, QS brings a clear improvement, by eliminating the slow-start phase after an RTT. This effect is observed in the case of no-DAMA, QS performance: after an initial increasing delay due to traffic accumulating in the transmit buffer during the first RTT (as a result of slow-start), the QS-Request is accepted and TFRC starts to transmit at faster than the full rate, thus emptying the transmit buffer. The same happens with DAMA QS, but this time the traffic from the transmit buffer accumulates in the uplink queue, until the DAMA request is granted to the requesting terminal.

The bandwidth allocation flexibility brought by DAMA comes at a cost. There is a much slower start-up phase for protocols using slow-start ,such as TCP [3] and TFRC, especially when there is a large DAMA allocation delay is big. This cost is significantly reduced by the use of QS, but remains non negligible for large allocation delays. This latter case is analysed in greater detail below.

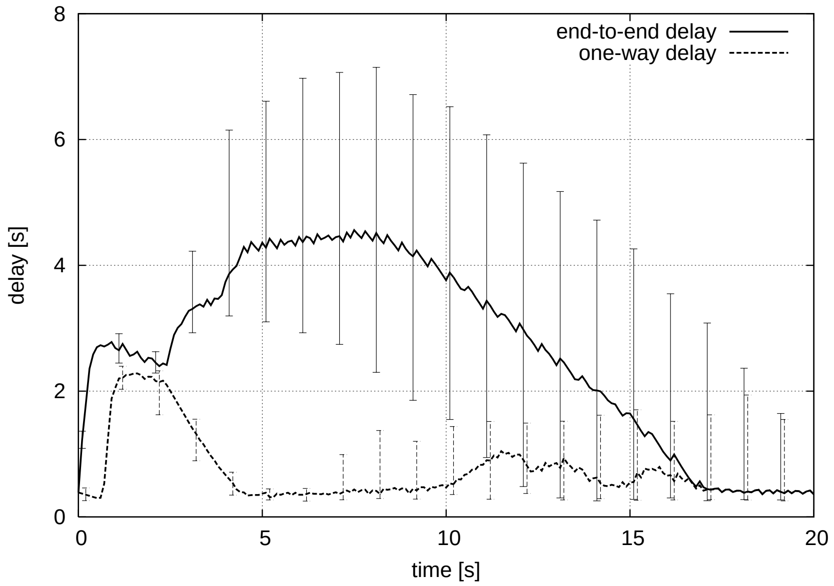

Fig. 6. Single 256 kb/s flow with DAMA and QS, large-frame configuration.

Figure 6 separates the E2ED into the time spent in the TFRC transmit buffer — represented by the difference between the E2ED curve and the OWD curve — and the OWD itself, due to the fixed propagation delay plus the time spent in the DAMA uplink queue. As in all figures in this section, the line is the median value over all possible time offsets of the start of the flow with respect to the super-frame, and the vertical bars show the variability between minimum and maximum values. In the following analysis, all times are mean times: exact times are dependent on the phase of the satellite frame at beginning of transmission.

The application starts sending at a full rate of 256 kb/s. A QS-Request of 320 kb/s is generated for this rate and TFRC starts its slow-start phase, resulting in packets queued at the transmit buffer, while the uplink queue stays empty because of the minimum DAMA allocation. After one *RTTm* (0.8s), TFRC receives the QS-Response and starts sending at the full rate. Since DAMA allocation is still at its minimum of 44 kb/s, packets are queued in the uplink queue. In the meantime, a DAMA request is issued for the current rate (generally less than the full rate) every *Tr* (0.8 s), the exact timing of the first request depending on the phasing of the request with respect to the start of an allocation frame. Packets are still queued at the uplink queue and the delay increases linearly for one allocation delay after which, at time 2.5 s, the allocation is received and the queues begin to empty. During this time, TFRC receives the first feedback report after the QS-Response, and ignores it [11]. The second feedback report arrives at 3*RTTm*, that is 2.3 s, but reports a rate that is much less than the full rate, so TFRC throttles the sending rate: the uplink queue empties, DAMA makes a request for a lower rate and the transmit buffer starts filling. From this point on, the interaction between the DAMA closed loop and the TFRC closed loop generates an oscillating behaviour whose exact dynamics depend on their relative phase (*phase effect*), leading to wide variability, as shown by the vertical bars. In-depth discussion can be found in Appendix B.

3.4 Improving QS performance in a satellite networks: QSD

A major problem of the sequence of events visible in Figure 6 is that the requested DAMA allocation arrives too late with respect to the QS Response. This causes TFRC at the receiver to issue a feedback report with a low perceived rate and in turn causes the TFRC sender to throttle its sending rate just when the requested bandwidth is going to be available on the satellite channel, causing bandwidth oscillation and an increased overall delay.

To improve the performance of TFRC over DAMA using QS, we describe a method that accelerates the availability of satellite bandwidth for the transmitter, and additionally evaluate the possibility of delaying the QS-Request so that the QS-Response and the DAMA bandwidth availability are synchronised at the transmitter. This is obtained through a cross-layer signalling method that we name QSD, short for QS-DAMA. Once the QS agent at the uplink traffic terminal receives a QS-Request and decides there is enough free capacity on the shared satellite channel (see Appendix A for details), it adds the QS-Request to the next DAMA bandwidth allocation request. More precisely, for each *Tr* interval, the request is computed as defined in (1) plus the sum of the QS Requests received during the *Tr* interval that can be satisfied by the currently unused shared satellite bandwidth.

To align the arrival of QS-Response at the requester with the availability of the DAMA request on the BTP, the QS agent may delay the QS-Request by a short time that we call the *QSD delay Tq*, before forwarding it along the data path. This choice has the disadvantage that the first data packet, which transports the QS request, is also delayed, leading to an initial delay increase equal to *Tq*.

All in all, QSD consists of two mechanisms:

- *Immediate DAMA request*: the QS agent at the terminal adds the requested QS rate to the next DAMA request .
- *Optional QS request delay*: the QS agent at the terminal waits for a QSD delay *Tq* before forwarding the QS-Request to the next hop in the network.

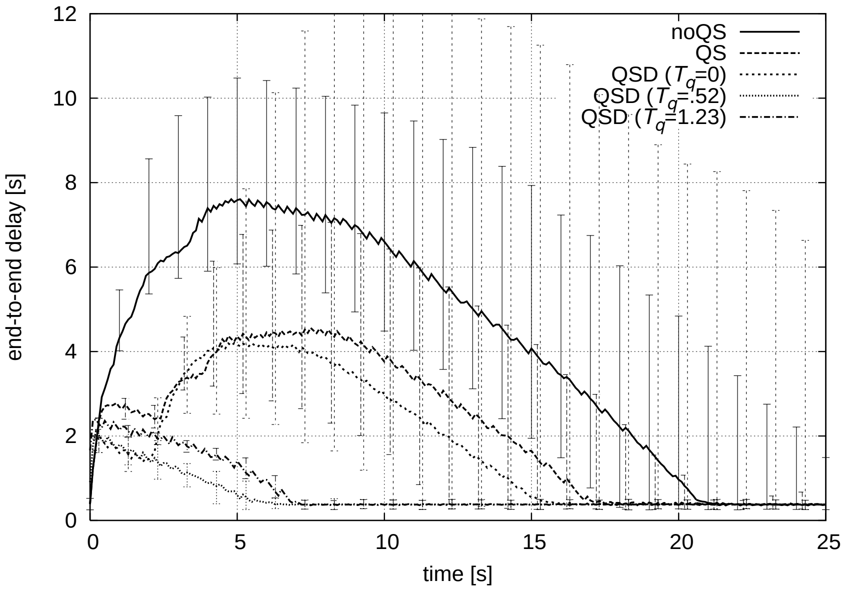

Fig. 7. Single 256 kb/s flow with DAMA and various QS strategies, large-frame configuration.

Figure 7 compares E2ED using DAMA with different QS strategies. The worst performer (largest E2ED) is the case without QS, the same depicted in Fig. 5a. The next worst is QS without any interaction with DAMA, depicted in Figure 6. The behaviour of the proposed QSD implementation without an added delay is represented by the line marked with $Tq = 0$, indicating that the QS request is forwarded immediately: the delay profile of QSD is similar to that of QS. The time elapsed from the issue of a QS Request by the transmitter to the moment it obtains an allocation on the satellite channel is equal on average to $Tr/2 + Ta$, which is 2051 ms. The time that elapses between the issue of the QS-Request and the reception of a QS-Response, assuming that delays due to the terrestrial networks are negligible, is equal to the mean RTT, which is 777 ms. This means that the approved QS- Response arrives at the sender at a time when no DAMA capacity has yet been allocated. This is the same problem that was observed in the QS case.

When this timing mismatch is compensated for by delaying the QS-Request by

$$Tqmax = Tr/2+Ta - RTTm, \qquad (2)$$

the performance significantly improves (indicated by the lowest curve)..

The reduction in maximum E2ED is shown in Figure 8, for each of the four cases of Figure 7 but at different flow rates.

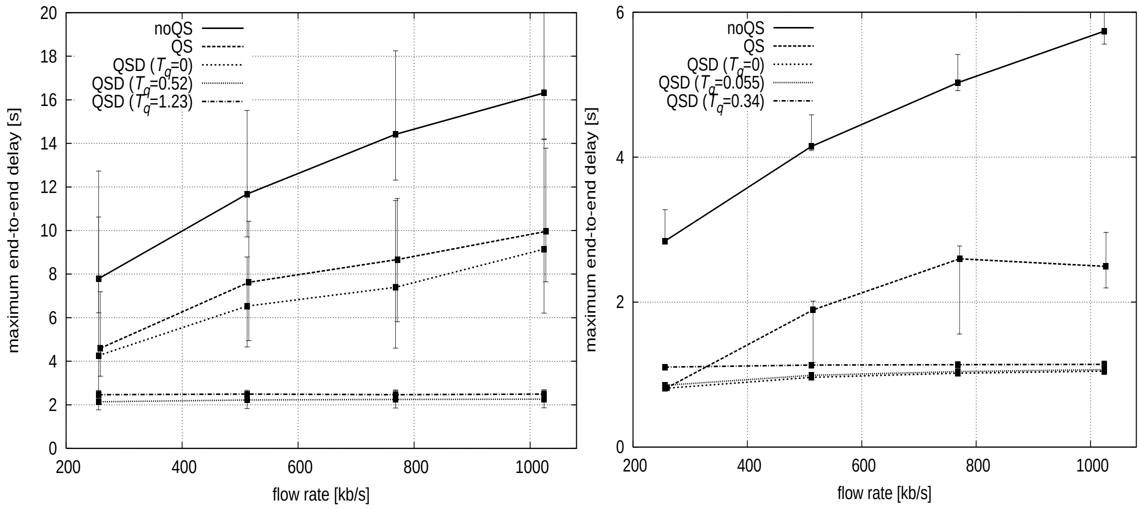

Fig. 8. Single flow with DAMA and various QS strategies.
(a) large-frame configuration, (b) small-frame configuration.

In these results the maximum delay increases with flow rate for all strategies except for QSD, for which it remains low and constant in both the large- and small-frame configurations. Simulations however show that values of *Tq* significantly smaller than *Tqmax* computed in (2) give very similar results, as depicted in Figure 8. It is therefore sufficient to set $Tq = 520$ms and $Tq = 55$ in the large- and small-frame configurations respectively. Appendix B discusses how these values are calculated.

Delaying an initial packet carrying a QS-Request could lead to a nofeedback timer expiry if the feedback does not arrive within the initialized RTO. TFRC as a standalone protocol specifies an initial RTO of two seconds [6]. Since a QS-Response must be received within this period, this places an upper bound on the total round trip delay that can be presented by the end-to-end path. While the proposed *Tq* delays in this paper are short enough to prevent any nofeedback timer expiry, the Datagram Congestion Control Protocol (DCCP), RFC 4340 [14], also supports a TFRC-based congestion control method whose specification states that the initial RTO is only one second. We believe that this is too low an initial RTO value for some Internet scenarios (including the one described in this paper), and justify this claim by citing the specified initial RTO value for TCP of three seconds [20].

4 Aggregated flows scenario

We have analysed the dynamics of TFRC over DAMA and have identified ways to improve performance by studying the simple specific case of a single connection on an otherwise idle satellite network. In this section we consider a satellite network and the parameters that are the same as described in Section 3.1, but with a satellite network that is lightly loaded, rather than unloaded.

We consider three set of five traffic flows streamed between randomly selected hosts connected to a satellite terminal. The traffic generated by the sources approximate VoIP calls, video phone calls and video clips. Each flow has a Poisson integration time, with the same parameters for each of the five traffic sources. The duration of the flows is exponentially distributed for simplicity (this duration only determines the overall background load, whereas the interesting dynamics involving TFRC occur in the first few seconds of flow life).

The parameters of the three types of source are shown in Table II. *Tolerable delay* is the maximum delay jitter that the transmitter considers acceptable, after which the packet is considered late and discarded. *Packet length* for 16 kb/s is chosen to be consistent with Skype; for 256 kb/s flows the length represents that used by the videoconferencing program VLC (equal to 7 MPEG packets). The *transmit buffer* is computed from the two previous parameters. *Mean duration* values are chosen for a typical flow duration for phone calls (VoIP), videophone calls and video clips (IPTV). The *mean interval* is the flow intergeneration mean time, chosen so that the *traffic share* of each type of traffic is equal to the one in the last column. The total load on the link is about 2 Mb/s, that is one third of the shared capacity.

TABLE II. STREAMING SOURCES

rate	**tolerable delay**	**packet length**	**transmit buffer**	**mean duration**	**mean interval**	**traffic share**
16 kb/s	0.1 s	160 B	2 pkt	60 s	1 s	50%
64 kb/s	0.1 s	512 B	2 pkt	180 s	40 s	15%
256 kb/s	5 s	1316 B	128 pkt	300 s	120 s	35%

In the simulation we used TFRC-SP for the 16 kb/s flows and TFRC for the 64 and 256 kb/s flows. This distinction is not significant in these scenarios, because the satellite network is underutilised, (far from congestion).

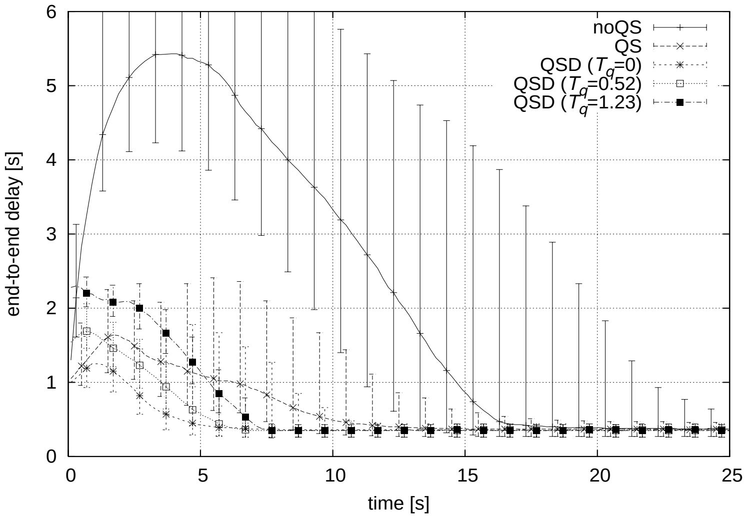

Fig. 9. Start-up behaviour for 256 kb/s flow in the aggregated flow scenario.

Figure 9 shows the E2ED of the IPTV connections (256 kb/s) in the aggregated flow scenario. Using QS brings considerable advantages in the E2ED performance, even when the network is moderately congested (one third of the total capacity). The error bars show the 5th and 95th percentiles over a total of 100000 simulated VoIP flows, 2500 video conferencing flows and 833 IPTV flows.

The improvement of QSD with respect to QS is less dramatic in the aggregated scenario compared to the single-flow scenario. This is because when a new IPTV flow starts sending packets at a high rate, most of the time the satellite terminal already has bandwidth allocated in response to requests from other flows traversing the same terminal. This bandwidth is shared between all the competing flows, each of which suffers a little, but much less than the single flow suffered in the completely unloaded scenario of the previous section. Additionally, the chosen DAMA algorithm does not immediately relinquish the assigned capacity when terminals reduce the requests. In a dynamic scenario, where connections continuously end, releasing bandwidth, this means that new connections often exploit some of the (unused) capacity still allocated to the terminal.

As a consequence, not delaying the QS request ($Tq = 0$) results in better performance than introducing a delay in this scenario. However, the single-connection scenario analysed in the previous section, while not general, is realistic, because it is an important situation when a single host or a small office is connected to the Internet through a satellite connection and a multimedia flow is started on an otherwise idle network. For this reason it is useful to find a value for *Tq* that is as small as possible, yet allows QSD to provide an advantage to the single-connection scenario. In Appendix B we compute this value, which we call *Tqopt*, as 520 ms and 55 ms in the large- and small-frame configs, respectively. Figure 9 shows that setting $Tq = Tqopt$ gives an E2ED performance that is midway between $Tq = 0$ and $Tq = Tqmax$.

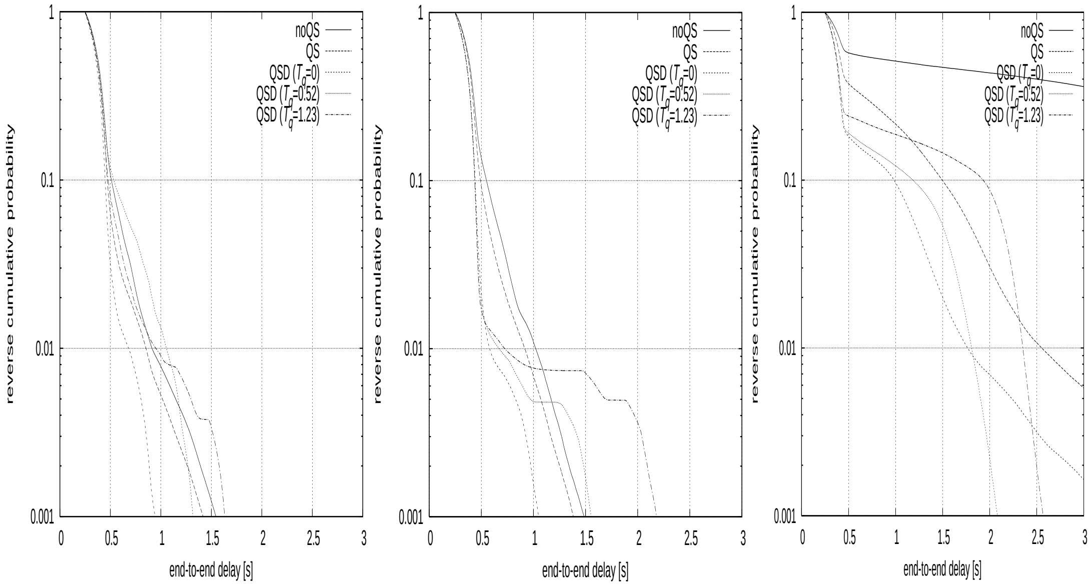

Fig. 10. Distribution of E2ED for
(a) VoIP, (b) Videoconference and (c) IPTV connections.

The variability of E2ED over time is important since this dimensions the required dejitter buffer at the receiver. Jitter is very high at the beginning of a flow, then it stabilises to a low value due to framing on the satellite link and to concurrent flows using the shared bandwidth. Figure 10 reports the distribution of the E2ED for the three classes of flows over the first 30 s after the beginning of the flow. In Figure 10 a it is shown how the different strategies have little influence on the delay behaviour of VoIP, due to the low bandwidth used, which generally allows VoIP to profit from any allocated DAMA excess bandwidth. The same happens for 99% of the videoconference call packets, as shown in Figure 10 b, where the only noticeable feature are the tails of the *Tq*=0.52s and *Tq*=1.23s cases, which end up being a copy of the *Tq*=0 case, but shifted to the right by *Tq*: the packets that create the tails are the first ones in the flows, those that are delayed by *Tq*. Figure 10 c shows the graph relative to IPTV flows, where we observe that performance of the noQS case is clearly worse than the others, and that *Tq*=0.52s and *Tq*=1.23s are the ones with the sharpest descent, being copies of one another apart from the second being shifted to the right (i.e., with highest delays) by 1.23s-0.52s. Notice how Figure 10 c shows that the choice of *Tq=Tqopt* is indeed the best among the proposed ones.

Table III shows the fraction of packets lost because of transmit buffer overflow when TFRC throttles the data rate output from the sending node during slow-start, while the data source transmits at full rate. An IPTV transmit buffer set to tolerate a 5 s delay, is able to absorb all delays when QS is used. No packets are lost because of buffer overflow in the large terminal uplink queue of 1000 packets, (the value used in Skyplex Data equipment).

TABLE III. FRACTION OF PACKETS LOST DUE TO TRANSMIT BUFFER OVERFLOW.

type	tot	VoIP	Video	IPTV
No QS	0.11	0.11	0.11	0.025
QS	0.044	0.045	0.044	0
QSD (*Tq*=0)	0.036	0.037	0.036	0
QSD (*Tq*=0.52)	0.020	0.020	0.020	0
QSD (*Tq*=1.23)	0.012	0.012	0.012	0

Conclusions and future work

This paper has discussed the impact introduced by a satellite link that uses Demand Assignment Multiple Access (DAMA) on the performance of streaming multimedia flows that utilise the congestion-control offered by the TCP Friendly Rate Control (TFRC) algorithm. TFRC can improve co-existance of TCP and UDP streaming flows and therefore avoid congestion collapse. However, the paper shows that the varying delay introduced by DAMA can significantly impact the performance. This effect is a result of interactions between the end-to-end congestion control loop and the control loop used on the satellite link to request/allocate capacity to a specific terminal. Similar effects have previously been observed for TCP sessions.

To mitigate these effects this paper proposes and analyses the use of an experimental method, known as Quick-Start. Quick-Start provides a way for the network-layer to inform the link-layer about the expectations of a new media flow, and more efficiently schedule capacity to a new flow. In our proposal, we modify the QS agent at the uplink satellite terminal to inform the DAMA system as soon as a QS Request is received.

One side-effect of implementing QSD is that the satellite uplink router commits capacity as a result of receiving a QS-Request. Hence if the sender does not utilise its approved QS rate, then this additional allocation could be wasted (or allocated to other traffic flows). In this paper we assume directly connected end-hosts, or end hosts connected via QS-enabled networks whose transit delay is significantly less that the total satellite link delay. If the sender were to be connected via an intermediate network with appreciable delay, this would delay the QS packets that arrive at the satellite terminal resulting capacity being allocated before the media flow starts. This reserved capacity could be wasted.

Results are presented to show the performance with and without using the QS method. It is shown the QS can significantly improve performance, especially when the total allocated capacity to a terminal is much less that the volume of new traffic being added by the multimedia flow (e.g. when the link is idle, or has little existing traffic).

We observe that when using QS, the performance improves when the QS agent delays forwarding the QS-Request. We call this method QS-DAMA (QSD). This may partially align the reception of a QS Approval with the reception of the allocated DAMA capacity. Through fluid-flow analysis we have derived appropriate values for this delay that improve performance when a link has little existing allocated capacity. However, delaying the QS-Request also has the disadvantage that in delaying the QS-Request packet, it also delays the first packet of the flow, which may result in the flow measuring an inflated initial Round Trip Time (RTT), which paradoxically can result in preventing the flow taking full advantage of the QS capacity (since the flow must pace its transmission according to the measured RTT). We suggest that judicious choice of when to employ the extra delay can alleviate this problem.

Although the analysis of performance employed parameters appropriate to the Skyplex platform, we expect the techniques to be applicable to other similar systems, including DVB-RCS.

Appendix A: Estimation of unused capacity on the satellite channel

When the QS agent at the uplink terminal receives a QS-Request, it must evaluate whether there is enough capacity on the shared satellite channel to be allocated to satisfy the received request. Terminals may become aware of current bandwidth utilization using a centralised approach A discussion of both approaches follow:

A.1 Centralised estimation

A.1.1 Continuous broadcast

The RNCC broadcasts the amount of free capacity. This could be the BTP, requiring an appropriate field to be added to the DVB standard, or by sending a packet/frame to all the terminals in a private format.

A.1.2 On demand

The RNCC answers queries from the terminals requesting the amount of free capacity in a private format. This method implies an additional delay equal to *RTTm* for the QS-Request to be forwarded.

A.2 Distributed estimation

Each terminal individually evaluates the amount of free capacity. There are two possibilities, depending on DAMA using FCA (free capacity assignment).

A.2.1 DAMA does not use FCA

Estimation is accomplished by examining the BTP and summing the capacity of all unassigned slots.

A.2.2 DAMA uses FCA

The only possibility left is to listen for other terminal's transmissions. Depending on the topology and capabilities of the network, it may be impossible to use this method.

A.2.2.1 Single spot without fade countermeasure

The estimation is possible but inaccurate, because a terminal may not transmit in a slot it has requested, or could be transmitting in a slot it has not requested, but was assigned by FCA.

A.2.2.2 Single spot with fade countermeasure

If up-power control or adaptive modulation and coding are used as fade countermeasure systems, a terminal in deep fading may be unable to listen for transmissions directed to other terminals. This would lower the accuracy of estimation during fades.

A.2.2.3 Multiple spots

If the satellite network is multi-spot, each terminal can only listen for transmissions directed to its spot, so it can only have an estimate of the unused capacity for transmissions directed to its own spots.

Our simulation assume that the free capacity can be determined without delay, so we consider case A.1.1. In [10] two methods are considered for evaluating the unused bandwidth: considering a moving average of a number N of equally spaced observations and considering their moving peak. In our setting, we make a new observation at each BTP, so the estimation interval is *Tb*. We considered only the last estimation, so we set $N = 1$. Different values for N could be considered in future research.

The algorithm used by the QS agent for deciding whether to approve a request is the Target algorithm defined in [10]. The agent considers the channel to be underutilised if the free capacity is at least 20% of the shared bandwidth, which in [10] means setting *qs_thresh* = 0.8. In order to avoid over allocating bandwidth to several consecutive requesters, the QS agent remembers the requests approved in the last *M time intervals*. In our simulations M is set to 2, and the variable *state_delay* contains the length of the time interval. We ensure that the agent remembers approval of QS-Request for the maximum time that can elapse from the reception of the request to the time it is reflected in the received BTP. This time is equal to *Tr* + *Ta*, so we set *state_delay* = (*Tr*+*Ta*)/2.

Appendix B: Optimal value for the QSD delay *Tq*

Equation (2) provides a simple way to compute a value for the QS delay *Tq*. However, simulations show that the value obtained is significantly higher than necessary. Lowering *Tq* has beneficial effects in the multiple-connection scenario analysed in Section 4. The DAMA algorithm we adopted deallocates bandwidth slowly, hence a terminal has often more available capacity than requested. When additional capacity is available, it is used by newborn connections: for these the smaller the *Tq* delay, the better.

It would be useful to compute the minimum *Tq* value *Tqopt* for which the advantages shown in Figure 7 hold. The optimal value is thus in the range from 0 to the value from (2), that is 1.23 s and 0.34 s in the large- and small-frame configs, respectively.

Tqopt should be computed as the minimum value of *Tq* such that the delay keeps decreasing after the initial local maximum, rather than showing the hump visible in Figure 7 for the case of *Tq*=0. We developed a fluidic model for the single-connection scenario to study the delay behaviour for the two considered configurations. This is available for interactive use and download at [19]. Figure B1 shows the default configuration computed by the model in the large-frame configuration with *Tq* set to 400 ms. The four monotonely increasing curves are the cumulative outputs in kb/s at four different points in the network model as shown in Figure 3: at the output of the traffic generator (the generator's rate), at the output of the transmit buffer (the TFRC rate), at the output of the uplink queue (the DAMA rate) and at the output of the receiving terminal (the received rate). Horizontal distances from one curve to the next are equal to the amount of buffered traffic in the transmit buffer (the TFRC buffer) the uplink queue (the DAMA buffer) and the radio link, respectively. Vertical distances are equal to the respective delays, that is the delay introduced by TFRC in the transmit buffer, the delay introduced by DAMA in the uplink buffer and the propagation plus framing delay *RTTm*, respectively (see Table I). The E2ED is the vertical distance between the first and the fourth curve, that is, the total delay between the generator and the receving terminal; in Figure B1 it is plotted versus the sending time.

The main difference of the fluidic model with respect to the simulation are that packetisation at the generator's output is not considered, framing at the uplink queue output is not considered and the receiver's estimate of the round trip time is fixed and set equal to *Frstep* = *RTTm*+*Tq*.

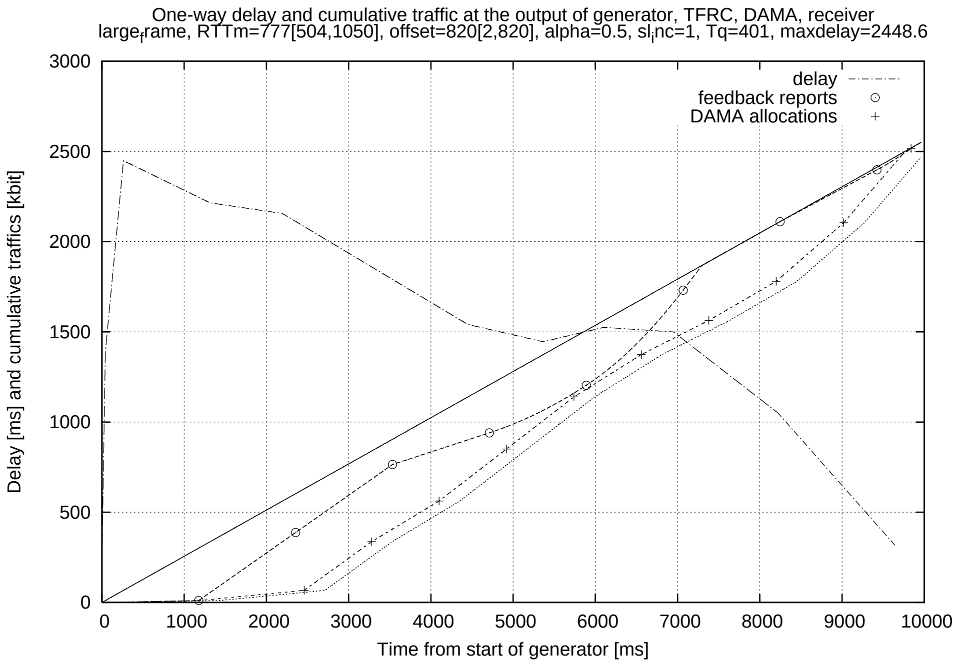

Fig. B1. Fluidic modelling of traffic queuing in the large-frame configuration.

Figure. B1 presents the results of the fluid model for a single connection scenario. The generated traffic accumulates in the transmit buffer until the first feedback report is received (the QS Approval) at time *RTTm*+*Tq* = 1178 ms. From this point on, the transmit buffer empties linearly, because its output rate is 320 kb/s while the rate of the generator is 256 kb/s. From the reception of the first feedback report to the reception of the first DAMA allocation at time *Tr*+*Ta* = 2460 ms the uplink queue increases, after which it empties at the same rate as the transmit buffer, i.e., 320-256 kb/s. The second feedback report is ignored [11], and the sending rate is maintained until the third feedback report is received at 3*Frstep* = 3534 ms. The 3rd feedback report has been sent by the receiver at time *qen* = 3*Frstep*-*RTTm*/2 and is based on traffic received in the interval from *qen*-*Frstep* to *qen*, that is, traffic that exited the uplink queue during the time interval from *qen*-*Frstep*-*RTTm*/2 = *RTTm*+2*Tq* = 1579ms to 2*RTTm*+3*Tq* = 2757ms. During this time interval the first DAMA allocation is received in our scenarios. This means that the traffic received during this time interval is initially very low (the minimum DAMA allocation) and then is equal to the QS request (because of the QSD request).

Since the 3rd feedback report is significantly lower than the QS-Request (unless we set a very high *Tq*), TFRC slows down to the reported rate until it receives the 4th feedback report. At this point, it grows gradually its rate up to the maximum of twice the 4th feedback report and twice the previous rate (that is, the 3rd feedback report). In Fig. B1 this rate is reached when the 5th feedback report is received; at that moment, the uplink queue empties and the E2ED stops growing, because the TFRC rate is the same as the rate of the generator. Experimenting with various values for the involved parameters shows that the behaviour shown in Fig. B1 is typical of the situation when *Tq* is near its optimal value, that is when increasing it further does not significantly lower the maximum E2ED.

To compute *Tq* while approximating on the safe side, we consider the initial received rate as 0 – rather than the minimum DAMA allocation rate – and the final received rate as the rate of the generator of 256 kb/s – rather than the QS-Request. From the above discussion, we want the TFRC rate at the 5th feedback report to be equal to the generator's rate. This means that the rate of the 3rd feedback report should be half the generator's rate, which in turn means that the first DAMA allocation should be received in the middle of the interval above defined, leading to the equation

$$Tr+Ta=RTTm+2Tqopt+\frac{RTTm+Tq}{2}\Rightarrow Tqopt=\frac{2(Tr+Ta)-3\mathrm{RTTm}}{5} \quad . \tag{B1}$$

Equation (B1) yelds *Tqopt* = 520 ms and *Tqopt* = 55 ms for the large- and small-frame configurations, respectively. As said, these values are computed with many approximations, the most important being the fact that the QS-Request (320 kb/s) in our case is significantly larger than the rate of the generator(256 kb/s). Since the approximations are all done on the safe side, the values obtained for *Tqopt* are larger than those that minimise the maximum E2ED in the worst case as seen in the simulations.

The purpose of computing *Tqopt* as the value of *Tq* is to find the minimum *Tq* value for which the packet delay keeps decreasing after the initial maximum. The packet that suffers the initial maximum is sent at a time *th_maxt* approximately equal to

$$th\_maxt = ((Tr+Ta)-(RTTm+Tqopt))*mindama/full,$$

where *mindama* is the minimum DAMA allocation given to terminals even when they request no traffic, and *full* is the sending host streaming rate. The delay *th_maxd* experienced by the packet sent at *th_maxt* is approximately equal to

$$th\_maxd = (Tr+Ta)+RTTm/2-th\_maxt,$$

so one can say that the maximum delay experienced by a flow is for packets sent after *th_maxt* after the flow, that it is equal to *th_maxd*, and that this delay is not exceeded by other packets in the flow if *Tq* = *Tqopt*. The expressions given above for these quantities are approximately valid for a wide range of parameters.

References

[1] Z. Xiang; Q. Zhang; W. Zhu; Z. Zhang; Y.-Q. Zhang, "Peer-to-peer based multimedia distribution service", IEEE Transactions on Multimedia, Vol. 6, Issue 2, pp. 343- 355, April 2004

[2] M. Handley, S. Floyd, J. Padhye, J. Widmer, "TCP Friendly Rate Control (TFRC): Protocol Specification", draft-ietf-dccp-rfc3448bis-01, Internet draft, draft in Progress, March 2007.

[3] A. Gotta, R. Secchi, F. Potortì, "An analysis of TCP startup over an experimental DVB-RCS platform", proceedings of the International Workshop on Satellite and Space Communications (IWSSC), pp. 176-180, Madrid (ES), September 2006.

[4] A. Sathiaseelan, G. Fairhurst, "Using Quickstart to Improve the Performance of TFRC-SP over Satellite Networks". , IEEE International workshop on Satellite and Space Communications (IWSSC) 2006, Madrid.

[5] E. Feltrin, E. Weller, E. Martin and K. Zamani. Design, Implementation and Performance Analysis of an On Board Processor-Based Satellite Network. In proc. of International Conference on Communications (June 2004), vol. 6, pp. 3321–3325.

[6] S. Floyd, E. Kohler, "TCP Friendly Rate Control: the Small Packet (SP) Variant", IETF RFC 4828, April 2007.

[7] S. Floyd, M. Allman, A. Jain, P. Sarolahti, "Quick-Start for TCP", IETF RFC 4782, January 2007.

[8] TDMA-DAMA

[9] G. Almes, S. Kalidindi, M. Zekauskas, "A One-way Delay Metric for IPPM", IETF RFC 2679, September 1999.

[10] P. Sarolahti, M. Allman, S. Floyd, "Determining an Appropriate Sending Rate Over an Underutilized Network Path". To appear in Computer Networks V51, N7, May 2007.

[11] G. Fairhurst, A. Sathiaseelan, "Quick-Start for DCCP", IETF Internet draft draft-fairhurst-tsvwg-dccp-qs-00.txt , work in progress, March 2007.

[12] EN 301 192 "Digital Video Broadcasting (DVB), Specification for Data Broadcasting".

[13] G. Fairhurst, C. Nocker, “Unidirectional Lightweight Encapsulation (ULE) for Transmission of IP datagrams over an MPEG-2 Transport Stream (TS)”, IETF RFC 4326, December 2005.
[14] E.Kohler, M.Handley, S.Floyd, “Datagram Congestion Control Protocol (DCCP)”, IETF RFC 4340, March 2006.
[15] S. Floyd, E. Kohler, J. Padhye, “Profile for Datagram Congestion Control Protocol (DCCP) Congestion Control ID 3: TCP-Friendly Rate Control (TFRC)”, IETF RFC 4342, 2006.
[16] EN 301 421, "DVB Modulation and Coding for DBS satellite systems at 11/12 GHz", ETSI 1999.
[17] EN 301 790, "DVB Interaction Channel for Satellite Distribution Systems", ETSI, 2003
[18] EN 302 307, “DVB; Second Generation Framing Structure, Channel Coding and Modulation Systems", ETSI, 2006
[19] Fluid Simulator, <http://wnet.isti.cnr.it/software/damatfrc.m>
[20] V.Paxson, M. Allman, “Comupting TCP's Retransmission Timer”, IETF RFC 2988, November 2000.
[21] SATSIX project, www.ist-satsix.org