

This is a preprint of an article published in
the International Journal of Satellite Communications and Networking (IJSCN)
<http://www.interscience.wiley.com/>

Vol. 27, No. 3, pp. 141-161, May-June 2009
<http://dx.doi.org/10.1002/sat.929>

Using Quick-Start to enhance TCP-friendly rate control performance in bidirectional satellite networks

Raffaello Secchi^{1,*}, Arjuna Sathiseelan², Francesco Potorti¹, Alberto Gotta¹ and Gorry Fairhurst²

¹*Wireless Network Laboratory, CNR-ISTI, Via Moruzzi 1, 56124 Pisa, Italy*
²*Electronics Research Group, University of Aberdeen, U.K.*

SUMMARY

In the not so distant future, we envisage an Internet where the biggest share of capacity is used by streaming applications. To avoid congestion collapse from unresponsive flows calls for a robust and ubiquitous end-to-end multimedia congestion control mechanism, such as TCP-friendly rate control (TFRC), which provides fair sharing with the other Internet traffic. This paper therefore analyses the implications of using rate-adaptive congestion control over satellite links that utilize demand allocation multiple access (DAMA) to maximize satellite transponder utilization. The interaction between TFRC and DAMA is explored using simulations supported by fluidic flow models. The analysis shows that DAMA slows down the start-up phase of TFRC, causing non-negligible delays. To mitigate this problem, we propose a new cross-layer method based on the Quick-Start mechanism. This can accelerate the start-up of multimedia flows by a judicious allocation of additional capacity derived from cross-layer signalling. Copyright © 2009 John Wiley & Sons, Ltd.

Received 19 June 2008; Revised 23 October 2008; Accepted 22 December 2008

KEY WORDS: multimedia; TFRC; satellites; DAMA; Quick-Start

1. INTRODUCTION

The advent of services such as YouTube and the emerging market for Internet television lead us to envision a future where an important portion of Internet traffic originates from streaming applications. While emerging TV over IP services are typically provisioned within the boundaries of an ISP domain, we foresee applications that do not rely on network-centric quality of service (QoS) mechanisms. These applications will provide acceptable quality without requiring network guarantees. Despite an anticipated growth in traffic [1,2], such applications

*Correspondence to: Raffaello Secchi, Wireless Network Laboratory, CNR-ISTI, Via Moruzzi 1, 56124 Pisa, Italy.

†E-mail: raffaello.secchi@isti.cnr.it

Contract/grant sponsor: European Commission; contract/grant number: IST-27393

1 will be able to avoid posing a congestion threat to the Internet, by using congestion-responsive
3 methods that limit their sending rate to a share of the network capacity that is fair to both
streaming media and TCP flows. TCP-friendly rate control (TFRC) [3,4] is one such IETF-
defined congestion control method that has recently been standardized for streaming traffic.

5 At the same time of an increase in streaming media, we anticipate a continuation of the
current growth in bidirectional packet communications over satellite, especially systems based
7 on the digital video broadcasting (DVB) [5,6]. Bidirectional satellite communication will make
satellite Internet affordable, through the use of dynamic bandwidth allocation. Demand
9 allocation multiple access (DAMA) allows cost-efficient sharing of the satellite bandwidth and
supports bandwidth on demand strategies for bandwidth provisioning.

11 We consider the implication of employing TFRC over a satellite network using link-layer
DAMA. We accept the motivation behind using TFRC [7] for future Internet media services
13 including triple-play services [8] and analyse its consequence on a DAMA-based satellite
network, such as the DVB with return channel by satellite (DVB-RCS) system [9].

15 A previous study [10] identified a significant performance degradation during start-up of a
new TFRC-controlled multimedia flow or resumption after a pause in the media (with respect to
17 a satellite link using fixed bandwidth allocation). The experimental Quick-Start (QS) mechanism
[11,12] is also therefore considered. This provides opportunities for the network to accelerate
19 start-up and simulations have shown that this mitigates the problem induced by TFRC when
voice is used over a fixed-allocation satellite link [13]. In this paper, we broaden the scope of this
21 research by integrating QS and satellite DAMA. We later propose and analyse the performance
of a new method that we call QS DAMA, QSD, and show that this can almost completely offset
23 the performance degradation.

25 The remainder of this paper is organized as follows: In Section 2, we review the relevant
technologies: DVB, DVB-RCS, DAMA, Skyplex Data, TFRC, and QS. In Section 3 we use
simulation to assess the performance of TFRC over a fixed-allocation satellite network and
27 compare this with the performance for a DAMA system. We examine the impact of using QS
for a single TFRC streaming flow. Section 4 examines a scenario where several TFRC flows of
29 different types traverse a DAMA system. We analyse different methods of interaction between
DAMA and QS and compare their relative merits and identify possible future developments to
31 extend this analysis. The conclusions summarize the implications of our work.

33 2. TECHNOLOGY OVERVIEW

35 2.1. The DVB-RCS architecture

37 DVB is a mature technology used to send packet data over wireless media. It is particularly
popular in the DVB for satellite (DVB-S) format, where IP packets are encapsulated using
39 multi-protocol encapsulation [14] or the unidirectional lightweight encapsulation [15], before
being delivered over a satellite link as a stream of 188-byte cells.

41 We consider a bidirectional network serving a geographic group of terminals whose use of the
shared satellite capacity is controlled by a network control centre (NCC). The NCC periodically
43 broadcasts a terminal burst time plan (TBTP) defining a frame structure and identifying which
terminal may use each time slot in a frame. Terminals transmit using multi-frequency time
45 division multiple access (MF-TDMA) under the control of the TBTP. Several methods exist for
deciding how to allocate bandwidth to terminals. The simplest uses constant rate assignment

(CRA), where the rate for each terminal is determined by a human operator. More complex dynamic methods fall under the umbrella of DAMA methods. DVB-RCS defines two basic types of dynamic capacity request: rate-based dynamic capacity (RBDC) and volume-based dynamic capacity (VBDC) [16]. In addition, an NCC usually can distribute the unallocated capacity using a free capacity assignment (FCA) algorithm. Combining these methods, the rate requested by a terminal is

$$R_{\text{req}}(t) = \max \left\{ R(t) + \frac{q(t)}{T_s} + F, C \right\} \quad (1)$$

where the RBDC rate, $R(t)$, is the mean traffic rate entering the terminal from the terrestrial network over a time interval $[t-T_r, t]$ and T_r is the super-frame duration; the VBDC volume, $q(t)$, is the uplink queue length at time t ; T_s is the time; C is the constant CRA rate; and F is the instantaneous FCA rate allocated to the terminal. This equation is compliant with the request methods described in [9] in that it is a combination of all the four capacity request categories of DVB-RCS. Once the requested bandwidth is allocated, the terminal is able to empty its uplink queue in a time equal to T_s , while serving the traffic arriving at the terminal.

2.2. The Skyplex Data Platform

The commercial Skyplex Data system [17] forms the reference for this work. Though the results provided throughout the paper are based on Skyplex configuration, we believe that our method is applicable to any other DVB-RCS compliant system. Skyplex utilizes the Eutelsat HotBird6, which carries four Ka band transponders, each with a Skyplex Platform with an on-board processor (OBP) to perform stream multiplexing. The downlink employs DVB-S and the return channel (uplink) provides some features of DVB-RCS, forming a satellite mesh IP network supporting point-to-point or multicast transmission. Terminals can therefore communicate directly without needing a double satellite hop.

Each Skyplex Platform manages several channels configured in either a low rate (LR, 2.112 Mb/s) or a high rate (HR, 6.226 Mb/s) mode. Single channel per carrier is possible, although we consider DAMA in TDMA mode. The total number of channels is between 6 (all HR channels) and 18 (all LR channels). The time domain on the return channel is defined by the NCC in each super-frame in the TBTP. A super-frame consists of multi-frames divided into frames. We considered two configurations, the *small-frame configuration* (24 slots per frame) and the *large-frame configuration* (144 slots per frame). The latter is compliant with the Skyplex Data Modem.

Each frame is composed of six time slots (a burst of eight MPEG packets). Each terminal is assigned at least one slot per multi-frame, with additional slots requested using RBDC [10]. FCA is not used. This allows a maximum of six users, with a bandwidth assignment granularity of 352 kb/s for both LR and HR. The number of users can be extended by grouping 8 LR frames or 24 HR frames in a multi-frame, permitting $6 \times 8 + 6 \times 24$ users, with a bandwidth assignment granularity of 44 kb/s for both LR and HR. Finally, three multi-frames are grouped into an 820-ms long super-frame.

The capacity of the shared satellite channel is 6.336 Mb/s; therefore, a multi-frame is 45.6 ms for a small frame and 273 ms for a large frame. The allocated rate reduces exponentially following a reduction in the requested rate. The allocation delay is computed differently in the

1 two configurations, because in one case the propagation delay is greater than the TBTP period,
2 while it is smaller in the other case.

3 2.3. TFRC algorithm

4 At the transport layer, we consider a rate-adaptive multimedia application using TFRC. This
5 rate-paced equation-based congestion control algorithm requires the participation of both ends
6 of a connection [18]. A layer-4 control loop is therefore formed, where the sender receives
7 periodical feedback reports indicating the loss event rate measured at the receiver. The sender
8 uses this, together with the estimated round trip time (RTT) and loss rate (p), to calculate the
9 equivalent throughput that would have been achieved by a TCP connection under the same
10 conditions. The calculated throughput is used to adjust the sending rate every RTT [4].
11 Compared with TCP, this mechanism produces a much lower variation of throughput with time,
12 though at the expense of a slower responsiveness to fast changes in the available bandwidth.

13 Since the TCP throughput is largely dependent on the packet size, the TFRC algorithm may
14 fail to predict the fair share for a flow sharing the capacity with other flows using different
15 packet sizes. This happens, for instance, for VoIP flows, which tend to send small packets with
16 short inter-departure times. Our simulations therefore also consider the small packets variant of
17 TFRC (TFRC-SP) [18].

18 TFRC-SP is a minor variant of TFRC that seeks to achieve the same bandwidth as a TCP
19 flow using packets of up to 1460 B (the most common Internet packet size). To prevent a single
20 flow from sending small packets arbitrarily frequently, TFRC-SP also enforces a minimum
21 interval between packets.

22 Streaming applications make use of *transmit buffers* to queue the encoded media frames sent
23 by the application as packets are awaiting transmission by the transport layer. These may be
24 provided by the application or underlying network stack. Streaming audio (e.g. radio over IP)
25 may require a large transmit buffer to avoid packet losses, which may lead to a transmission
26 delay up to many seconds, whereas VoIP applications typically employ a transmit buffer of a
27 few media frames to constrain delay (e.g. less than 100 ms). A mismatch between the media rate
28 and the sending rate leads to packets being queued in the transmit buffer and could result in
29 buffer overflows. A TFRC receiver also usually uses a *de-jitter buffer* to remove the effects of
30 jitter before play-out.

31 2.4. The QS mechanism

32 At the network layer, we consider the addition of QS, which is an IETF experimental protocol
33 designed to provide lightweight signalling of the availability of unused capacity from the routers
34 to the communicating hosts [11,12]. QS was conceived to improve the performance of TCP bulk
35 transfers over lightly loaded network paths, but we suggest that it could also be useful for
36 multimedia flows, where we show that it can mitigate the effect of slow starting to the media
37 rate, either during the start-up phase or after an idle period [19]. All routers along a path must
38 approve a QS-Request and therefore all must support QS. While this is not necessarily an issue
39 with equipment directly connected to a satellite terminal or connected via a (private) network
40 that implements QS, there is as yet little support for QS in the general Internet.

41 Using QS, a host occasionally (e.g. at activation of a media flow) creates a QS-Request
42 indicating its desired sending rate (measured in bytes per second). This is sent using a QS option
43 in the IP header. Each router along the path may approve the requested QS-Rate, reduce it, or
44

1 indicate that the QS-Request is not approved. QS is not a reservation protocol: routers do not
 2 reserve bandwidth, nor do they specially treat subsequent packets from the same flow; when a
 3 router approves a request, it only asserts that at that moment there is sufficient free capacity for
 4 the requester to send at the requested rate without risking congestion. If the QS-Request is
 5 approved by *all* routers along the path, the sender receives a transport-level QS-Response,
 6 allowing it to send at the approved rate for an entire path RTT. QS only temporarily affects the
 7 flow to speed it up, then releases control to the default congestion mechanism. If the request is
 8 not accepted, QS does nothing, and the start-up algorithm of the default congestion mechanism
 9 is used. This creates a layer-3 control loop.

10 To participate in QS, we propose that the satellite terminal takes the role of a QS-enabled
 11 router. We also note that a TFRC receiver generally sends a feedback report as soon as the first
 12 packet is received during the start of a connection [18]. Therefore, a QS sender supporting
 13 TFRC must ignore the receiver rate reported in the first feedback report from the receiver, and
 14 stream at the accepted QS-Rate until the sender exits a QS validation phase [19]. Together, this
 15 enables a sender to use QS to probe the satellite sub-network about the available capacity and
 16 possibly to pre-allocate resources before a media flow starts. This involves cross-layer
 17 communications between TFRC at the transport layer, QS at the network layer, and the DVB
 18 allocation mechanism at the MAC layer.

21 3. PERFORMANCE OF TFRC OVER DAMA

22 This section analyses the interaction between TFRC operating at the transport layer and the
 23 service provided at the link layer using the DAMA mechanisms. It examines the reasons why
 24 TFRC performance is degraded by DAMA and proposes a scheme to offset this performance
 25 degradation.

27 3.1. The DAMA scheme

28 DAMA is considered using RBDC [9,16], with a request rate computed from the arrival
 29 rate at the uplink terminal [10]. The request and allocation are computed in the following
 30 way:

- 33 1. A terminal computes the rate of the incoming traffic over the period of a super-frame, T_r .
- 34 2. In the next multi-frame, the terminal sends a capacity request.
- 35 3. The NCC receives the requests from all terminals, and uses the received information to
 36 compute a bandwidth allocation. Bandwidth is assigned proportionally to each request,
 37 with a minimum equal to the allocation granularity. The NCC does not guarantee
 38 priority between terminals nor service quality. If the sum of the requests exceeds the
 39 available bandwidth, the value is normalized to the maximum capacity.
- 40 4. In the next super-frame the bandwidth allocation pattern is multicast in the TBTP.
- 41 5. The terminal uses the received TBTP to schedule its transmissions for the next super-
 42 frame (three multi-frames).

43 When a terminal requests an allocation smaller than the previous, the allocation is reduced
 44 exponentially. If $N_B(t_n)$ is the number of bursts allocated with the n th TBTP at the instant t_n and
 45 the local traffic detected in the T_r interval (t_n, t_n+1) corresponds to a number $N_L(t_n)$ of traffic

bursts, then at the $(n+1)$ th TBTP the bursts allocated are

$$N_B(t_{n+1}) = aN_B(t_n) + b[N_L(t_n) - N_B(t_n)] \quad (2)$$

This behaviour is described in [9], but no details are provided concerning the values of a and b . We assume $a = 1$, $b = 0.5$, thus obtaining an exponential decay with a time constant of two super-frames.

Figure 1 shows the sequence of messages exchanged between a terminal and the NCC for the bandwidth assignment. This forms a layer-2 control loop. The *allocation delay* (T_a) defines the period from the time a request is ready, at the end of the T_r period, and the availability of the desired bandwidth in the TBTP. During this period, the terminal uses an uplink queue to absorb any excess incoming traffic until the allocation is next updated.

3.2. The simulation scenario

Simulations were performed using *ns-2* with the TDMA-DAMA module described in [20]. The scenario modelled the HotBird6 satellite with Skyplex OBP and a propagation delay $T_p = 250$ ms. The network consisted of five terminals (including the NCC) with directly connected hosts that transmit media under the control of TFRC. A QS Agent at each terminal implements QS by processing a QS-Request, evaluating the available free capacity on the shared satellite channel, and then updates/forwards the QS-Request towards the client.

To assess the TFRC performance, we considered the *end-to-end delay* (E2ED) of packets forming a single TFRC flow. This is the difference between the time the sending application generates a packet and the time the packet is delivered to the de-jitter buffer of receiving application. Note that E2ED is larger than the *one way delay* (OWD) defined in [21], in that it also accounts for the queuing delay in the TFRC transmission buffer (see Figure 2).

In this paper we analysed only the case for constant media rate. Though actual application dynamics may depend upon the codec design, application behaviour, and encoded media content, constant rate applications are common. To have a complete picture of the dynamics of delay during the start-up period, each simulation was repeated, starting the flow with all possible offsets, in 1.8-ms steps, with respect to the super-frame.




Figure 1. Messages for bandwidth request and assignment.




Figure 2. One-way delay (OWD) and end-to-end delay (E2ED).

3.3. Performance of TFRC and QS

In this section we consider a 256 kb/s flow with a packet length of 1316 B, which is consistent with the typical output of the video LAN client (VLC) videoconferencing programme. We start our analysis by comparing the TFRC performance using a pre-assigned bandwidth of one-fifth of an HR frame (more than 4 times the bandwidth needed by the flow) and DAMA.

Figure 3 depicts the median of E2ED in *large-* and *small-frame* configurations with and without QS. The results confirm the findings in [10]: performance is acceptable with CRA because the constant RTT does not slow the TFRC self-clocking mechanism. When DAMA is used, the RTT increases due to the demand-assignment method, causing the TFRC-controlled sender rate to increase slowly. As a result, the application takes a much longer time to reach its target rate, than with the constant RTT presented by pre-allocated bandwidth. Since the sending application continues to generate data at the fixed media rate, a mismatch occurs between the media rate and the TFRC-controlled sending rate; in turn, this causes a queue building at the TFRC transmit buffer and a corresponding increase in the observed E2ED. In the CRA case, this leads to decreasing delay; in the DAMA case, however, traffic accumulates at the *uplink queue* until the DAMA request is granted to the requesting terminal.

The QS mechanism improves the performance by eliminating the sender's TFRC slow-start phase after the first RTT. During the first RTT, the delay increases due to traffic accumulating in the transmit buffer. When the QS-Request is accepted, TFRC starts to transmit faster than the application's media rate, thus emptying the transmit buffer.

The above considerations show that the bandwidth allocation flexibility of DAMA comes at a cost: there is a much slower start-up for protocols using slow-start, such as TCP and TFRC, especially when there is a large DAMA allocation delay. This cost is significantly reduced when using QS, but surprisingly remains non-negligible for large allocation delays.

Understanding the causes of inefficiency of QS with DAMA in the large-frame scenario requires a detailed analysis of the dynamics of the E2ED and OWD, shown in Figure 4. In this analysis the OWD reflects the time spent by packets in the uplink queue, while the difference between the E2ED and the OWD is the time a packet spends in the transmit buffer:

- (a) The application starts sending with a media rate of 256 kb/s and the QS Agent forwards a QS-Request for 320 kb/s towards the receiving host. TFRC starts its slow-start phase:




Figure 3. Single 256 kb/s flow with and without DAMA, with and without QS: (a) large-frame configuration and (b) small-frame configuration.

during this period, the TFRC sending rate is not sufficient to meet the application media rate, and packets are queued in the sender buffer.

- (b) After one RTT (0.8 s), the sender receives the QS-Response and increases its sending rate to the QS-Rate. Since the current DAMA allocation (44 kb/s) cannot satisfy the application media rate, a DAMA request is issued for the current rate, and any excess packets are queued in the uplink queue.
- (c) After discarding the first feedback report, the TFRC sender receives the second feedback after three RTT (at time 2.3 s). This report indicates a rate much less than the sending rate. TFRC therefore throttles the sending rate to the reported received rate.

The combination of a higher DAMA allocation and a lower sending rate causes the uplink queue to empty, triggering the DAMA to make a request for a lower rate. This causes the sender buffer to start filling up; at about the same time, the DAMA allocation is received.




Figure 4. Single 256 kb/s flow with DAMA and QS, large-frame configuration.

The subsequent interaction between the DAMA control loop and the TFRC control loop generates an oscillating behaviour whose exact dynamics depend on their relative phase (*phase effect*). The extent of oscillations is shown by the vertical bars that show the maximum and minimum delay for all possible relative phases.

The main problem with the QS implementation is therefore due to the late arrival of the requested DAMA allocation with respect to the QS-Response. This causes the TFRC receiver to issue a feedback report with a low perceived rate. In turn, this causes the TFRC sender to throttle its sending rate, just at the time when the requested bandwidth becomes available on the satellite channel, causing bandwidth oscillation and an increased overall delay.

Figure 4 also shows that the E2ED exhibits a significant delay variation (between 5 and 9 s) after a few seconds from the beginning of the transmission. Small delay variations at the beginning of a session lead to large variations in the E2ED. This is a result of the way the TFRC receiver estimates the received rate based on the number of packets received per RTT. This estimator suffers from a quantization effect when a small number of packets are transmitted. Both the E2ED and its variability were significantly reduced in the small-frame configuration.

3.4. Improving QS performance in satellite networks: QSD

To improve the performance of TFRC over DAMA using QS, we propose a new method that accelerates the availability of bandwidth to a terminal that receives a QS-Request. This method delays the QS-Request for a short time (T_q) before forwarding it along the data path so that the QS-Response and the DAMA bandwidth availability are synchronized at the transmitter. We call this cross-layer signalling method QSD. A QSD Agent at the uplink terminal that receives a QS-Request decides if there is sufficient free capacity on the shared satellite channel (Appendix A), it then adds the requested QS-Rate to the next DAMA request. More precisely, for each T_r interval, the request is computed as defined in (1); then, each QS-request received during the T_r interval is added to the request, provided enough satellite bandwidth is available. One problem with this method, analysed in the following section, is that the first data packet,

1 which transports the QS-Request, is also delayed. This causes an increase of T_q seconds in the
 2 initial RTT estimation at the sender, which may potentially degrade the TFRC start-up phase.

3 Figure 5 compares the E2ED using DAMA with different QSD strategies. The worst case
 4 (largest E2ED) is that without QS. The next worst is QS without any interaction with DAMA.
 5 The behaviour of the proposed QSD implementation with $T_q = 0$, that is, the QS-Request is
 6 forwarded immediately, is similar to that of QS. The time elapsed from the issue of a QS-
 7 Request to the time an allocation is received on the satellite channel is $T_r/2 + T_a$ on average,
 8 which corresponds to 2051 ms. The time that elapses between the issue of the QS-Request and
 9 the reception of a QS-Response, assuming that the delay due to the terrestrial networks is
 10 negligible, is the mean RTT, which is 777 ms. Therefore, in this case, the approved QS-Response
 11 arrives at the sender at a time when no DAMA capacity has yet been allocated, as observed in
 12 the QS case.

13 3.5. Analysis and optimization of QSD

14 The lowest curve of Figure 5 shows how the performance significantly improves when delaying
 15 the QS-Request compensates for the timing mismatch:

$$16 T_{q_{\max}} = T_r/2 + T_a - \text{RTT}_m \quad (3)$$

17 where RTT_m is the mean RTT seen by the TFRC connection. Figure 6 illustrates that this
 18 choice is effective for a range of flow rates: the maximum delay increases with flow rate for all
 19 strategies except for QSD, for which it remains low and constant in both the large- and small-
 20 frame configurations. Simulations, however, show that a value of T_q significantly smaller than
 21 $T_{q_{\max}}$ produces almost the same results (i.e. $T_q = 520$ and 55 ms in large- and small-frame
 22 configurations). Lowering T_q has beneficial effects especially in the multiple-connection scenario
 23 analysed in the following section since it reduces the probability of timeout. Indeed, delaying the
 24 packet carrying a QS-Request could result in the sender not receiving a response before the
 25 TFRC no-feedback timer expires. RFC 3448 specifies an initial RTO of 2 s. While the proposed
 26 value of T_q is short enough to prevent a no-feedback timer expiry, we believe that this initial
 27 value of T_q is short enough to prevent a no-feedback timer expiry, we believe that this initial
 28 value of T_q is short enough to prevent a no-feedback timer expiry, we believe that this initial
 29 value of T_q is short enough to prevent a no-feedback timer expiry, we believe that this initial




Figure 5. Single 256 kb/s flow with DAMA and various QSD strategies, large-frame configuration.

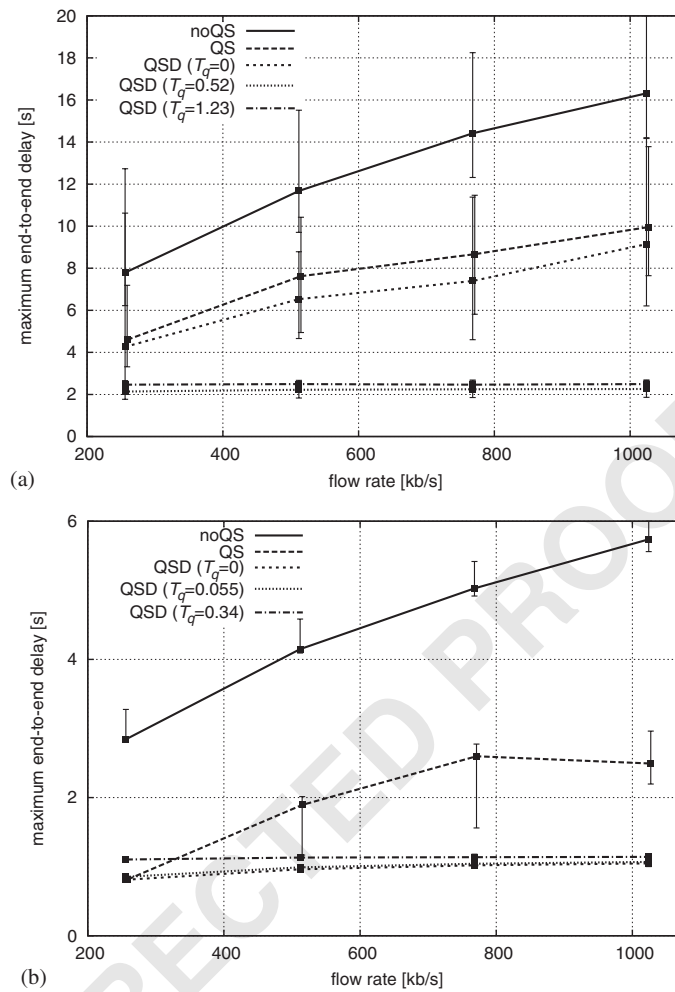


Figure 6. Single flow with DAMA and various QSD strategies: (a) large-frame configuration and (b) small-frame configuration.

RTO value is too low 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 3 s [22].

To determine a proper value for T_q , we developed a fluidic model (freely available in [23]) that reproduces the single-connection scenario. Using the fluidic model we have evaluated the proposed mechanism for a wide range of system parameters, but we have reported only the results that are more specific to the context of this paper. Figure 7 illustrates the output of the fluidic simulator for the large-frame configuration with $T_q = 400$ ms. The four monotone increasing curves show the cumulative outputs at four different points in the network: the traffic generator, the output of TFRC transmit buffer, the output of uplink queue, and the receiver. Vertical distances from one curve to the next indicate data in the TFRC buffer, the DAMA buffer, and the radio link. Horizontal distances indicate the delay introduced by TFRC in the transmit buffer, the delay introduced by DAMA in the uplink buffer, and the propagation plus framing delay RTT_m , respectively. The E2ED is the horizontal distance between the first and

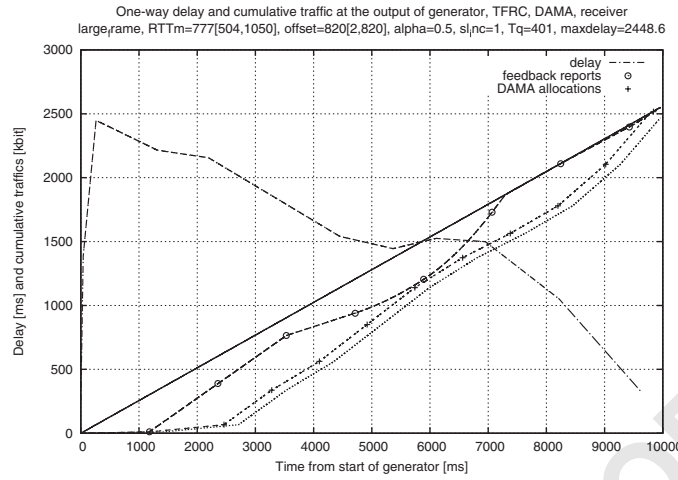


Figure 7. Fluidic modelling of traffic queuing in the large-frame configuration.

fourth curves. Though the fluidic model does not consider the effect of packetization and framing and the RTT estimate is fixed at $RTT_m + T_q$, the results provided by this model are very close to the results provided by the more detailed simulator.

By experimenting with the fluidic model for various T_q , we were able to identify a condition on T_q such that the E2ED keeps decreasing after the initial local maximum, rather than showing a second hump (visible in Figure 5 for the case of $T_q = 0$). We found that this condition is close to the optimality, hence it can be used to derive a simple expression to optimize T_q ($T_{q_{opt}}$).

The behaviour of the system using QSD is similar to that of QS, which is described in Section 3.3; the system passes through the steps (a)–(c) though now $RTT = RTT_m + T_q$. The second feedback report is sent at $t_{en} = 3RTT - RTT_m/2$. Since it is based on the traffic exiting the uplink queue during the interval

$$[t_{en} - RTT_m/2 - RTT, t_{en} - RTT_m/2] \quad (4)$$

and DAMA allocation is a very low during this period, and it reports a low receive rate. Thus, the TFRC sender reduces its rate until it receives the fourth feedback report. Then, as dictated by TFRC, it gradually increases the rate to the maximum of twice the fourth feedback report and twice the previous rate (third feedback report). Eventually, the media rate is reached after receiving the fifth feedback report. At this point, the uplink queue empties and the E2ED decreases.

To compute a safe value for T_q , we consider the initial received rate as zero, rather than the minimum DAMA allocation rate, and the final received rate as the rate of the sender, rather than the QS-Rate. The goal is for the TFRC rate at the fifth feedback report to equal the sender rate. This means that the rate of the third feedback report should be a half the sender rate, and therefore the first DAMA allocation should be received in the middle of the above interval (4), leading to

$$T_r + T_a = \frac{1}{2}[(t_{en} - RTT_m/2 - RTT) + (t_{en} - RTT_m/2)] \quad (5)$$

Thus, $T_{q_{opt}}$ is the minimum T_q for which the packet delay continues to decrease after the initial

1 maximum. Simplifying (5), we get

$$T_{q_{opt}} = \frac{1}{5}[2(T_r + T_a) - 3RTT_m] \quad (6)$$

3 Substituting configuration parameters (Table I), we get $T_{q_{opt}} = 520$ ms and $T_{q_{opt}} = 55$ ms for the
5 large-frame and the small-frame scenarios, respectively.

7 4. A SCENARIO WITH AGGREGATED TFRC FLOWS

9 The previous section analysed the dynamics of TFRC over DAMA by studying a single
11 connection on an otherwise idle satellite network. In this section, we continue to use the satellite
13 network and parameters described in Section 3.1, but now examine the case where the network
15 is lightly loaded, rather than unloaded. We consider three sets of five traffic flows streamed
17 between randomly selected hosts connected to a satellite terminal. The traffic generated by the
19 sources approximates VoIP calls, videophone calls, and video clips. Each flow has a Poisson
integration time, with the same parameters for each of the five traffic sources. For simplicity the
duration of the flows is exponentially distributed (this duration only determines the overall
background load, whereas the interesting dynamics involving TFRC occur in the first few
seconds of a flow's life).

21 The parameters of the three types of sources are given in Table II. Tolerable delay is the
23 maximum delay jitter that the transmitter considers acceptable, after which the packet is
25 considered late and discarded. VoIP applications employ a de-jitter buffer of a few media frames
27 to constrain delay (100 ms), whereas streaming audio and video require a large de-jitter buffer.
The packet length for 16 kb/s is chosen to be consistent with Skype; for 256 kb/s flows the length
represents that used by VLC. The transmit queue is derived from the two previous parameters.
Mean duration values are chosen for typical phone calls (VoIP), videophone calls, and video
clips (IPTV). The mean interval is the flow intergeneration mean time, chosen so that the share

29 Table I. Symbol definitions.

Parameter	Symbol	Small frame (ms)	Large frame (ms)
Allocation period	T_f	45.6	273
TBTP period	$T_b = 3T_f$	137	820
Measurement interval	$T_r = T_b$	137	820
Propagation delay	T_p	252	252
Allocation delay	$T_a = 6T_b, 2T_b$	820	1641
Mean RTT	$RTT_m = T_r + 2T_p$	550	777
Max. QSD delay	$T_{q_{max}} = (T_r/2 + T_a) - RTT_m$	339	1273

39 Table II. Streaming sources.

Rate (kb/s)	Tolerable delay (s)	Packet length (B)	Transmit buffer (pkt)	Mean duration (s)	Mean interval (s)	Traffic share (%)
16	0.1	160	2	60	1	50
64	0.1	512	2	180	40	15
256	5	1316	128	300	120	35


1 of each type traffic of traffic is equal to the value in the last column. The total link load is about
2 Mb/s, which is one-third of the shared capacity.

3 The simulation used TFRC-SP for the 16 kb/s flows and TFRC for the 64 and 256 kb/s flows.
4 This distinction is not significant in these scenarios, because the satellite network is underutilized
5 (far from congestion). Figure 8 shows the E2ED of the IPTV connections (256 kb/s) in the
6 aggregated flow scenario. QS brings considerable advantages in the E2ED, even for
7 a moderately congested network (one-third of the total capacity). The error bars show the 5th
8 and 95th percentiles over a total of 100 000 simulated VoIP flows, 2500 videoconferencing flows,
9 and 833 IPTV flows.

10 The improvement from using QSD instead of QS is smaller in this aggregated scenario
11 compared with the single-flow one; when the source starts sending packets at a high rate, the
12 terminal may already have bandwidth allocated in response to those requests that are triggered
13 by other traffic traversing the same terminal. This bandwidth is shared between all competing
14 flows, each of which suffers a little, but much less than the single flow suffered in the completely
15 unloaded scenario of the previous section. Additionally, the chosen DAMA algorithm does not
16 immediately relinquish the assigned capacity when a terminal reduces its request. In a dynamic
17 scenario, where connections continuously end, releasing bandwidth, this means that new
18 connections can often exploit unused capacity still allocated to the terminal.

19 As a consequence, this scenario offers better performance when not delaying the QS-Request
20 ($T_q = 0$). However, the single-connection scenario analysed in the previous section, while
21 not general, reflects an important situation when a single host or a small network is connected
22 to the Internet through a satellite connection and a multimedia flow is started on an otherwise
23 idle network. For this reason it is useful to find a value for T_q that is as small as possible,
24 yet allows QSD to provide an advantage to the single-connection scenario. The one provided by
25 (5) provides a good compromise as the E2ED measure is midway between $T_q = 0$ and
26 $T_q = T_{q_{max}}$.

27 The variability of E2ED is a critical parameter to size the de-jitter buffer at the receiver. After
28 the first transient phase in which high delay variations are experimented, the jitter stabilizes to a
29



31
33
35
37
39
41
43
45
Figure 8. Start-up behaviour for 256 kb/s flow in the aggregated flow scenario.

1 very low threshold that depends only on the delay required to schedule concurrent flows over
 3 the frame. Figure 9 reports the distribution of the E2ED for the three classes of flows in the first
 5 30s of their lifetime. Figure 9(a) shows that the choice of strategy has little influence on the




Figure 9. Distribution of E2ED for: (a) VoIP; (b) videoconference; and (c) IPTV connections.

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 ($T_q = 0$)	0.036	0.037	0.036	0
QSD ($T_q = 0.52$)	0.020	0.020	0.020	0
QSD ($T_q = 1.23$)	0.012	0.012	0.012	0

delay behaviour of VoIP, due to the low bandwidth used, which generally allows VoIP to profit from any excess allocated DAMA bandwidth. The same happens for 99% of the videoconference call packets (Figure 9(b)), where the only noticeable differences are the tails for $T_q = 0.52$ and 1.23 s. This resembles the case for $T_q = 0$ shifted by T_q seconds: the packets that create the tails are the first packets in the flows, those that are delayed by T_q . Figure 9(c) shows the graph relative to IPTV flows, where the performance when not using QS is worse than others, and $T_q = 0.52$ and 1.23 s are the curves with the sharpest descent, being copies of one another, apart from the second being shifted to the right (i.e. with highest delay) by $1.23 - 0.52$ s. Notice how Figure 9(c) shows that the choice of $T_{q_{opt}}$ is indeed the best among the proposed choices.

Finally, Table III gives the fraction of packets lost due to transmit buffer overflow when TFRC throttles the output rate of the sender during slow-start, while the media arrives at fixed rate. A transmit buffer set to tolerate a 5-s delay (e.g. for IPTV) is able to absorb all delays when QS is used in the simulated scenario. No packets were lost as a result of buffer overflow in the terminal uplink queue of 1000 packets (the value used in Skyplex Data equipment).

4.1. Further scenarios

Sections 3 and 4 presented a detailed analysis of some typical TFRC applications connected via QS-enabled networks whose transit delay is significantly less than the total satellite transit delay. This subsection identifies additional topics that we have not explored in this paper.

We did not model the effect of jitter in the packets arriving at the uplink terminal. In practical networks, such jitter could arise from implementation constraints in TFRC, statistical multiplexing of the links before the terminal, or the effects of an adaptive variable rate media codec that tunes performance to available capacity. These effects may result in QS packets arriving in bursts by a terminal. The effects of allocation strategies and traffic multiplexing at the terminal could further contribute to burstiness and could potentially aggravate downstream congestion. More detailed application, transport, and network models could be used to examine these effects.

We did not model the effects of significant delay in the terrestrial network. One side-effect of using QSD is that the QS-Request triggers a request for capacity in advance of actual data arriving at the terminal: if the sender does not immediately utilize its approved QS-Rate, the early additional allocation could be wasted. Our analysis has examined hosts connected via QS-enabled networks whose transit delay is significantly less than the total satellite transit delay. We observed that in this case the performance using QS improves when the QS Agent delays forwarding the QS-Request. However, using this optimization with a sender or a receiver connected via an intermediate network with appreciable delay would also delay the arrival of the media flow (QS packets) at the terminal, which could result in the additional capacity being

1 **AQI** allocated before the media flow starts, which is not being utilized. Future work will refine the
3 mechanism to address these issues.

7 5. CONCLUSIONS

9 This paper has analysed the effect of using satellite DAMA on the performance of streaming
11 multimedia flows that utilize the TFRC algorithm. The use of TFRC congestion control is
13 motivated by the desire to improve co-existence of TCP with UDP streaming flows, and
15 therefore avoid congestion collapse. The varying delay introduced by DAMA is shown to
significantly impact the performance of TFRC, even in the absence of congestion losses. This
results from interactions between the end-to-end congestion control loop and the control loop
used by the satellite link to request/allocate capacity to a specific terminal. Similar effects have
previously been observed for TCP sessions.

17 A method is proposed using Quick-Start (QS) to mitigate these interactions. This does not
19 require explicit QoS reservation across the network path, but instead allows the network
21 layer to inform the link layer about the expectations of a new media flow. We have
23 shown that this provides an opportunity for a terminal to request an increase in allocated
capacity to accommodate a new flow. QS can significantly improve performance,
especially when the existing capacity allocated to a terminal is much less than the volume
of new traffic being added by the multimedia flow (e.g. an idle terminal or one with little
existing traffic).

25 We define a new cross-layer method known as QS DAMA (QSD), which updates the use of
27 QS at the terminal by purposely delaying forwarding of the QS-Request. This partially aligns
29 reception of a QS-Approval with the reception of the additional capacity allocated by the
31 DAMA. Fluid-flow analysis has been used to derive appropriate values for this delay, showing
33 that it can improve performance. However, delaying the QS-Request also has a disadvantage
that it delays the end-to-end QS mechanism, resulting in the sender measuring an inflated
measured initial RTT. This can paradoxically result in preventing the flow taking full advantage
of the QS capacity (since a QS flow must pace its transmission according to the RTT). We
suggest that judicious choice of when to employ the extra delay can alleviate this problem in
common scenarios (e.g. hosts connected to the satellite network via low-latency terrestrial
networks).

35 Although our analysis employed parameters appropriate to the Skyplex Platform,
37 we do expect these techniques to also be applicable to other similar systems, including
DVB-RCS.

41 APPENDIX A: ESTIMATION OF UNUSED CAPACITY ON THE SATELLITE 43 CHANNEL

45 When the QS Agent at the terminal receives a QS-Request, it must predict whether there will be
enough capacity on the shared satellite channel that could be allocated to satisfy the received
request. Terminals may become aware of current bandwidth utilization using a centralized
approach. A discussion of the approaches follows:

1 *A.1. Centralized estimation*

3 *A.1.1. Continuous broadcast.* The NCC broadcasts the amount of free capacity, e.g. by adding
5 an additional field to the TBTP, or by sending a packet/frame to all the terminals in a private
7 format.

9 *A.1.2. On demand.* The NCC answers queries from the terminals requesting the amount of free
11 capacity in a private format. This method implies an additional delay equal to RTT_m for the
13 QS-Request to be forwarded.

15 *A.2. Distributed estimation*

17 Each terminal individually evaluates the amount of free capacity. There are two possibilities
19 depending on whether the DAMA uses FCA.

21 *A.2.1. DAMA does not use FCA.* Estimation is performed by examining the
23 TBTP and summing the capacity of all unassigned slots (requiring these to be identified in
25 the TBTP).

27 *A.2.2. DAMA uses FCA.* A terminal may listen to another terminal's transmissions
29 (limited by the topology and capabilities of the network and the link power budget of the
31 terminal).

33 *A.2.2.1. Single spot without fade countermeasure.* The estimation is possible but inaccurate,
35 because a terminal may not transmit in a slot it has requested, or could be transmitting in a slot
37 it has not requested, but was assigned by FCA.

39 *A.2.2.2. Single spot with fade countermeasure.* If up-power control or adaptive modulation
41 and coding are used as fade countermeasure systems, a terminal in deep fading may be unable to
43 listen for transmissions directed to other terminals. This would lower the accuracy of estimation
45 during fading.

A.2.2.3. Multiple spots. If the satellite network is multi-spot, each terminal can only listen for
transmissions directed to its spot, and so can only have an estimate of the unused capacity for
transmissions directed to its own spots.

Our simulation assumes that the free capacity can be determined without delay; therefore, we
consider case A.1.1. In [12] 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. We make a new observation each TBTP; therefore, the estimation interval is
 T_b . We considered only the last estimation; therefore, 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 [12]. The agent considers the channel to be underutilized if the free capacity is at least
20% of the shared bandwidth; this means setting $qs_thresh = 0.8$ [12]. To avoid overallocation
of bandwidth to several consecutive requesters, the QS Agent remembers requests that were
approved in the last M time intervals. M was set to 2, and the variable $state_delay$ contains the
length of the time interval. The agent remembers approval of a QS-Request for the maximum

1 time that can elapse from the reception of the request to the time it is reflected in the received
 2 TBTP ($T_r + T_a$; therefore, we set $\text{state_delay} = (T_r + T_a)/2$).

5 ACKNOWLEDGEMENTS

7 This work was supported by the European Commission under the European Satellite Communications
 8 NoE (SatNEx II, IST-27393) within the Sixth Research Framework Programme.

13 REFERENCES

- 15 1. Gill P, Arlit M, Li Z, Mahanti A. YouTube traffic characterization: a view from the edge. *ACM Internet Measurement Conference*, San Diego, CA, 2007.
- 17 **AQ2** 2. Cherkasova L, Gupta M. Analysis of enterprise media server workloads: access patterns, locality, content evolution, and rate of change. *IEEE/ACM Transactions on Networking* 2004; **12**(5).
- 19 3. Floyd S, Fall K. Promoting the use of end-to-end congestion control in the Internet. *IEEE/ACM Transactions on Networking* 1999; **7**(4).
- 21 4. Floyd S, Kohler E, Padhye J. Profile for datagram congestion control protocol (DCCP) congestion control ID 3: TCP-friendly rate control (TFRC). *IETF RFC 4342*, 2006.
- 23 5. EN 301 421. *DVB Modulation and Coding for DBS Satellite Systems at 11/12 GHz*. ETSI, 1999.
- 25 6. EN 302 307. *DVB: Second Generation Framing Structure, Channel Coding and Modulation Systems*. ETSI, 2006.
- 27 7. Kohler E, Handley M, Floyd S. Designing DCCP: congestion control without reliability. *ACM SIGCOMM*, Pisa, Italy, 2006.
- 29 8. SATSIX Project. www.ist-satsix.org.
- 31 9. EN 301 790. *DVB Interaction Channel for Satellite Distribution Systems*. ETSI, 2003.
- 33 10. Gotta A, Secchi R, Potorti F. An analysis of TCP startup over an experimental DVB-RCS platform. *Proceedings of the International Workshop on Satellite and Space Communications (IWSSC)*, Madrid, Spain, September 2006; 176–180.
- 35 11. Floyd S, Allman M, Jain A, Sarolahti P. Quick-Start for TCP. *IETF RFC 4782*, January 2007.
- 37 **AQ3** 12. Sarolahti P, Allman M, Floyd S. Determining an appropriate sending rate over an underutilized network path. *Computer Networks* 2007; **7**(51), to appear.
- 39 13. Sathiaselvan A, Fairhurst G. Using Quickstart to improve the performance of TFRC-SP over satellite networks. *IEEE International Workshop on Satellite and Space Communications (IWSSC)*, Madrid, Spain, 2006.
- 41 14. EN 301 192. *Digital Video Broadcasting (DVB), Specification for Data Broadcasting*, 2003.
- 43 15. Fairhurst G, Nocker C. Unidirectional lightweight encapsulation (ULE) for transmission of IP datagrams over an MPEG-2 transport stream (TS). *IETF RFC 4326*, December 2005.
- 45 16. Lexow HP, Navekvién TA, Paxal VM. Satellite resource management combining distributed cooperative methods and central control using BoD variants of DVB-RCS. *ICSC*, Rome, Italy, 2005.
- AQ4** 17. Feltrin E, Weller E, Martin E, Zamani K. Design, implementation and performance analysis of an on board processor-based satellite network. *International Conference on Communications*, vol. 6, 2004; 3321–3325.
18. Floyd S, Kohler E. TCP friendly rate control: the small packet (SP) variant. *IETF RFC 4828*, April 2007.
19. Fairhurst G, Sathiaselvan A, Quick-Start for DCCP. *IETF Internet Draft*, draft-ietf-dccp-quickstart-01.txt, September 2008, work in progress.
20. Gotta A, Secchi R, Potorti F. Simulating dynamic bandwidth allocation on satellite links. *Proceedings of the Valuetools*, Pisa, Italy, October 2006. *WNS2 Workshop*.
21. Almes G, Kalidindi S, Zekauskas M. A one-way delay metric for IPPM. *IETF RFC 2679*, September 1999.
22. Paxson V, Allman M. Computing TCP's retransmission timer. *IETF RFC 2988*, November 2000.
23. Fluid Simulator. <http://wnet.isti.cnr.it/software/damatfrc.m>.
- AQ5** 24. Kohler E, Handley M, Floyd S. Datagram congestion control protocol (DCCP). *IETF RFC 4340*, March 2006.

AUTHORS' BIOGRAPHIES



Raffaello Secchi received his Laurea degree in Telecommunication Engineering in July 2002 from the University of Pisa, Italy. From January 2003 to December 2005 he was a PhD student at the University of Pisa and received his PhD in May 2006, by defending a thesis on Traffic Modelling and Control in High Speed Networks. Since 2006 he was a postdoctoral fellow with the Information Science and Technology Institute (ISTI) of National Research Council in Pisa, and is currently a research fellow at the University of Aberdeen.



Arjuna Sathaseelan (PhD 2005, BE 2000) is a research fellow at the University of Aberdeen. His current research interests are in the integration of satellite broadband access in next-generation networks. He leads work on network simulation, design of multimedia transport protocols, TCP protocol optimization, and performance analysis. He is the author of several Internet drafts in the IETF DCCP working group.

Francesco Potorti received his electronic engineering degree from the University of Pisa in 1991. He is currently a full-time researcher at the ISTI-CNR institute (formerly CNUCE) in Pisa, Italy, where he worked since 1989 in the fields of satellite communication protocols and fade countermeasure systems, and where he is a member of the Wireless Networking Laboratory. He was Italian representative in the European COST action 285.



Alberto Gotta received MS degree from the University of Genoa in 2002. From 2002 to 2004 he was a researcher with the Research Unit of the Italian National Consortium for Telecommunications (CNIT) at the University of Genoa. Since 2004 he has been a researcher with the ISTI Institute of the Italian National Research Council (CNR). He has worked in the SatNEx project funded under the Network of Excellence. He is currently involved in the SatNEx II project and in the EU-funded project RINGrid—Remote Instrumentation in Next-generation Grid.

1
3
5
7
9
11
13
15
17
19
21
23
25
27
29
31
33
35
37
39
41
43
45



Gorry Fairhurst is currently a reader in Communications at the School of Engineering at the University of Aberdeen, Scotland. He leads a research group, whose research includes Internet engineering, link ARQ, satellite-specific tuning of communication protocols, and development of reliable multicast transport protocols. He is also an active participant in the specification and engineering of standards for Internet transport and design of IP-friendly links. He contributes to the work of the IETF and ETSI and currently chairs the IETF IPDVB and DCCP working groups.

UNCORRECTED PROOF