

Multi-Connection TFRC Video Streaming in a Concatenated Network: Latency and Video Quality

Salah S. Al-Majeed¹ and Martin Fleury²

¹London School of Commerce, United Kingdom, salah.saleh@lsclondon.co.uk

²University of Essex, Colchester, United Kingdom, fleum@essex.ac.uk

Abstract. Sending a single video stream over multiple TCP-Friendly Rate Control (TFRC) connections is a promising lightweight way of coping with wireless channel losses *and* traffic congestion in a concatenated network (one consisting of a broadband wireless link and a wired all-IP network). Multi-connections introduce the need for data re-ordering at the receiver. This paper considers the potential delay at an uplink destination on the Internet. It also considers the impact on video quality of packet drops due both to channel loss and router buffer overflow, when the TFRC congestion controller is applied. Results for an IEEE 802.16e (mobile WiMAX) link show a worthwhile gain in video quality from using three or more connections over a single connection but with start-up delay in the multi-connection case due to the need to avoid possible buffer underflow.

1 Introduction

In this paper, we employ a form of MULTTFRC [1] with multiple TCP-friendly Rate Control (TFRC) [2] connections to stream video across a concatenated network. Such all-IP or Next Generation Networks are being widely developed (for instance in the UK in BT's 21CN) as a cost-effective replacement for traditional telephony networks. A concatenated network combines an access network with a core network that may consist of heterogeneous sub-networks. We assume a broadband wireless access link with a core wired-network. Specifically, IEEE 802.16e (mobile WiMAX) [3] is modeled as the broadband wireless link. In such a network, a video stream is subject to packet loss due to wireless channel conditions *and* to traffic congestion on the wired network, as well as congestion on the access network.

In video streaming across an all-IP network (one in which the IP packet format is universal though the Multi-Protocol Labeling System may provide circuit-switched routing in the core), unreliable UDP transport serves to reduce delay at the expense of some packet loss, while application-layer TCP emulation [4], such as TFRC, acts as a form of cooperative congestion control (assuming most other traffic is carried through TCP transport). However, TCP emulation by the application is *not* the same as TCP. TCP itself is unsuitable for delay-variation intolerant video streaming, because it introduces unbounded delay in support of a reliable service. Instead, TCP emulation mimics the average behavior of TCP, but is not 'reliable' and does not result in the 'saw-tooth'-like rate fluctuations that arise from TCP's aggressive congestion control algorithms. The latter can cause disconcerting quality fluctuations at an end-user's display if the streaming quality is varied according to the congestion level.

In multi-connection TFRC video streaming, a *single* video source is multiplexed onto several connections across the wireless link in order to improve the wireless channel utilization, resulting in an increase in throughput. TFRC's main role when congestion occurs is to reduce the video streaming data rate across the wired portion of the concatenated network. It does this in response to packet drops at intermediate routers, which signal the presence of contending traffic. Unfortunately, TFRC can misinterpret as congestion packet losses due to wireless interference and noise. Though cross-layer approaches to avoid misinterpretation are possible, these are complex to implement and inflexible. By multiplexing a video stream across multiple connections it is hoped that the impact of packet loss on one or more of these connections will be mitigated by the rate across the remaining connections.

There is widespread interest in interactive IPTV (it is a goal of BT's 21CN). In Brazil already, mobile WiMAX is the basis of a networked digital TV service and uplink interactive services are in active development [5]. We ask what would occur if multiple TFRC connections were opened in the uplink (UL) from a WiMAX subscriber station (SS) to base station (BS) in the presence of cross traffic from other mobile SSs. Thus congestion also occurs on the uplink as well as fluctuating wireless channel conditions. In this situation, congestion will occur at the WiMAX real-time polling service (rtPS) queue and packet loss will occur over the wireless channel. We also consider the effect of packet loss as the multiple connections pass over the core IP-network when other traffic sources contend for access to buffers at intermediate routers. Propagation over the wired network is realistically assumed to be error-free, as it may well consist of optical fiber links.

This paper's main contribution is the finding that as the number of connections increases, reduced packet loss leads to improved video quality, because of the reduced sending time in sending the same video data. In contrast, in [1] improved video quality comes by increasing the quantity of video data that can be sent over the multiple connections. Of course, increased video data implies a lower compression ratio and, hence, higher quality. Unfortunately, if the number of connections varies, as it does in [1, 6, 7] then sending rate oscillations occur. If the compression ratio was varied at the source (either by changing the quantization parameter at the codec if live video or through a bit-rate transcoder) then oscillations in rate again run the risk of disconcerting changes in displayed video quality. However, we show that the quality increases anyway without the need to change the compression ratio and by keeping the number of connections constant. This is because with multiple TFRC connections, TFRC is better able to control its sending rate. In fact, TFRC [2] was designed for a high number of streams and has special measures if the number of streams is not high. Possibly, the difference in findings occurs because in work on MULTTFRC [1, 6, 7], apparently no account of the impact of cross-traffic occurs except to test the fairness of the scheme to coexistent traffic.

In our approach, video data is statically multiplexed onto the TFRC connections. The unit of multiplexing was taken to be a Group-of-Pictures (GOP) [8], with an Intra-refresh rate of 15. Just as in Peer-to-Peer video streaming, when video is delivered as chunks from a number of sources, there is a need to employ a reordering buffer. As a result, the start-up delay in the scenario tested was about 6 s, but the gain

in video quality (PSNR) compared to using a single connection was over one dB, a worthwhile gain. Start-up delay may be attributable to features of TFRC itself, which implies that a modified TFRC or an alternative congestion controller may reduce the delay. Again earlier work did not give much consideration to the effect of congestion in the TFRC feedback path, which we also now consider.

2 Scenario investigated

The scenario tested in this paper is shown in Fig. 1. The following describes the WiMAX part and this description is followed by a description of the inset, showing traffic sources and sinks within the core IP network.

2.1 WiMAX system

In Fig. 1, once a BS has allocated bandwidth to each SS, each SS must manage its queue according to the data arrival rate from user applications. In WiMAX Point-to-Multipoint (PMP) mode, there is no SS-to-SS communication unless it is via the BS. WiMAX networks support multiple service classes to accommodate heterogeneous traffic with varying requirements. WiMAX's rtPS is most suitable for real-time video services, particularly for Variable Bitrate Video (VBR), which is employed to maintain delivered video quality but may lead to 'bursty' arrival rates. Other congesting traffic is assumed to enter the non-real-time Polling Service (nrtPS) queue at the SS. In our experiments for both queues, a drop-tail queuing discipline was

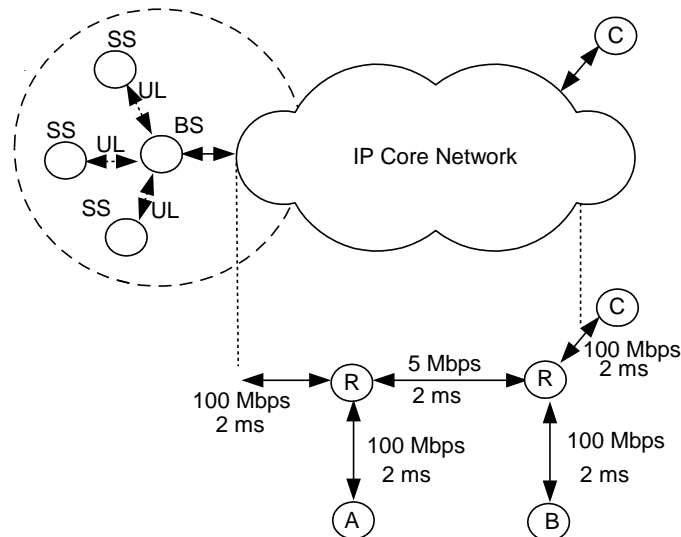


Fig. 1. Concatenated network with inset showing routing across the core network, A, B and C being sources and sinks, and R = router

simulated. Queue sizes were all set to fifty packets. This value was selected as it seems appropriate to mobile, real-time applications for which larger buffer sizes might lead both to increased delay and also greater active and passive energy consumption at the buffer's memory.

The WiMAX system operating in PMP mode was simulated by well-known ns-2 simulator (v. 2.29) augmented by a WiMAX module [10]. The simulator is allowed to reach steady-state over 20 s with other traffic passing over the network. The PHY settings selected for WiMAX simulation are given in Table 1, with additionally the MAC settings defaulted from [10]. The DL/UL ratio is not intended to be realistic but to aid in testing multiple-connection TFRC, as in practice the DL would be allocated the majority of the bandwidth. The antenna is modeled for comparison purposes as a half-wavelength dipole. The Gilbert-Elliott 'bursty' channel model is further explained in Section 2.5. The frame length is significant, as a longer frame reduces delay at the MS by permitting more data to be removed from any queues at each polling time. The value of 20 ms is at the high end of the available durations in the Standard [3] in order to reduce this source of queuing delay for real-time video streaming.

Table 1. Simulated WiMAX settings, OFDMA = Orthogonal Frequency Division Multiple Access, QAM = Quadrature Amplitude Modulation, TDD = Time Division Duplex

<i>Parameter</i>	<i>Value</i>
PHY	OFDMA
Frequency band	5 GHz
Duplexing mode	TDD
Frame length	20 ms
Max. packet length	1024 B
Raw data rate	10.67 Mbps
IFFT size	1024
Modulation	16-QAM 1/2
Guard band ratio	1/8
DL/UL ratio	1:3
Path loss model	Two-ray ground
Channel model	Gilbert-Elliott
MS transmit power	250 mW
BS transmit power	20 W
Approx. range to MS	0.7 km
Antenna type	Omni-directional
Antenna gains	0 dBD
MS antenna height	1.5 m
BS antenna height	32 m

2.2 WiMAX traffic characteristics

There were three SSs communicating to the BS, with one of the SS sending a VBR video sequence encoded with the H.264/Advanced Video Codec (AVC) [11] and split between the multiple TFRC connections. The other SSs are simply introduced as sources of competing traffic across the wireless link and do not indicate the likely size of a WiMAX network, which obviously could be larger. A trace file was input to ns-2 and packet losses recorded in the output. The output serves to calculate the PSNR. Video quality comparisons were made under the EvalVid environment [12]. As a test, we used the ‘Paris’ clip H.264 VBR-encoded at 30 Hz (frame/s) at Common Intermediate Format (CIF) (352 × 288 pixel/frame) with initial quantization parameter set to 26 (from a range 0 to 51). The slice size was fixed at the encoder to be a maximum of 900 B. Paris consists of two figures seated around a table in a TV studio setting, with high spatial coding complexity. H.264’s Baseline profile was selected, as this is more easily supported by mobile devices because of its reduced computational overhead. The Intra-refresh rate was every 15 frames with IPBB...I structure, i.e. the GOP size was 15. 1063 frames were transmitted. Simple Previous Frame Replacement (PFR) was set for error concealment at the decoder.

Table 2 records the simulated traffic characteristics for the three SSs communication with the BS. Network Adaptation Layer Units (NALUs) from the H.264 codec were encapsulated with Real Time Protocol (RTP) headers. After the addition of IP headers, these in turn formed a single WiMAX MAC Packet Data Unit (MPDU), which are variable-sized WiMAX packets. For simplicity, a WiMAX MPDU is now referred to as a packet.

For TFRC, the inter-packet sending time gap was varied according to the TFRC equation [2], not the simplified version reported in [7]. As described in [2], TFRC is a receiver-based system in which the packet loss rate is found at the receiver and fed-back to the sender in acknowledgment messages. The sender calculates the round-trip time from the acknowledgment messages and updates the packet sending rate. A throughput equation models TCP New Reno to find the sending rate:

$$TFRC(t_{rtt}, t_{rto}, s, p) = \frac{s}{t_{rtt} \sqrt{\frac{2bp}{3}} + t_{rto} \min\left(1, 3\sqrt{\frac{3bp}{8}}\right) p(1 + 32p^2)} \quad (1)$$

where t_{rtt} is the round-trip time, t_{rto} is TCP’s retransmission timeout, s is the segment size (TCP’s unit of output) (herein set to the packet size), p is the normalized packet loss rate, w_m is the maximum window size, and b is the number of packets acknowledged by each ACK. b is normally set to one and $t_{rto} = 4t_{rtt}$. It is important to notice that t_{rto} comes to dominate TFRC’s behavior in high packet loss regimes [2], which is why it is unwise to use a simplified form of (1). General inspection of (1) indicates that if the round-trip time and/or the packet loss rate (the two independent variables in the denominator of (1)) increase then the throughput reduces.

Table 2. Simulated WiMAX traffic characteristics

<i>SS-UL</i>	<i>Service type</i>	<i>Traffic type</i>	<i>Protocol</i>	<i>Packet Size (B)</i>
1	rtPS	VBR (video)	Multiple TFRC	Variable
		CBR	UDP	1000
2	nrtPS	FTP	TCP	
	rtPS	CBR	UDP	1000
3	nrtPS	FTP	TCP	
	rtPS	CBR	UDP	1000
<i>SS-DL</i>	nrtPS	FTP	TCP	
	1,2	rtPS	CBR	UDP
3	nrtPS	FTP	TCP	

In our variant to standard TFRC, the packet size, s , in the TFRC equation was dynamically altered according to the EvalVid-created trace file sizes. This variant makes for more responsive control rather than the mean packet length employed in the original TFRC formulation [2]. TFRC was originally intended for video-on-demand applications, when it is feasible to calculate the mean packet length. Setting a mean packet length is inappropriate for interactive multimedia applications. The underlying TFRC transport protocol was set to UDP, as is normal.

Coexisting rtPS queue CBR sources were all sent at 1500 kbps, i.e. at a similar rate to the video source. The inter-packet gap was 0.03 s for the CBR traffic. The FTP applications, which continuously supplied data according to available bandwidth, were set up out of convenience as a way of occupying the nrtPS queues; otherwise a Best-Effort (BE) queue might be more appropriate. Likewise, the DL traffic is simply selected to fully occupy the DL link capacity.

2.3 Core network traffic characteristics

In Fig. 1, all links except a bottleneck link within the core network are set to 100 Mbps to easily accommodate the traffic flows entering and leaving the network. The link delays are minimal (2 ms) in order to avoid confusing propagation delay with re-ordering delay. A bottleneck link with capacity set to 5 Mbps is set up between the two routers. The buffer size in each router was set to 50 packets. This arrangement is not meant to physically correspond to a network layout but to represent the type of bottleneck that commonly lies at the core network edge before entry into a corporate or campus network.

Node A sources to node B a CBR stream at 1.5 Mbps with packet size 1 kB and sinks a continuous TCP FTP flow sourced at node B. Node B also sources an FTP flow to the BS and a CBR stream at 1.5 Mbps with packet size 1 kB (see Table 2 downlink). Other SS sources apart from the video connections do not pass over the core network shown but are assumed to be routed elsewhere after passing the WiMAX BS. Node C in Fig. 1 is the sink for the TFRC multiple connections.

2.4 Management of connections

To systematically test the effect of multiple TFRC connections the number of TFRC connections was incrementally stepped up in successive experiments. In MULTTFRC itself, the number of connections is changed over time according to the average round-trip time of all the connections, but this hides the interpretability of results. As remarked earlier, it is also unclear from [1, 6, 7] how a single video stream would be apportioned between a varying number of connections. In our experiments, a single queue was segmented into GOPs (15 frames). Each connection was statically allocated its GOPs, which are taken in interleaved manner from the video sequence. As previously mentioned, this assumes that a re-ordering buffer is available at the receiver.

2.5 Channel model

A Gilbert-Elliott two-state, discrete-time, ergodic Markov chain [13] modeled the wireless channel error characteristics at the ns-2 physical layer. The probability of remaining in the good state was set to 0.95 and of remaining in the bad state was 0.94, with both states modeled by a Uniform distribution. The packet loss probability in the good state was fixed at 0.01 and the bad state default was 0.05. However, the bad state packet loss probability, P_B , was also varied as [0.01, 0.02, ..., 0.1]. In this way, we were able to judge the effect of worsening burst error channel conditions.

3 Evaluation

Initial investigations considered the WiMAX link alone in Fig. 1. Table 3 shows the average data-rate over time when transmitting the Paris clip over multiple connections, for two different WiMAX frame sizes: the default from Table 1 and 5 ms (frame duration code 2 in the Standard [3]). Clearly, TFRC is able to multiplex more data onto a link as the number of connections increases, though observation of a time-wise plot of throughput shows that during transmission TFRC sharply reduces its overall sending rate in response to packet loss. Because the sending period for one connection is more than the display period of Paris with the shorter frame duration, the longer frame duration is clearly preferable.

Fig. 2 plots the video stream packet drop rate relative to channel packet error rate. Included in the percentages in Fig. 2 are any additional packet losses arising from buffer overflow at the SS caused by the SS packet scheduler being otherwise occupied servicing the rtPS queues in the three SSs. In this Figure, the shorter frame size is

Table 3. Sending periods and throughputs from the video streaming MS to the WiMAX BS.

<i>No. of connections</i>	<i>SS to BS (s)</i>		<i>Throughput</i>	
	<i>frame size 5 ms</i>	<i>(kbps)</i>	<i>frame size 20 ms</i>	<i>(kbps)</i>
1-conn	71.4	217	33.5	467
2-conn	35.8	437	20.5	754
3-conn.	23.3	663	17.7	874
4-conn.	17.4	889	14.6	1059

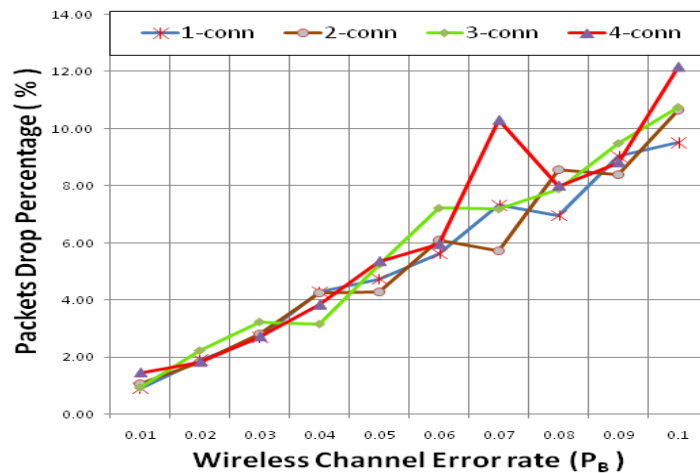


Fig. 2. Average packet drop rate for an increasing number of connections, according to channel error rate.

employed, which will give more favorable results than the longer frame size. As will be observed, no strong effects result from increasing the number of connections. Moreover, for all but the highest error rates the packet loss rate is below 10%.

An interesting comparison is with the throughput when the core network is included, Table 4. There is a similar pattern to the throughputs in Table 3 but the rates are reduced to when streaming only over the WiMAX link. We interpret this effect as not due to TFRC's response to packet loss but due to its response to the increased round trip time caused by queuing delay in the buffer prior to the bottleneck link in Fig. 1. This is confirmed by the increase in per slice/packet end-to-end delay as more connections are added, Table 3. In effect, the packets from other connections intervene in the router buffers causing an increase in latency.

More significantly for reconstruction of the video stream is the GOP ordering, which for four connections is shown in Fig. 3. Notice that the first GOP contains parameters that are fixed throughout the sequence, a feature of the H.264/AVC codec. Therefore, this GOP is transported more quickly. To avoid a sudden injection of traffic into the network, connection starting times were offset by 0.5 s.

A noticeable feature of this Figure is the lengthier start-up periods in sending initial GOPs on each of the connections. This does mean that about 6 s of frames (amounting to 90 frames) should be stored in the reordering buffer, to avoid the possibility of subsequent underflow in the decoder's playout buffer. As the destination is on the fixed network the reorder buffer is not expected to be a drain on energy resources, as it might be on an SS. 6 s is longer than an ideal start-up time of around 2 s but not too large to be objectionable to the user. Interestingly, when comparing with the throughput reported in [1] and repeated in [7], for MULTTFRC there are periods of at least three seconds when the throughput is approximately over half the peak rate. The aggregate throughput also may oscillate. In fact, on finding this

problem, we compared with [1] and discovered that allowance was made for 10 s start-up buffering before beginning decode, also to avoid buffer underflow. However, that work [1] used data from MPEG-4 at a lower 10 frame/s to test buffer occupancy.

The cause of the initial lengthier start-up periods may be a combination of factors. However, the early response of TFRC appears to be implicated. The initial rate of TFRC is set to one packet/s and no default settings for round-trip time or packet loss rate are used in the throughput equation (1). Normally, if no acknowledgement arrives within two round trip times then TFRC reduces its sending rate by half and goes into a slow start, similarly to TCP. However, the initial default value of the no-feedback timer is set to 2 s, which implies that TFRC's initial rate may be prolonged if acknowledgments are lost or delayed. If acknowledgement drops or delays still occur then it is possible that the rate will be halved again before slow-start. However, the timeout interval will be shorter as it is now given by:

$$timeout = \max(4r, (2s)/TFRC) \quad (2)$$

where r is the estimated round-trip time, as before s is the packet size, and $TFRC$ is the sending rate given by (1). Though in [7] it was acknowledged that drastic reductions in sending rate could occur due to the onset of slow-start, this was attributed to heavy packet loss and not to the loss or delay of acknowledgments, without data packet loss necessarily occurring.

Table 4. Sending periods and throughputs from the video streaming SS to the core network destination (node C in Fig. 1).

<i>No. of connections</i>	<i>MS to node c frame size 20 ms</i>	<i>Throughput (kbps)</i>
1-conn	35.2	444
2-conn	22.4	690
3-conn.	21.6	716
4-conn.	15.6	991

Table 5. Mean per slice/packet end-to-end delay

<i>No. of connections</i>	<i>Mean end-to-end delay (s)</i>
1-conn.	0.035
2-conn.	0.036
3-conn.	0.039
4-conn.	0.062




Fig. 3. Example arrival sequence at the receiver (node C in Fig. 1) showing the start and end times. GOP 1 contains the parameter-set for the sequence. Connection start times are staggered by 0.5 s.

Corresponding to Fig. 3, Fig. 4 plots individual throughputs and the aggregate throughput. As might be expected from Fig. 3, throughput gradually climbs until a plateau is reached. There is evidently some unfairness between the TFRC flows as connection 3 needs to prolong its delivery because of lower throughput at an earlier stage. However, there are less oscillations in rate than reported for MULTTFRC [1, 7], which is explained by the static scheduling scheme employed by us.




Fig. 4. Example run showing throughput over time for individual connections and the aggregate throughput.

Packet loss over time displays an oscillatory pattern for the example in Fig. 5, which is why it is unwise to rely on mean loss statistics alone. Based on the packet

loss patterns the average PSNR was found when increasing the number of connections, as recorded in Table 6. The frame sizes are adjusted in Table 6 to account for the buffer underflow that would occur were the shorter frame size to be used throughout the network path. However, counter-intuitively, employing a shorter frame size for a few connections over the WiMAX link alone results in lower video quality than when sending over the complete path. This is best explained by buffer overflow at the SS, caused by the short WiMAX frame size, rather than packet losses on the WiMAX wireless channel. When the number of connections increases, TFRC is better able to regulate its rate and the video quality increases over the single wireless link. Notice, however, that using smaller frame size even when the video quality is high can lead to excessive delay at the SS buffers if traffic is heavy.

4 Conclusion

This paper has conducted a relatively realistic investigation of multiple TFRC connections for uplink video streaming over a concatenated network, consisting of a WiMAX access network and a fixed network with a bottleneck at the network edge. The study has shown that with static scheduling of the video stream over the connections, increased throughput results. Reducing the video send time reduces the risk from wireless channel error. However, it also implies that reordering at the receiver is required. The resulting start-up delay was about 6s for the reasonably complex test video. With a moderate number of connections (four were used) video quality improved by over one dB for streaming across the modeled network. Using a smaller WiMAX frame size can lead to further improvements across the wireless link itself but there is a risk of excessive queuing at the subscriber station devices causing unacceptable delays. The role of the feedback channel is important, as loss or delay of acknowledgment packets seems to be implicated in the TFRC congestion controllers' slow start-up, one of the potential causes of buffer underflow. Further investigation will consider the role of acknowledgments and whether a reduction in the acknowledgment rate may improve performance further. It may also be possible to 'warm-up' the TFRC connection handlers by sending non-video data to start with, which could be discarded thereafter. This can reduce the size of the reordering buffer, if such a reduction were required.

Table 6. Video quality (PSNR) according to number of connections.

<i>No. of connections</i>	<i>PSNR (dB) recorded at WiMAX base station (frame size 5 ms)</i>	<i>PSNR (dB) recorded at node C in Fig. 1 (frame size 20 ms)</i>
1-conn.	26.72	31.84
2-conn.	31.32	32.34
3-conn.	35.92	33.15
4-conn.	35.32	33.34




Fig. 5. Aggregate packet loss numbers for connections for a sample run over time.

References

1. Chen, M., Zakhor, A.: Rate Control for Streaming Video over Wireless. *IEEE Wireless Comms.* 12(4), 32–14, (Aug. 2005)
2. Handley, M., Padhye, J., Floyd, S., Widmer, J.: TCP-Friendly Rate Control (TFRC): Protocol Specification. RFC 3448 (2003)
3. IEEE, 802.16e-2005. IEEE Standard for Local and Metropolitan Area Networks. Part 16: Air Interface for Fixed and Mobile Broadband Wireless Access Systems (2005)
4. Widmer, J., Denda, R., Mauve, M.: A Survey on TCP-friendly Congestion Control. *IEEE Network.* 15(3), 28–37 (May/June 2001)
5. Meloni, L. G. P.: A New WiMAX Profile for DTV Return Channel and Wireless Access. In: Chen, K.-C., de Marca, J. R. B. (eds.) *Mobile WiMAX*, pp. 291–392. Wiley & Sons, Chichester, UK (2008)
6. Chen, M., Zakhor, A.: Rate Control for Streaming Video over Wireless. In: *IEEE INFOCOM*, pp. 1181–1190 (2004)
7. Chen, M., Zakhor, A.: Multiple TFRC Connection Based Rate Control for Wireless Networks. *IEEE Trans. Multimedia.* 8(5), 1045–1062 (Oct. 2006)
8. Sadka, A.: *Compressed Video Communications*. Wiley & Sons, Chichester, UK (2006)
9. Balkrishnan, H., Padmanabhan, V., Seshan, S., Katz, R.: A Comparison of Mechanisms for Improving TCP Performance over Wireless Links. *IEEE/ACM Trans. on Networking.* 5(6), 756–769 (2007)
10. Tsai, F. C.-D. et al.: The Design and Implementation of WiMAX Module for NS-2 Simulator. In: *Workshop on NS2: The IP Network Simulator*. article no. 5 (2006)
11. Wiegand, T., Sullivan, G. J., Bjontegaard, G., Luthra, A.: Overview of the H.264/AVC Video Coding Standard. *IEEE Trans. Circuits Syst. Video Technol.* 13(7), 560–576 (July 2003)
12. Klaue, J., Rathke, B., Wolisz, A.: EvalVid - A Framework for Video Transmission and Quality Evaluation. In: *Int. Conf. on Modeling Techniques and Tools for Computer Performance*, pp. 255–272 (2003)
13. Haßlinger, G., Hohlfeld, O.: The Gilbert-Elliott model for packet loss in real time services on the Internet. *14th GI/ITG Conf. on Measurement, Modelling, and Evaluation of Computer and Commun. Sys.*, pp. 269–283 (2008)