Multimedia Congestion Control for Broadband Wireless Networks

Arjuna Sathiaseelan, Gorry Fairhurst, Electronics Research Group, University of Aberdeen
{arjuna,gorry}@erg.abdn.ac.uk

Abstract Multimedia applications have seen an exponential growth within the Internet. Most applications currently use UDP as their underlying transport protocol and transmit at the maximum rate allowed by the application, irrespective of other traffic. This raises fairness issues when used in the general Internet, and the possibility of starvation of other packet flows. The applications therefore need to evolve to use a congestion control protocol. A new protocol called the DCCP was therefore proposed by the IETF to allow efficient delivery of multimedia applications while also being able to detect and respond to network congestion. This paper investigates this topic relating this to multimedia transmission over broadband wireless (including satellite) networks, especially when such networks are integrated in the next generation Internet.

***Index Terms*—Internet Multimedia, TFRC, DCCP, Wireless.**

I. INTRODUCTION

Recent years have seen an exponential growth in multimedia applications within the Internet as a whole, with a general move towards support for triple-play services that combine voice, video (e.g. TV, teleconference) and data services over a common network. Efficient multimedia delivery requires broadband networks that are not currently available in all regions within the European Union (EU) due to geography (e.g. mountains), small rural communities, or peripheral regions, creating a “digital divide”, and will necessitate the use of broadband wireless technology to integrate these users within the global internet.

The most widely used transport protocol is the Transmission Control Protocol, TCP. However, TCP introduces an arbitrary delay because of its reliability and in-order delivery requirements, making it unsuitable for real-time multimedia. Most currently deployed multimedia applications therefore use the Real Time Protocol (RTP) with UDP [3] as the underlying transport, in preference to TCP. However, current UDP-based applications can (and often do) transmit at a constant rate, irrespective of the available capacity [9]. Although this suites the existing design of multimedia codec, the growth of such long-lived non-congestion-controlled traffic posed a real threat to the overall health of the Internet [6], and in particular forms an obstacle to the deployment of triple-play services over bandwidth-limited technologies such as broadband wireless.

An end to end Internet path often traverses a range of networks each having their own characteristics. Hence integrating a broadband wireless network with the Internet requires that the networking protocol used to carry multimedia services works in harmony with existing protocols (eg. TCP). A congestion control protocol is therefore required [10]. Furthermore a standards-based approach, such as the Datagram Congestion Control Protocol (DCCP) is needed to ensure widely deployed methods that have a good confidence of interoperability.

Broadband wireless networks span a large range of properties from short-range high-speed personal networks, to long-reach wide-area networks. This paper focuses on the latter. It does not constrain itself to a single wireless network technology, indeed the authors suggest that future networks will need to operate in environments combining several wired and wireless technologies to offer a seam-less end to end delivery of triple-play services. However, as a concrete example of wide-area wireless technology we have considered a satellite network model built on work in the SATSIX project utilising a two-way satellite network using the DVB-RCS standard [1].

One of the challenges faced by multimedia is the limited capacity offered by many wide-area wireless technologies. In general these networks offer less capacity than available in wired networks, and typically utilize sophisticated Radio Resource Management (RRM) methods to dynamically share the available radio bandwidth between a population of users. This can result in networks with limited (often variable) capacity, and appreciable delay (waiting for radio capacity to be requested by the RRM subsystem and then allocated to a specific terminal). This is known to impact TCP [2], we explore if this also impacts multimedia flows that use DCCP.

The remainder of this paper investigates congestion control for multimedia flows to allow closer integration of broadband wireless networks within the multimedia-enabled Internet. Section II discusses link QoS issues and the need for congestion control; Section III discusses congestion control protocols; Section IV provides simulation results; and

This work was supported by the EU Information Society Technologies SATSIX project, IST-2-26950. Gorry Fairhurst is funded by the University of Aberdeen, UK, where he leads the Communications and Imaging Research Group

Section V provides a discussion and section VI provides a conclusion.

II. QoS and Internet Multimedia

Performance of the end-to-end Internet service may be controlled by invoking network-layer QoS [14] using either the Integrated Services (IntServ) or Differentiated Services (DiffServ) model. The IntServ approach requires support throughout the network and is suited to “walled-garden” approaches used for corporate applications and to deliver the first TVoIP offerings. This approach has however not gained widespread use within the general Internet, due to the need to establish policing policies between networks and the need to deploy complex control protocols in routers. The Diffserv approach assigns flows to a set of traffic aggregation classes, and associates each class with appropriate QoS attributes. Traffic flows share network capacity, and congestion control techniques are required above the IP-layer to fairly distribute the available capacity between the flows.

Even when QoS is not used end-to-end, link-layer QoS can be highly desirable for wireless networks to ensure that flows belonging to different service classes have guaranteed delivery depending on the Service Level Agreements (SLA). QoS is sufficient to allocate the required bandwidth for every flow depending on their class and police them accordingly.

Link-layer QoS methods in the wireless network can enable radio capacity to be reserved for individual packet flows, or at a least can permit cross-layer RRM so that radio capacity can be assigned on the basis of expected traffic characteristics and requirements, and traffic may be appropriately queued at the transmit interface (e.g. using Class-Based Queues, CBQ).

In the absence of QoS signaling, a proxy may be used at the edge of a wireless network to snoop the control messages (e.g. SDP signaling for a VoIP call), and then to infer a link-layer QoS based on the encoding rate chosen by the application. Although, the advent of new audio codecs that support silence suppression, can make traffic bursty, which means that if radio capacity were allocated based on expected peak rate, the allocated capacity would be not fully utilized. Moreover, when multimedia flows using UDP share the same QoS class with data services using TCP, they may degrade the performance of the data services during congestion.

III. Congestion Control

A. TCP Friendly Rate Control (TFRC)

TFRC [13] was designed to allow multimedia applications to provide congestion control when using UDP with RTP [4]. In TFRC, the receiver periodically sends a feedback report informing the recent packet loss event rate for the session (this requires bi-directional communication). Using this the sender calculates the allowed sending rate for the next Round Trip Time period by using an equation that models the equivalent throughput that would have been obtained by a TCP flow. TFRC has been suggested for multimedia applications such as TVoIP and VoIP, since it allows the sending rate to vary more smoothly by decreasing and increasing the sending rate gradually. It can be used for both unicast and multicast traffic.

TFRC was designed to be reasonably fair when competing for capacity with TCP connections that use a fixed packet size. This has consequences when used for voice services that utilize small IP packets: a low rate TFRC flow sharing a network capacity (e.g. the satellite link) with high-bandwidth TCP flows (large-packets) may need to slow down, even though the nominal rate of the TFRC flow is less than the rate achieved by the TCP flows.

To enable TFRC to be used for voice traffic, TFRC-SP [15], a Small-Packet (SP) variant of TFRC was designed. TFRC-SP attempts to use the same capacity as a TCP flow using packets of up to 1500 bytes. It also enforces a minimum interval of 10 ms between packets, to prevent a single flow from sending packets too fast. TFRC and TFRC-SP are sufficient to provide congestion control to UDP multimedia; however the application designer needs to implement these protocols correctly. They also seek to provide only congestion control, and not other important services such as feature-negotiation, firewall traversal and mobility support.

B. Datagram Congestion Control Protocol (DCCP)

To provide a standard way to introduce congestion control into multimedia applications, the Internet Engineering Task Force (IETF) defined the Datagram Congestion Control Protocol [6] in 2006, as an alternative to UDP for unicast multimedia applications that prefer timeliness of data to reliability. This new transport protocol is designed for deployment as a standard feature in end hosts (PCs, VoIP codecs, and other internet-enabled multimedia appliances).

Media/Session Control	Media Codecs
SIP | H.323 | RTSP	RTP+RTCP
Transport Layer Interface	
CCID 2 | CCID 3 | CCID 4 | …	
DCCP	
IPv4	IPv6

Table 1: DCCP Multimedia Stack

DCCP provides a comprehensive set of features:

- Although many DCCP sessions are expected to transfer data in a uni-directional flow, a DCCP connection also contains a flow of acknowledgement traffic in the return direction. This informs a sender whether the transmitted packets have arrived, and the path RTT. This allows DCCP to provide a standard way to implement congestion control. DCCP recognizes that not all multimedia applications require the same service, it therefore provides a choice of congestion control mechanisms for each half-connection, according to three profiles:
    - CCID 2 [11] TCP SACK-like Congestion Control, appropriate for flows that like to receive as much as

possible over the long term (e.g. download of streaming content).
  - o CCID 3 [12] TFRC Congestion Control, appropriate for flows that prefer to minimize abrupt changes in the sending rate, including conversational streaming.
  - o CCID 4 [16] TFRC Congestion Control for Small Packets, appropriate for flows that prefer to minimize abrupt changes in the sending rate, but that generate small packets at a low rate.
- DCCP provides an interface to other standards-based network techniques such as Explicit Congestion Notification (ECN), and Path MTU Discovery.
- DCCP implements reliable connection setup and teardown enabling firewall traversal.

DCCP also provides facilitates to negotiate connection properties between end hosts and copious space for options, allowing new features to be added as the needs of applications becomes better understood.

IV. SIMULATION

This section presents results to assess the need for congestion control, and how DCCP may be used to provide congestion in an environment where path characteristics and path RTT vary. The network simulator ns-2 was used [8]. The experiments simulated a scenario that includes a broadband wireless link connected to the Internet, which forms a congestion bottleneck with variable capacity (up to 2 Mbps) and delay (a base propagation delay of 300ms delay). This incorporates an RRM model (based on that in [2]) and one QoS class for all traffic.

A. Fairness

We first show how DCCP CCID 3 could improve fairness when capacity is shared between multiple UDP streams and a single (long-lived or bulk) FTP TCP flow. We modeled a multimedia flow using a constant bit rate (CBR) traffic.

Each UDP flow sent 50 packets per second, each packet with a fixed 1000 B of data (400 kbps – e.g. a low-rate video codec). The single TCP flow carried packets of size 1460 bytes with a maximum congestion window (cwnd) of 40 packets per RTT (60 KB). Initial tests used one single UDP flow and one single TCP flow. Tests were then repeated by gradually increasing the number of UDP flows, up to a total of 50 flows.

Figure 1shows that as the number of UDP flows increased, the throughput of the TCP flow is gradually reduced. When there are multiple UDP flows, the single TCP was deprived of its fair share of the capacity and was unable to increase its cwnd appropriately. This demonstrates that multimedia applications that use UDP do not therefore fairly cooperate with TCP flows (unless using a congestion control mechanism, such as TFRC).

The tests were repeated by replacing the UDP flows by flows using DCCP CCID 3. **Figure 1** shows that the TCP flow was able to get more than its fair share of the capacity. Multimedia flows using CCID 3 decreased more gradually than TCP, but still allowed the TCP flow to achieve a fair share of the available capacity. We also note that the feedback generated by CCID 3 is once per RTT, rather than once per packet (or pair of packets) as in TCP.

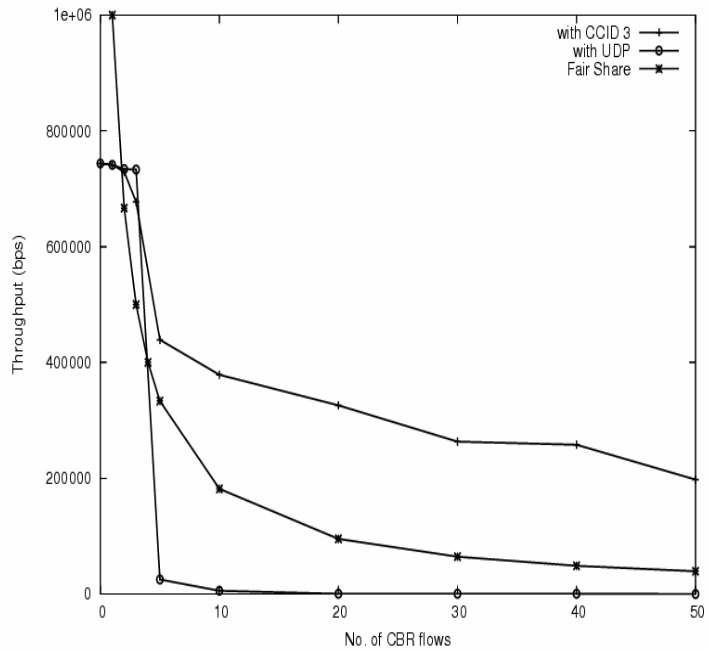

Figure 1: Throughput of a single TCP flow (vs) Varying no of CBR flows

These results demonstrate that DCCP can enable multimedia flows to coexist with TCP traffic within the same QoS class without causing unfair sharing of network capacity. This fairness is important when deploying a mixture of voice, video and data services, and particularly important for links with limited and varying capacity, such as those offered by broadband wireless. As such, the deployment of DCCP can act as an enabling technology to the growth of widespread support for triple-play services.

B. VoIP over DCCP

To assess the quality of a VoIP call using DCCP CCID 4 for a range of delays, simulations were performed by modeling the G.711 (64kbps) and G.729 (8kbps) codecs, both with and without Voice Activity Detection (VAD). When each call is setup, it starts sending media packets when the sending rate reaches the desired encoding rate of the codec. The model from [18] was used to model the effect of VAD with conversational data. This model is consistent with the 3GPP assessment of VoIP codecs [20].

A range of average burst and silence parameters were used based on exponential distributions of 1.0s and 1.5s. The transmit buffer was set to 5 packets. To assess quality we use the E-Model defined in [17] to determine a Quality Factor “R” for each voice call. For VoIP traffic, the R score is calculated based on subfactors, that affect the impact of loss and delay.

This R Score can then be translated to the Mean Opinion Score (MOS) using
Table 2.

R Score	MOS	Perceived Quality
90-100	4.34-4.5	Best
80-90	4.03-4.34	High
70-80	3.6-4.03	Medium
60-70	3.1-3.6	Low
50-60	2.58-3.1	Poor

Table 2: R Score, MOS and Perceived Quality

Tests were performed with background TCP web traffic for various network delays. Figure 2 plots the performance of the codecs with and without VAD. Even with increased delays, medium perceived quality could be achieved without VAD. This shows that using DCCP as the underlying transport protocol does not hamper the performance when the call is sent at constant rate.

When VAD is employed, the media is encoded to form bursty traffic. The conservative nature of DCCP CCID 3, can delay the ability to utilize capacity at the start of a talk spurt, with a resulting impact on the performance of the application and leading to a poor perceived quality.

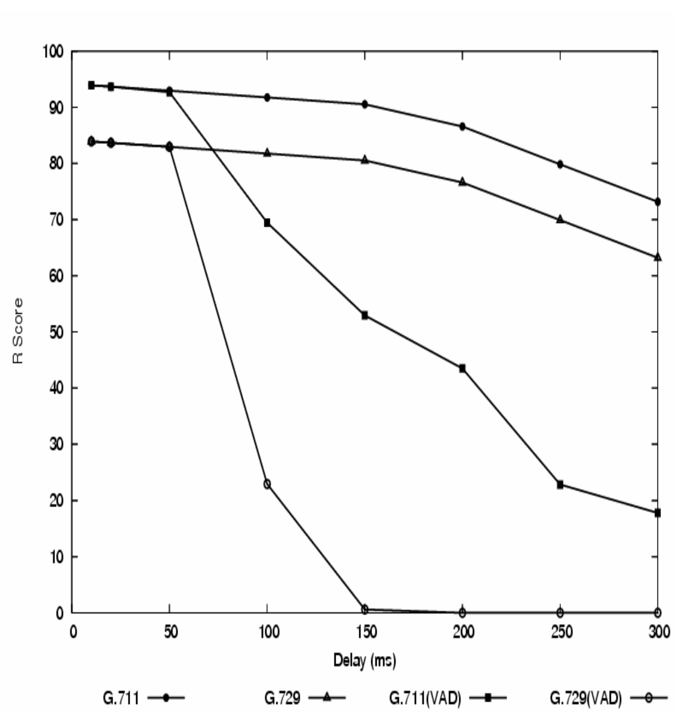

Figure 2: R Scores v Network delay

C. Performance Interactions with RRM

Using DCCP without any QoS in networks that combine high delay and dynamic assignment of capacity (RRM) can hamper the performance of delay-sensitive flows such as VoIP. To asses the need for link-QoS along with DCCP, we analyzed the performance of a VoIP call (G.711 codec, 50 packets per second) over a link with a propagation delay of 300 ms and one with a propagation delay of 600 ms (Figure 3). We also compared a simple queuing method using no priority and one using Class-based Priority Queuing (CBQ). We analyzed both Transparent Satellite (TS) and Regenerate Satellite (RS) scenarios.

Without link QoS, there is no prioritization of traffic, and the available capacity is mostly used by incoming TCP web traffic. Furthermore, the RRM subsystem can not predict traffic characteristics when requesting capacity assignments. The increased and varying end to end delay for the voice packets impairs the voice quality (as also observed by other researchers [18]). The voice quality decreases with increasing background web traffic.

In contrast, a QoS-aware link can classify and prioritize flows [14]. Priority queuing allows voice traffic to take priority allowing the RRM to transmit these packets with a higher precedence using the received capacity allocations. An increase in total traffic (web or VoIP) arriving at a single transmitter, leads to a higher rate of request for capacity to the RRM, increasing the number of available timeslots and this reducing the delay seen by the higher priority VoIP traffic.

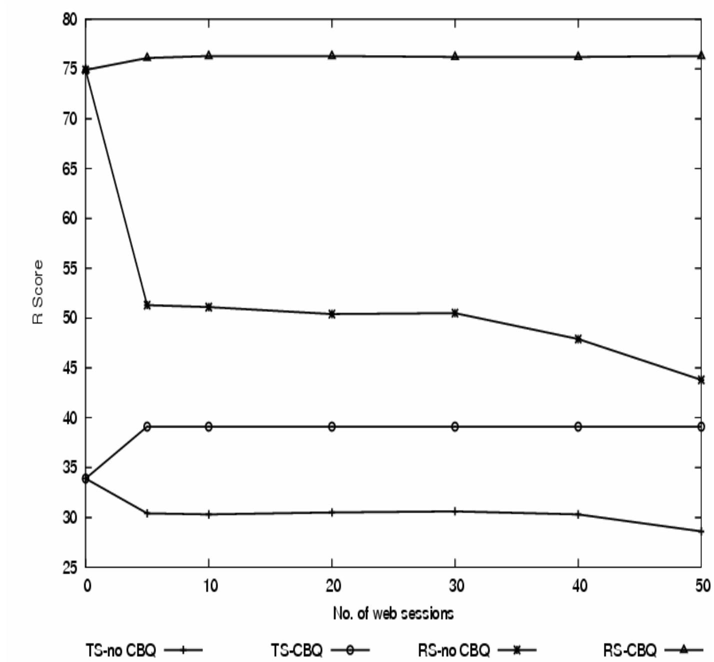

Figure 3: R Scores for VoIP Traffic

The previous figure illustrates the effect of a 600 ms propagation (TS) delay, making the network poorly suited for interactive traffic such as VoIP. This is evident from the low R Scores when compared to the case for 300 ms case (RS). The QoS-aware approach is advantageous and enables multimedia applications to use DCCP, enabling them to maintain their quality as well as cooperating fairly with other flows.

V. DISCUSSION

Next generation of networks are expected to be both multi-service, but also to incorporate a wide range of wireless and wired technology. This complicates development of real-time multimedia applications [9]. Despite widespread deployment of VoIP and IPTV systems in wired IP networks, there are still no effective solutions for multimedia congestion control or deployable inter-domain QoS [10].

This paper has investigated how DCCP could be used to offer multimedia congestion control. Initial simulations were presented in section IV to demonstrate that DCCP can offer acceptable solutions. The authors note several important limitations in their current simulations that need to be addressed prior to detailed analysis:

- We only analyze the case for simple media codecs. Real-world media codecs are known to result in much more complex traffic patterns (the investigation of VAD in section IV highlights the issues this may raise when used in wireless networks that employ sophisticated RRM).
- We do not take into consideration the adaptive waveforms used in many wireless technologies. These can operate over periods similar to those of congestion control protocol, possibly introducing interactions.
- Finally, the results assume a simple interface between the transport protocol and the media codec, while this may be typical of current fixed-rate codecs, adaptive multimedia codecs can benefit from much more sophisticated interfaces, changing the dynamics of the way in which the transport protocol is used.

This simulation work proceeds in parallel with implementation of DCCP in the Linux Kernel. Future work at the University of Aberdeen will explore DCCP in an Internet environment that includes both fixed and wireless networks and study the implications of a variety of multimedia codecs. This is seen as an important step towards a framework to support real-time multimedia applications that can be used throughout the range of networks that comprise the Internet.

VI. CONCLUSION

This paper demonstrates that congestion-controlled multimedia streams may be delivered over broadband wireless networks that utilize sophisticated radio resource management protocols. It presents a brief overview of the need for congestion controlled multimedia and in particular describes the use of the DCCP transport protocol. The paper provides some initial simulation results for two DCCP Congestion Control IDs and shows that when matched to suitable link-layer QoS methods acceptable performance may be achieved in terms of fairness with other traffic and user perception of the multimedia content.

REFERENCES

[1] EN 301 790: *DVB Interaction Channel for Satellite Distribution Systems*, ETSI, 2003.

[2] M. Sooriyabandara, G. Fairhurst *Dynamics of TCP over BoD Satellite Networks,* Int. J. of Satellite Communications and Networking, 21, p427-4449, 2003.

[3] Postel J., User Datagram Protocol, RFC 768, 1980.

[4] Schulzrinne H., et al. *RTP: A Transport Protocol for Real Time Applications*, RFC 1889, January 1996.

[5] Floyd S., Handley M., Kohler E., *Problem Statement for the Datagram Congestion Control Protocol (DCCP),* RFC 4336, March 2006.

[6] Kohler E., Handley M., Floyd S., *Datagram Congestion Control Protocol (DCCP),* RFC 4340, March 2006.

[7] Jacobson V., *Congestion Avoidance and Control*, ACM SIGCOMM, pp. 273-288, 1988.

[8] McCanne S., Floyd S., *Network Simulator*. http://www.isi.edu/nsnam/ns/

[9] D.Wu, Y. T. Hou, Z-Q Zhang, *Transporting Real-Time Video over the Internet: Challenges and Approaches* Proc. IEEE, 33(12), 2000.

[10] S. Floyd, K. Fall *Promoting the Use of End-to-End Congestion Control in the Internet*, IEEE Trans on Networking, 7(4), p 458-472 1999.

[11] Floyd S., Kohler E., *Profile for DCCP Congestion Control ID 2: TCP-like Congestion Control,* RFC 4341, March 2006.

[12] Floyd S., Kohler E., Padhye J., *Profile for DCCP Congestion Control ID 3: TCP-Friendly Rate Control (TFRC),* RFC 4342, March 2006.

[13] Handley M., et al., *TCP Friendly Rate Control (TFRC) Specification*, RFC 3448, January 2003.

[14] Mort R. J., Combes S., Cruikshank H., *QoS architecture for broadband satellite multimedia (BSM) networks*, IST Mobile Summit 2006.

[15] Floyd S., Kohler E., *TFRC: the Small-Packet (SP) Variant*, Work in progress, 2006.

[16] Floyd S., Kohler E., *Profile for DCCP Congestion ID 4: TCP-Friendly Rate Control for Small Packets (TFRC-SP)*, IETF Work in Progress, 2006.

[17] G.107, *The E-model, a computational model for use in transmission planning*, ITU-T, 2005.

[18] H. Skinemoen, et al., *VoIP over DVB-RCS with QoS and Bandwidth on Demand*, IEEE Wireless Communications, October, p46-53, 2005

[19] K. Sriram, W. Whitt, *Characterizing superposition arrival processes and the performance of multiplexers for voice and data*, IEEE Globecom, December 1985.

[20] TR 26.976 *Performance characterization of the. Adaptive Multi-Rate Wideband (AMR-WB) speech codec (Release 6*), 3GPP, ETSI, 2004.

[21] E. Kohler, M. Handley, and S. Floyd *Designing DCCP: Congestion Control Without Reliability*, ACM SIGCOMM, September 2006.