

Performance of VoIP using DCCP over a DVB-RCS Satellite Network

Arjuna Sathiaseelan Gorry Fairhurst
Department of Engineering,
Electronics Research Group,
University of Aberdeen, United Kingdom
Email: {arjuna,gorry}@erg.abdn.ac.uk

Abstract—The Datagram Congestion Control Protocol (DCCP) is a new IETF-defined transport protocol for Internet multimedia. The DCCP Congestion Control Identifier 4 (CCID4) was specifically designed to support the stream of small packets generated by applications such as Voice over IP (VoIP). This uses equation based congestion control. The long round trip delay when the path includes a satellite link could therefore degrade the performance of VoIP. In this paper, we analyze using simulation the performance of VoIP using CCID4 over a DVB-RCS network. We also analyze the different variants of CCID4 that have been proposed. The paper concludes that the current algorithm of CCID4 suffers for paths with appreciable delay when used to support voice calls with a bursty nature. The results show that these effects are only partially mitigated by the currently proposed variants of CCID4.

I. INTRODUCTION

Multimedia (streaming audio and video, Voice over IP (VoIP) and interactive video) is seen as a key service in most next generation networks. Delivery of multimedia services over the satellite network demands that the satellite network should be easily integrated with the fixed and mobile Internet. This requires the networking protocols used to carry the multimedia services to cooperate in harmony with the existing protocols that are used for data services.

UDP [15] is a connectionless protocol that offers non-guaranteed datagram delivery between end hosts. UDP gives applications direct access to the datagram service of the IP layer. An application running over UDP must provide their own mechanisms to deal with retransmission for reliable delivery, packetization and reassembly. Current multimedia applications use UDP [5] as the underlying transport protocol. UDP is usually chosen in preference to TCP [11], because:

- Startup Delay: the three-way handshake before initiating data transfer induces a delay. UDP avoids this delay.
- Statelessness: TCP holds connection state. This increases the risks of potential state holding attacks. UDP avoids holding connection state.
- Trading Reliability against Timing: multimedia data is timely, if it is not delivered by some deadline (typically a small number of round trip times (RTTs)), the data will not be useful at the receiver. TCP can introduce an arbitrary delay because of its reliability and in-order delivery requirements, making it unsuitable for real-time media.

Current UDP-based applications can (and often do) transmit at a constant rate, irrespective of the available capacity. Growth

of such long-lived non-congestion-controlled traffic posed a real threat to the overall health of the Internet [5], and in particular form an obstacle to the deployment of triple-play services over bandwidth-limited technologies such as mobile Internet, wide-area wireless and satellite systems. Concerns about the impact of congestion that could result from multimedia services have lead to investigation of congestion-responsive transport protocols. This is essential for building next generation multimedia applications to ensure continued operation of the Internet with a wide variety of multimedia/data applications. These issues have been recognized by the IETF and are addressed by a new standard, the Datagram Congestion Control Protocol (DCCP).

In this paper, we analyze using simulation the performance of VoIP using DCCP Congestion Control Identifier 4 (CCID4) over a DVB-RCS network. We also analyze the different variants of CCID4 that assess the performance of VoIP. The paper concludes that the current algorithm for CCID4 is currently inefficient to support voice calls that are highly bursty, and the existing variants of CCID4 do improve the performance slightly, but are still not sufficient. The rest of the paper is organized as follows. Section II presents a brief overview of the DCCP protocol. Issues pertaining to satellite networks are provided in Section III. Section IV presents the variants of the DCCP CCID4 protocol. Simulation setup and results are provided in Section V followed by conclusions and acknowledgements.

II. DATAGRAM CONGESTION CONTROL PROTOCOL

DCCP [5] is a packet stream transport protocol. Applications that can use DCCP include those that prefer timeliness of data to reliability. These applications benefit from the flow-based semantics of TCP, but do not want TCP's in-order delivery and reliability semantics. Applications include streaming media and voice over IP. The primary motivation for the development of DCCP is to provide a way for applications to gain access to standard congestion control mechanisms without having to implement them at the application layer and to fairly cooperate with other transport protocols namely TCP. DCCP provides an alternative to UDP to efficiently deliver multimedia content over satellite and at the same time fairly cooperating with other transport protocols.

DCCP provides a standard way to implement congestion control and congestion control negotiation for multimedia

applications. It provides a choice of congestion control mechanisms for each half-connection. CCID 2 is based on the TCP SACK-like congestion control protocol [6], is appropriate for DCCP flows that would like to receive as much bandwidth as possible over the long term, consistent with the use of end-to-end congestion control, and that can tolerate the large sending rate variations characteristic of Additive Increase Multiplicative Decrease (AIMD) congestion control, including halving of the congestion window in response to a congestion event. CCID 3 (Congestion Control Identifier 3) is based on the TCP Friendly Rate Control (TFRC) protocol [7]. TFRC [9] is a rate-paced equation-based congestion control protocol that requires both the sender and receiver to participate in determining the allowed sending rate. TFRC exhibits a much lower variation of throughput over time compared to TCP making it more suitable 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, while ensuring that it competes fairly with TCP. However, this makes TFRC to respond slower to changes in available bandwidth compared to TCP. TFRC-SP [4], a Small-Packet (SP) variant of TFRC was designed for applications such as VoIP that send small packets. TFRC-SP seeks to achieve the same bandwidth in bps as a TCP flow using packets of up to 1500 bytes. TFRC-SP also enforces a minimum interval of 10 ms between packets, to prevent a single flow from sending small packets arbitrarily frequently. CCID 4 [8] is a suggested variant of CCID3 that is based on TFRC-SP.

III. SATELLITE BANDWIDTH ON DEMAND

Satellite bandwidth on demand (BoD) is an integral component in satellite resource management. It operates at the lower link layer (the Medium Access Control (MAC) sub-layer) on satellite access systems. The Digital Video Broadcast Return Channel System (DVB-RCS) standard [2] defines a framework for constructing a standard-based broadband satellite system. BoD acts between each Return Channel Satellite Terminal (RCST) and a particular hub/gateway terminal (often connected to the wider Internet). DVB-RCS currently specifies five different BoD capacity request mechanisms:

- Continuous Rate Assignment (CRA): a fixed number of slots are allocated and remain valid while the terminal is active (logged on to the network).
- Rate Based Dynamic Capacity (RBDC): Capacity is assigned after an explicit capacity request (CR) from an RCST. The rate can later be modified by another CR.
- Volume Based Dynamic Capacity (VBDC): a demand-based scheme that sends a CR for additional slots corresponding to the volume of packets in the transmit queue.
- Absolute VBDC (AVBDC): similar to VBDC, but each new CR specifies the current number of slots required to empty the RCST transmit queue.
- Free Capacity Assignment (FCA): this scheme assigns capacity to a specific RCST that has not been currently allocated to one or more RCSTs. This assignment is not explicitly requested by an RCST.

These mechanisms, used alone or in combination, form the basis of different bandwidth allocation schemes.

Previous research on VoIP over DVB-RCS [19], [21], [20], [23] show that transmission of VoIP traffic over a DVB-RCS system can give a very good to medium perceived quality depending on the allocation scheme. When the presence and characteristics of VoIP traffic is known, RBDC may be used to explicitly reserve time slots for this traffic. Priority scheduling is often also used, to ensure VoIP gains access to slots in preference to other traffic without incurring additional delays from the BoD system.

Most currently published results consider a simple transport protocol (UDP) that does not offer congestion control and many assume the voice traffic to have a constant rate (i.e. they did not consider the effects of voice activity detection (VAD)/silence suppression, which lead to the voice traffic being bursty). However, both congestion control (as in DCCP) and silence suppression can increase the efficiency of transmission, and can permit VoIP traffic to share a bandwidth-limited link with other Internet traffic (as required in triple-play scenarios). It is therefore likely that these methods will be deployed in the general Internet and therefore important to assess their performance over next generation satellite links.

IV. VARIANTS OF CCID 4

This section investigates the impact of the satellite delay on the performance of a number of DCCP CCIDs. DCCP can be used over a satellite link, but that the performance of some CCIDs are sensitive to path characteristics including the effects of path RTT, packet loss, and variation in the RTT (packet jitter). Some problems of carrying multimedia traffic using CCID 3 or CCID 4 on satellite networks are:

- Initial Slowstart: CCID 3 starts with a maximum of 4 packets per RTT. The initial rate (for the first few RTTs) is insufficient for multimedia applications that demand large encoding rates. For a small RTT, the start-up delay is acceptable, however an appreciable satellite delay can severely degrade the performance.
- Idle period: During an idle period, the sending rate may reduce to a minimum of 4 packets per RTT. Slowstarting to the encoding rate takes an excessive time when operating over satellite links.
- Sending rate limit: CCID 3's sending rate can be at most twice the current receiver rate. This growth rate is not sufficient to maintain the encoding rate when a VoIP application oscillates between silence and talk periods.

Two mechanisms have been proposed that may enhance the performance of CCID4, these are described below.

A. Faster Restart

The Faster Restart (FR) mechanism [10] was motivated by the assumption that after idle periods, slowstart may be safely accelerated (quadruple the sending rate every congestion-free RTT) based on the knowledge of past transmission rates. Faster Restart also specifies that the allowed sending rate is never reduced by an idle period below eight packets per RTT, for small packets.

Configuration Parameter	Value
Frame (duration)	69.632ms
Subframes	4
Superframe size	2
Timeslots per subframe (TRF)	6
MPEG2 packets per timeslot	4
Size of timeslot	752B
Total MPEG2 packets per subframe	24
Granularity	86kbps
Uplink Bandwidth per Carrier	2.073 Mbps
Downlink Bandwidth	33.18 Mbps

TABLE I
AMERHIS SIMULATION PARAMETERS

B. Quickstart

Quickstart (QS) [3] is a new method that was originally conceived to improve the performance of TCP bulk transfers over lightly-loaded network paths. QS may also be useful for multimedia flows. In this case it can alleviate the effect of slowstarting to the encoding rate, and after periods of silence [1]. QS is an IETF experimental protocol designed to provide lightweight signalling between routers and end hosts. Using QS the authors suggest that standard Internet protocols can effectively and efficiently work over a wide range of links - including those with satellite delay. The sender sends a QS request for its required sending rate. In the use described in this paper, the QS request uses the media encoding rate to select the requested sending rate. If a QS enabled router can support the requested rate, this approves the rate, otherwise it either reduces the rate or does not approve the request and removes the rate request header. On receipt of the QS request, the receiver sends a QS response (if the requested rate is approved) with the suggested sending rate back to the sender. Based on the approved rate, the sender appropriately increases the rate. QS maybe used when a sender faces the need for a sudden increase in transmission rate.

V. SIMULATION OF DCCP OVER SATELLITE

This section analyzes the performance of CCID4 and its variants for VoIP traffic when used over a DVB-RCS satellite network. The Amerhis regenerative DVB-RCS satellite system [14] was simulated by adapting the TDMA-BoD system [13] in ns-2 [12] to implement a bandwidth-on-demand allocation method appropriate to the Amerhis system. Priority link queuing was also added to ns-2, together with support for CCID4, Faster Restart (FR), and Quickstart (QS).

Table I presents the simulation parameters. Simulation tests were performed with both RBDC and CRA with and without priority queuing. Each terminal has a CRA allocation of one timeslot, ensuring that every terminal has at least one timeslot to send data. Simulations were started from 10.0s to ensure that each terminal has been initialized with a single slot.

When each call is setup, it starts sending media packets when the sending rate reaches the desired encoding rate of the codec. [18] introduces an ON/OFF model for conversational

Configuration Parameter	Value
Number of web nodes per terminal	Varying
Number of web sessions	10
Number of pages per session	10
Number of objects per page	10
Average size of object	10
Shape parameter for Pareto distribution of web size	1.05
Interpage parameter	1

TABLE II
BACKGROUND WEB TRAFFIC PARAMETERS

R Score	MOS	Perceived Quality
90 to 100	4.34 to 4.50	Best
80 to 90	4.03 to 4.34	High
70 to 80	3.60 to 4.03	Medium
60 to 70	3.10 to 3.60	Low
50 to 60	2.58 to 3.10	Poor

TABLE III
R SCORES, MOS AND PERCEIVED QUALITY

data having periods of voice activity of variable length with periods of silence. This model is consistent with the 3GPP assessment of VoIP codecs in [24]. In this paper, we have used various average burst and silence parameters using exponential distributions. The codec modeled was a G.711 codec sending 50 packets per second (sampling interval of 20ms), each packet having 160 bytes of data. The size of the transmit buffer was 5 packets.

Simulations were also performed with background web traffic (Table II). An average of twenty simulation runs were used for the analysis in this paper.

We use the E-Model defined in [17] to determine the Quality Factor "R" for each voice call. For VoIP traffic, the R score is calculated based on subfactors such as impact of loss (I_e) and impact of delay (I_d) using the formula $R = 94.2 - I_e - I_d$, where $I_e = \lambda_1 + \lambda_2 \ln(1 + \lambda_3 e)$ and $I_d = 0.024d + 0.11(d - 177.3)I(d - 177.3)$, $I(x)$ being a unity function. This factor "R" can then be translated to determine the Mean Opinion Score (MOS). Table III gives the correlation of the R Score, MOS and the perceived quality.

A. CCID4 and its variants with 10s silence period

In this section, we demonstrate the sending rate behavior of CCID4 and its variants during the startup period and after a brief idle period. There was no background traffic. The silence period was set to 10s to illustrate the effect of idleness (actual values vary and are a function of the codec and the way the media is used). Figure 1 shows that it takes up to 5s for CCID4 to reach the encoding rate of 8000 bytes per second (Bps). CCID4-FR also takes the same amount of time since CCID4-FR performs differently compared to CCID4 only after idle periods. CCID4-QS reduces the call setup phase (ramp-up period) by increasing the sending rate to the encoding rate

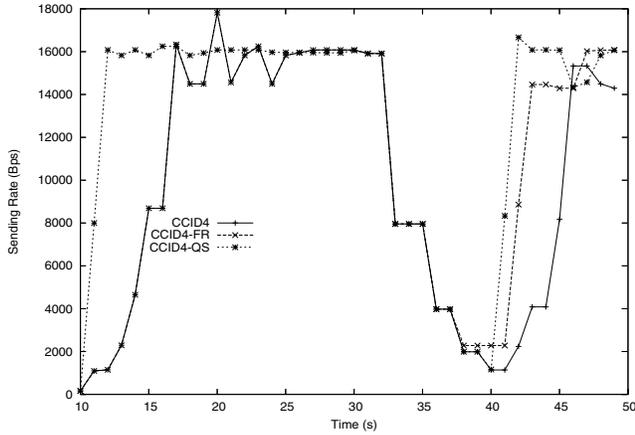


Fig. 1. Comparison of the sending rates of CCID4, CCID4 with FR and CCID4 with QS.

in 1s, compared to that of 5s for CCID4-FR. Similarly after the 10s idle period, CCID4 once again takes upto 5s to ramp up to the encoding rate, whereas CCID4-FR and CCID4-QS can ramp up within 2s and 1s respectively. CCID4's variants effectively ramp up the sending rate to the encoding rate much quicker compared to CCID4 after a brief idle period (more than 4 RTTs). CCID4-QS also enables the sender to ramp up during the initial slowstart. The sending rate can be upto twice the receiver rate, and this is reflected in the graph where the sending rate can grow upto 16000 Bps even though the sender effectively sends upto 8000 Bps of data.

B. CCID4 Performance: Applications with and without VAD

An encoder can reduce the capacity used by implementing Voice Activity Detection (VAD) with suppression. This changes the VoIP traffic from fixed rate (e.g. 50pps for G.711) to a bursty on/off source. Two cases are important: one with buffering, where the encoder buffers the packets generated during the ON period for a period of time and then transmits at a constant rate; the second with no buffering where the encoded packets during the ON periods are sent directly in a bursty fashion.

1) *Constant Rate Traffic:* To verify the perceived quality of a VoIP encoder sending at a constant rate using CCID4 over the Amerhis system, we simulated a Constant Bit Rate (CBR) source sending at 50pps (64kbps, VAD not set), and a CBR source sending at a rate of 25pps denoting a 50% speech activity (VAD with small amount of buffering). This buffering is sufficient to remove the effects of silence and maintain a constant rate of 32kbps at the sender. As the media rate is constant, CCID4 manages to maintain a steady state rate, thus no packets are lost due to transmit buffer overflow. Simulations were performed with RBDC (with and without priority queuing), and with CRA. Simulations were also performed with varying background web traffic. This was achieved by varying the number of web nodes connected to the terminal.

Figure 2 shows that when using RBDC (without priority) the perceived quality of a voice call decreases with an increasing

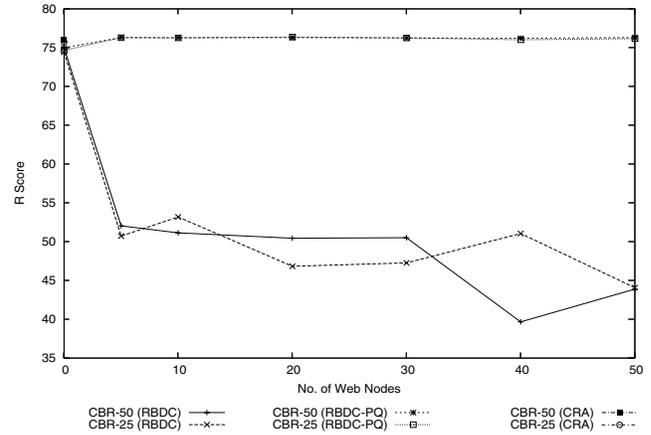


Fig. 2. R Score: Constant rate traffic with varying number of web nodes.

volume of background web traffic. With priority queuing (RBDC-PQ), the perceived quality of the voice call is steady and similar to that of quality of the call when CRA is used, irrespective of the encoding rate of the application. This is evident from the graph, where the performance of CBR sending 50pps and CBR sending 25pps is similar.

With priority queuing, the voice call has priority over the web traffic, and when the packets of the VoIP flow arrive at the MAC layer, these packets are given a higher precedence and sent using the slots already allocated in response to web traffic. Thus, with increasing web traffic, the packets of the VoIP flow have more slots available compared to no background web traffic. With priority queuing the R score is in between 70 to 80 thus having "medium" perceived quality. As the rate was steady, there was no loss of packets due to transmit buffer overflow. Hence the sub factor mainly affecting the R score is the increased delay of the satellite link.

Tests were also performed to determine jitter (both mean and std deviation) over the RCST network. Figure 3, shows that with a dynamic allocation scheme such as the RBDC along with priority queuing, the average jitter along with the standard deviation (std) is minimal (for CBR with 50 pps, the mean is between 14 and 16 ms, std is in the range 19 ms to 21 ms; for CBR with 25 pps, the mean is in the range 20 to 21 ms, std is in the range 35 ms to 50ms) compared to the large value when RBDC is used without priority queuing.

These results agree with the previous findings [19], that sending VoIP calls at a constant rate does give medium perceived quality. With priority queuing, the perceived quality of the voice call remains unaffected despite the underlying BoD scheme. Moreover, using DCCP as the underlying transport protocol, does not hamper the performance when the call is sent at constant rate.

2) *Bursty Traffic:* This section, evaluates the quality of voice calls that are bursty. When the encoder does not buffer the encoded packets during ON periods the media rate becomes bursty. [18] in their model suggest that the sender generates a voice call by interleaving bursts and silence with lengths generated by an exponential distribution having mean 0.352 s and 0.65 s respectively. Based on this model, we

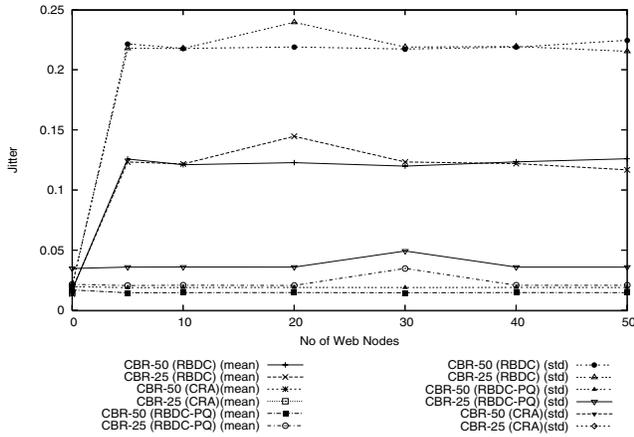


Fig. 3. Jitter: Constant rate traffic with varying number of web nodes.

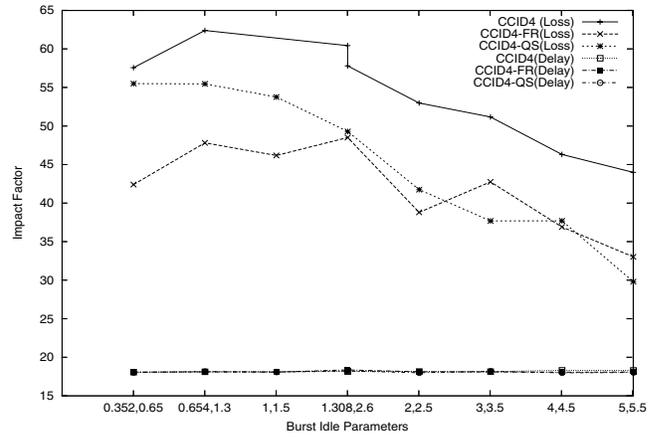


Fig. 5. Impact factor.

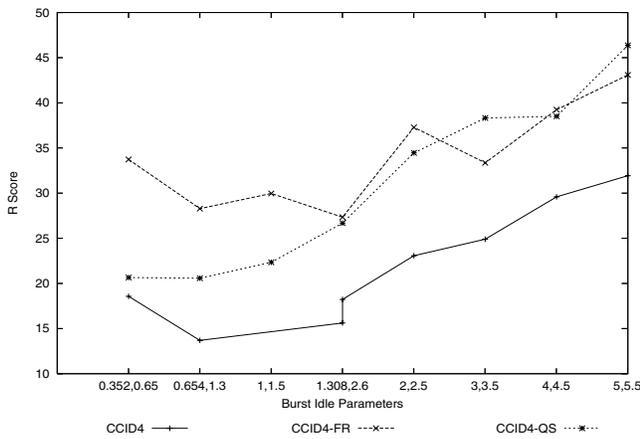


Fig. 4. R Score: Bursty traffic with different burst and silence parameters.

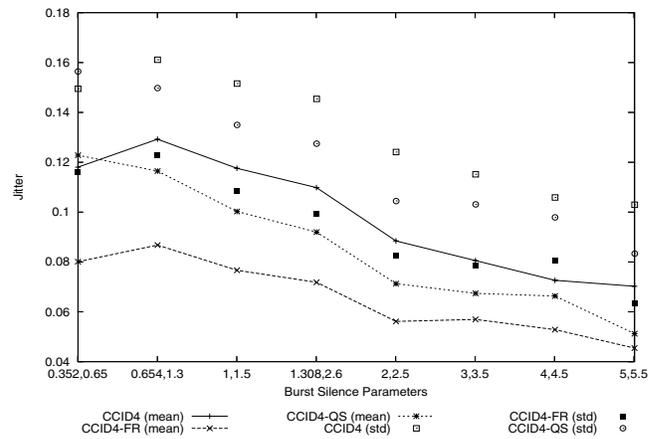


Fig. 6. Jitter: Bursty traffic with different burst and silence parameters.

used the parameters suggested in [18] and also increased the burst and idle periods to model varying speech activity. The simulation tests were also performed having five web nodes generating background web traffic.

Figure 4 presents the results of CCID4 and its variants in the presence of different burst and idle periods.

When the media rate is bursty, all variants of CCID4 perform badly giving poor perceived quality. CCID4 performs poorly compared to its variants for all different speech activity parameters. CCID4's sending rate can be at most twice the current receiver rate. This growth rate is not sufficient to keep up with the encoding rate when the application oscillates between silence and talk periods. CCID4-QS does not attempt to solve this problem either. CCID4-FR does solve this problem to a certain extent i.e. based on the previously highest achieved rate, the sender can quadruple up to that rate in the next RTT. Thus the performance of CCID4-FR is better compared to CCID4-QS. However the current FR proposal states that the sending rate can reduce to a minimum of 8 packets per RTT after a timeout, and does not have any recommendations on the minimum sending rate when the sender is datalimited. This means that if the application sends one packet per 2 RTTs, then the receiver rate could be less than the initial sending rate. This

rate may not be sufficient when the application starts sending more packets in the next RTT. During a large idle period (more than 4 RTTs), the nofeedback timer expires and the sending rate of CCID4 and CCID4-QS reduces to a minimum of 4 packets per RTT. Slowstarting from a rate of 4 packets per RTT to the encoding rate takes a long time. This mismatch between the application encoding rate and CCID4's sending rate could lead to packet drops in the transmit buffer queue. CCID4-FR's sending rate can reduce up to a minimum of 8 packets per RTT and can ramp up quickly. CCID4-QS can ramp up to the encoding rate in one RTT. Thus the performance of CCID4-QS and CCID4-FR are better compared to CCID4.

Figure 5 presents the impact subfactors (I_e and I_d) that affect the R score. The impact caused due to loss is severe compared to impact caused due to delay. CCID-FR exhibits lower impact due to loss compared to CCID4-QS since it can quadruple up to the maximum achieved rate in the next RTT. CCID4-QS can only ramp up quickly after a nofeedback timer has expired, whereas CCID-FR can ramp up quickly even without timeouts, if it has previously achieved a higher rate and there is currently no congestion. This improvement reduces packet loss in the transmit buffer when the application sends more packets and hence the subfactor I_e for CCID4-

FR is less compared to CCID4-QS. In some cases (3.0,3.5) and (5.0,5.5), there were more timeouts, and thus CCID4-QS performed better. Both variants have a smaller I_e compared to CCID4. The impact subfactor I_d is similar for CCID4 and its variants.

Tests were also performed to determine jitter (both mean and std deviation). Figure 6, shows that when the media rate is bursty, the average jitter along with the standard deviation (std) is large. This is due to CCID4's inability to handle bursty traffic and not due to the satellite link.

CCID4-FR exhibits lower jitter compared to CCID4-QS. This is because the CCID4-FR's sending rate is not reduced below eight packets per RTT during a timeout and also because of its ability to quadruple its sending rate. This is sufficient to transmit the packets buffered in the transmit buffer of 5 packets after an idle period. Whereas CCID4-QS and CCID4 specify a sending rate limit of four packets per RTT after a nofeedback timer expiry and do not have provision for ramping up quickly during shorter idle periods. This causes packets to be queued in the buffer for a longer time and hence the increased jitter. This raises issues in determining appropriate buffer management methods that need to be found to deploy DCCP for this type of application.

In summary, using DCCP as the underlying transport protocol for carrying VoIP traffic can give medium perceived quality when the media sends at a constant rate. However sending media at a constant rate without silence suppression wastes capacity and is not optimal when using satellite links, where capacity is often either a limited or an expensive resource. An application that uses silence suppression causes traffic that is bursty. Due to the overly conservative nature of the currently defined DCCP CCIDs, carrying bursty VoIP traffic over DCCP leads to a poor perceived quality.

The paper considered two mitigations, CCID4-FR and CCID4-QS. Even though these variants provide faster access to network capacity and therefore improve the performance of bursty VoIP traffic, there are still issues that need to be addressed. Jitter is small when the VoIP traffic is encoded using a constant rate codec, but increases for bursty VoIP traffic using silence suppression. Sender buffering policy is also a significant factor impacting user performance for bursty traffic sources (such as speech with VAD). These issues impact performance, particularly for high-delay paths, including satellite links, and demand further development of the DCCP transport protocol interface and CCIDs.

VI. CONCLUSIONS

In this paper, we analyzed the performance of VoIP traffic using DCCP CCID4 over a DVB-RCS network. This confirms that voice calls of medium perceived quality can be achieved over a DVB-RCS network. The result was achieved for constant rate traffic using DCCP. However, trends in the global Internet are seeing the emergence of VoIP traffic utilizing silence suppression and the introduction of congestion control methods for multimedia traffic. The resultant bursty traffic raises some fundamental design issues for CCID4 while operating over large delay links such as satellite. Future work

is needed to identify ways to address these design issues which may lead DCCP to a more appropriate protocol for transferring VoIP traffic over satellite links.

VII. ACKNOWLEDGEMENTS

We would like to thank Raffaello Secchi from ISTI for providing and helping us understand the intricacies of the TDMA-BoD code. This work was supported by the EU Information Society Technologies SATSIX project, IST-2-26950.

REFERENCES

- [1] A. Sathiseelan, G. Fairhurst, Use of Quickstart for Improving the Performance of TFRC-SP Over Satellite Networks, International Workshop on Satellite and Space Communications (IWSSC2006), Spain, 2006.
- [2] EN 301790, "Digital video broadcasting (DVB); interaction channel for satellite system distribution", ETSI.
- [3] S. Floyd, M. Allman, A. Jain, Quick-Start for TCP and IP, IETF Work in progress, draft-ietf-tsvwg-quickstart-02, 2006.
- [4] S. Floyd, E. Kohler, TCP Friendly Rate Control (TFRC): the Small-Packet (SP) Variant, IETF Work in progress, draft-ietf-dccp-tfrc-voip-05, 2006.
- [5] S. Floyd, M. Handley, E. Kohler, Problem Statement for the Datagram Congestion Control Protocol (DCCP), IETF RFC 4336, 2006.
- [6] S. Floyd, E. Kohler, Profile for Datagram Congestion Control Protocol (DCCP) Congestion Control ID 2: TCP-like Congestion Control, IETF RFC 4341, 2006.
- [7] S. Floyd, E. Kohler, J. Padhye, Profile for Datagram Congestion Control Protocol (DCCP) Congestion Control ID 3: TCP-Friendly Rate Control (TFRC), IETF RFC 4342, 2006.
- [8] S. Floyd, E. Kohler, Profile for DCCP Congestion Control ID 4: the Small-Packet Variant of TFRC Congestion Control, IETF Work in progress, draft-floyd-ccid4-00, 2006.
- [9] M. Handley, S. Floyd, J. Padhye, J. Widmer, TCP Friendly Rate Control (TFRC): Protocol Specification, IETF RFC 3448, 2003.
- [10] E. Kohler, S. Floyd, Faster Restart for TCP Friendly Rate Control (TFRC), IETF Work in progress, draft-ietf-dccp-tfrc-faster-restart-01, June 2006.
- [11] V. Jacobson, Congestion Avoidance and Control, ACM Sigcomm, pp. 273-288, 1988.
- [12] S. McCanne, S. Floyd, Network Simulator. <http://www.isi.edu/nsnam/ns/>
- [13] "Ns2 simulator for satellite platform", <http://votos.isti.cnr.it/simulator.htm>.
- [14] A. Yun, J. Casas, J. Prat, "AMERHIS, a multimedia switching node in the space", 9th Ka Band Broadband Communications Conference, Italy, 5-7 Nov 2003.
- [15] J. Postel, User Datagram Protocol, IETF RFC 768, 1980.
- [16] H. Schulzrinne, S. Casner, R. Frederick, V. Jacobson, "RTP: A Transport Protocol for Real-Time Applications", IETF RFC 1889, 1996.
- [17] G.107, "The E-model, a computational model for use in transmission planning", ITU-T, 2005.
- [18] K. Sriram, W. Whitt, Characterizing superposition arrival processes and the performance of multiplexers for voice and data", IEEE Globecom, New Orleans, December 1985.
- [19] V. Castro, G. Serrano, M. Fernandez, M. Moya, "Quality of Service of VoIP over DVB-RCS", Sixth Baiona Workshop on Signal Processing in Communications, 2003.
- [20] Nera, "VoIP over DVB-RCS: A Radio Resource and QoS Perspective", White Paper, http://satlabs.org/pdf/VoIP_over_DVB-RCS_by_NERA.pdf.
- [21] H. Skinnemoen, A. Vermesan, A. Iuoras, G. Adams, X. Lobao, "VoIP with QoS and Bandwidth-on-Demand for DVB-RCS", White Paper, http://satlabs.org/pdf/VoIP_with_QoS_and_BoD_for_DVB-RCS.pdf.
- [22] M. Lambert, "VoIP over Satellite: An EMS Technologies Canada Technical Notes", White Paper, http://satlabs.org/pdf/VoIP_over_DVB-RCS_by_EMS.pdf.
- [23] J. Ott, S. Prella, D. Meyer, O. Bergmann, "IPTTEL-via-Sat: Cookbook for IP Telephony via DVB-RCS", ESA Project Report, 2005.
- [24] 3GPP TR 26.976 v6.0.0 (2004-12), "Performance characterization of the Adaptive Multi-Rate Wideband (AMR-WB) speech codec (Release 6)".