



Experimental performance of DCCP over live satellite and long range wireless links

Golam Sawar, Emmanuel Lochin, Roksana Boreli

► To cite this version:

Golam Sawar, Emmanuel Lochin, Roksana Boreli. Experimental performance of DCCP over live satellite and long range wireless links. 2007 International Symposium on Communications and Information Technologies, Oct 2007, Sydney, Australia. pp.689-694, 10.1109/ISCIT.2007.4392105 . hal-02871152

HAL Id: hal-02871152

<https://hal.science/hal-02871152>

Submitted on 23 Jun 2020

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Experimental performance of DCCP over live satellite and long range wireless links

Golam Sawar^{1,2} Emmanuel Lochin¹ Roksana Boreli^{1,2,3}

¹ National ICT Australia Ltd, Australia,

² University of NSW, Kensington, Australia,

³ 7-ip, Eveleigh, Sydney, Australia,

{firstname.lastname}@nicta.com.au

Abstract—We present experimental results for the performance over satellite and long range wireless (WiMax) links of the new TCP-Friendly Rate Control (TFRC) congestion control mechanism from the Datagram Congestion Control Protocol (DCCP) proposed for use with real-time traffic. We evaluate the performance of the standard DCCP/CCID3 algorithm and identify two problem areas: the measured DCCP/CCID3 rate is inferior to the rate achievable with standard TCP and a significant rate oscillation continuously occurs making the resulting rate variable even in the short term. We analyse the links and identify the potential causes, i.e. long and variable delay and link errors. As a second contribution, we propose a change in the DCCP/CCID3 algorithm in which the number of feedback messages is increased from the currently standard of at least one per return trip time. Although it is recognised that the increase in control traffic may decrease the overall efficiency, we demonstrate that the change results in higher data rates which are closer to what is achievable with TCP on those networks and that the overhead introduced remains acceptable.

Index Terms—DCCP, rate-based protocols, TCP friendly protocols

I. INTRODUCTION

Datagram control protocol DCCP [20] is a recently standardized protocol filling the gap between TCP and UDP protocols. Unlike TCP, it does not support reliable data delivery and unlike UDP, it provides a TCP-friendly congestion control mechanism in order to behave in a fair manner with other TCP flows. DCCP includes multiple congestion control algorithms which can be selected in regards to the user needs.

DCCP identifies a congestion control algorithm through its Congestion Control ID (CCID). Two CCIDs are now standardized by the Internet Engineering Task Force (IETF). CCID2: a window based congestion control algorithm similar to TCP and CCID3 [17]: a TCP-Friendly Rate Control (TFRC) algorithm. A variant of CCID3 (CCID4) for small packets size is currently under discussion [16] and aims at improving VoIP communications.

CCID2 is appropriate for senders who would like to take advantage of the available bandwidth in an environment with rapidly changing conditions as bursty real-time traffic such as traffic from compressed encoded video and network games. CCID3 is suitable for traffic with smooth changes in sending rates, such as telephony or video streaming. The CCID3 is based on the TCP throughput equation [19] and is designed to be reasonably fair when competing with TCP flow.

Satellite networks have long been a viable communications alternative for rural and remote areas and a number of new systems has been recently launched ([1], [8], [2]). Long range wireless networks like WiMax [9] are also increasingly deployed worldwide. The purpose of introducing DCCP was to better handle multimedia traffic and provide a degree of transmission control for real time traffic, not provided by UDP. The increasing deployment of satellite and wireless networks, which have different transmission characteristics in regards to delay and errors when compared to fixed networks, poses a question about suitability of new protocols like DCCP for use in those networks. As the performance of TCP over satellite and wireless links has been extensively evaluated with a number of proposed protocol enhancements [10], [12], it is likely that the TCP-like (CCID2) or TCP-friendly (CCID3) DCCP algorithms will also have performance difficulties with these networks.

In this paper, we present results of experimental evaluation of the performance of DCCP/CCID3 (DCCP with TFRC congestion control) over IPSTAR satellite network [2] and Unwired wireless network [6], both of which are operational in Australia. To the best of our knowledge, this study is the first one presenting results from real live network performance measurements of the new implementation of DCCP/CCID3 protocol [7].

The rest of the paper is structured as follows. Section II provides an overview of the related work, with a summary of DCCP protocol explained in Section III. Section IV outlines the experimental setup and the properties of the satellite and wireless links used for DCCP evaluation. The main body of our work is described in Section V: we present experimental results on the selected networks, analyse the results and identify the problem areas. In Section VI we propose an improvement to the DCCP protocol feedback which provides increased reliability and accuracy of the bandwidth estimate. Conclusions and a discussion of future work is presented in Section VII.

II. RELATED WORK

TFRC [19] is an equation-based congestion control mechanism involving both sender and receiver to compute an optimal sending rate following the network conditions. One of the main properties of TFRC is to remain friendly when competing

for the bandwidth with others TCP flows. These days, the definition of the TCP-friendliness is still being debated [11]. In this study, we will follow the definition from RFC3448 saying: “(...) a flow is “reasonably fair” if its sending rate is generally within a factor of two of the sending rate of a TCP flow under the same conditions.”. This definition concerns instantaneous values meaning equivalent rates should be achieved on average.

TFRC is suitable for multimedia streaming and VoIP applications as it allows a smoother sending rate and thus, much lower variation of throughput over time compared to TCP. This congestion control mechanism is currently implemented inside DCCP/CCID3 [20], [17] and its variant for small packets size in DCCP/CCID4 [16] protocols. However, this makes TFRC respond slower to changes in available bandwidth when compared to TCP. The impact of long delay on the DCCP/CCID3 and CCID4 performances in the context of satellite delay has been recently studied through simulation [14]. The authors evaluate in [14], [13], using ns-2 [4] simulation, the impact of recent mechanisms such as Quick Start [21] and Fast Restart [15] on the overall DCCP performance and show that these new mechanisms do not provide a consistent improvement to poor TFRC performance over a long delay network.

In this paper, we concentrate on experimental results over live and operational satellite and wireless networks and propose modification to the way DCCP control behaves in those networks.

III. DCCP/CCID3 PROTOCOL

DCCP was designed as a packet stream transport protocol. Applications which would want to use TCP features like congestion control and want to take advantage of flow based semantics of TCP without other TCP features like orderliness and strict reliability would find DCCP most appropriate. One of the primary goals of DCCP was to have congestion control in the transport layer which helps removing the necessity of having congestion control in the application layer as is the case when using UDP. Additionally, all the applications which would use DCCP will have the same standard congestions control mechanism. As another positive point, DCCP will nicely collaborate with the other existing congestion controlled transport protocol in the network like TCP.

RFC4342 provides a proficient and detailed outline on how TFRC can be incorporated into DCCP. TFRC is basically regulated by the rule which strictly says, it’s a receiver-based mechanism, where the sender’s transmission rate is purely guided by the feedback with TCP traffic sent to the receiver. The receiver provides feedback to the transmitter via control packets which contain two measured parameters: round trip time (RTT) and the loss event rate.

To evaluate throughput X , DCCP/CCID3 uses the following equation (1):

$$X = \frac{s}{(RTT \cdot \sqrt{\frac{p \cdot 2}{3}} + RTO \cdot \sqrt{\frac{p \cdot 27}{8}} \cdot p \cdot (1 + 32 \cdot p^2))} \quad (1)$$

Where:

- X is the transmit rate in bytes/second;
- s is the packet size in bytes;
- RTT the round trip time in seconds;
- p is the loss event rate, between 0 and 1.0, to represent the number of loss events as a fraction of the number of packets transmitted;
- RTO is the TCP retransmission timeout value in seconds.

As the equation based mechanism is RTT and loss dependant, the calculated values will plateau for high error rates and high RTTs. Additionally, the number of feedback packets is by default one per RTT which in lossy and high delay links could provide unreliable information (delayed i.e. incorrect by the time it is received, or dropped). Together with the increasing proliferation of satellite and wireless networks, these observations have also motivated our experimental work in DCCP over satellite and WiMax wireless networks.

IV. EXPERIMENTAL ENVIRONMENT

The experimental setup is presented Figure 1 for the satellite connection. In case of WiMax experiments, the satellite terminal is replaced by a WiMax modem and satellite gateway by a WiMax base station. The DCCP-capable application using a DCCP protocol client was connected to the DCCP protocol server and the application server, through a satellite or wireless link.

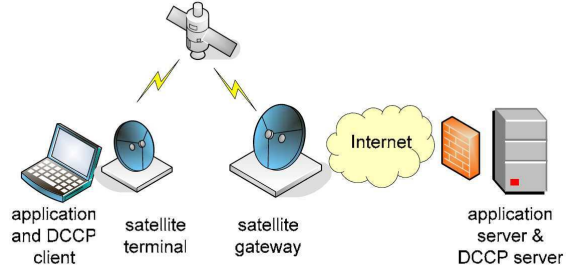


Fig. 1. Experimental setup

Both the satellite and wireless links operate through a local (Australian based) gateway, which is connected to the public Internet.

A. Network characteristics

To have a better understanding of the potential issues, we give an overview of the network round trip time and loss characterised during our experiments.

The satellite service used was IPSTAR plan with a nominal uplink and downlink bandwidth of 256kbit/s and 1Mbit/s. The Unwired wireless service plan used has the same nominal uplink and downlink bandwidth.

The RTT and packet error rate (PER) figures for both the satellite and wireless network are presented in Table I. RTT was measured with a series of ping measurements with packet size equivalent to standard DCCP packet size [17] and error

	RTT	PER
	mean/std.dev. (ms)	
IPSTAR	1206.3 / 35.7	0.24% - 0.42%
Unwired	519.2 / 128.3	0.48%

TABLE I
NETWORK CHARACTERISTICS

rate was measured both by IPERF with UDP protocol for different data rates and by a comparison of transmitted and received data streams in TCP and DCCP experiments which are presented in Section V.

To roughly estimate the likely performance of DCCP and provide an initial sanity check, we have calculated the range of values resulting from the equation (1) and the packet error rate and delay typical for IPSTAR and Unwired networks.

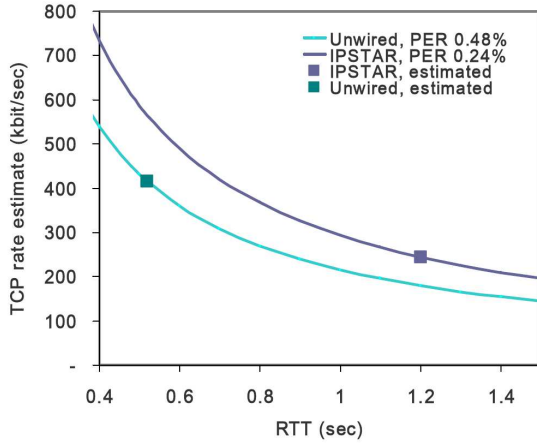


Fig. 2. DCCP/CCID3 based on equation (1) with fixed PER values and RTT ranging from 400ms to 1400ms.

Figure 2 shows the calculated data rate estimated by equation (1) for RTT values between 0.4 and 1.4 sec and for the lowest value of packet error rates (PER) listed in Table I.

B. Experimental configuration

We are using the DCCP/CCID3 Linux kernel implementation from Waikato University [7]. As the current implementation does not handle network address translation which is used in the IPSTAR modem/router, we are also using UDP encapsulation of DCCP packets [5]. The UDP packet size used is 1500B, which ensures there is no fragmentation of the DCCP packets (1000B sized). This adds overhead onto the DCCP data stream and hence decreases the observed data rate, however we have done all experiments with both DCCP and TCP in the same configuration and hence the comparative values which are of interest are not affected. IPERF 2.02 with DCCP patch [7] was used for all experiments.

Standard DCCP/CCID3 default parameters, as recommended in [7], were used for all tests.

	mean (kbit/s)	std.dev. (kbit/s)
TCP upload rate	216.16	16.1
TCP download rate	629.8	48.5

TABLE II
TCP UPLOAD AND DOWNLOAD VALUES OVER IPSTAR

V. RESULTS AND PERFORMANCE ANALYSIS

We have performed a series of tests to evaluate the performance of DCCP over IPSTAR and Unwired networks. The objective was to compare the DCCP achievable rate and rate variations with TCP. For all TCP experiments, we use the standard Linux TCP (kernel 2.6.20 with SACK enabled by default). Please note that the TCP results were sub-optimal as a standard TCP was used, rather than either a satellite-specific TCP variant like [18], [3] or the TCP proxy included in the IPSTAR terminal. All tests were done with no additional network loading, i.e. the tested application was the only traffic source on the link.

This section presents representative results of our experiments, provides a summary of all experimental results and analyses the DCCP performance. The upload and download measurements relate to the measurements of the uplink and downlink performance from the data receiver side point of view.

A. Performance over the IPSTAR satellite network

Received data rate for DCCP upload and download measured on the IPSTAR satellite network is presented Table III. To illustrate the DCCP behaviour, we have randomly chosen one upload and one download results shown in Figures 3 and 4. For comparison purpose, the mean and standard deviation values for corresponding received rate for ten TCP experiments are presented in Table II.

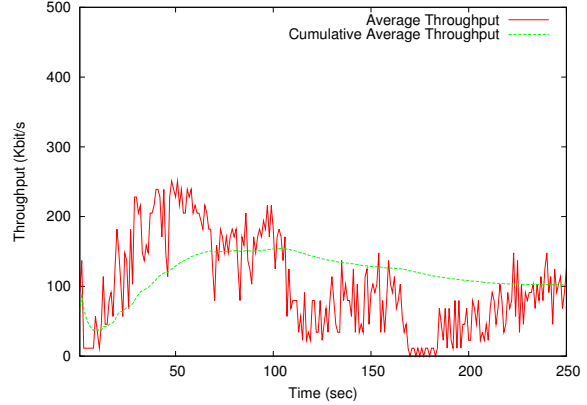


Fig. 3. DCCP upload over IPSTAR

The average performance of the ten experiments over IPSTAR is given in the last row of Table III.

It can be observed that TCP significantly outperforms DCCP over the same link and identical traffic loading, particularly

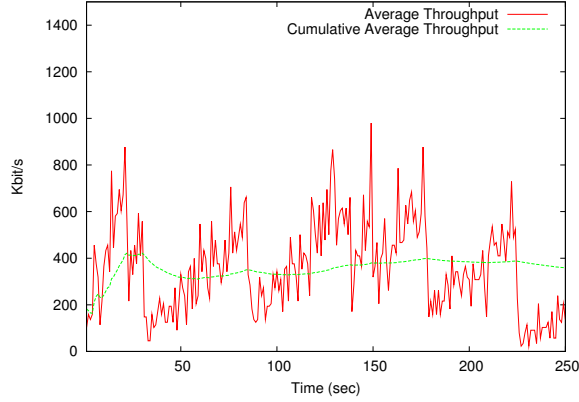


Fig. 4. DCCP download over IPSTAR

Run #	Avg. Rate upload (kbit/s)	Std.Dev. upload (kbit/s)	Avg. Rate download (kbit/s)	Std.Dev. download (kbit/s)
1	179.0	24.4	374.2	78.7
2	175.0	41.6	370.2	32.8
3	164.6	23.5	348.6	30.9
4	172.8	34.0	339.4	46.3
5	172.9	29.6	375.9	46.4
6	170.6	24.4	378.9	49.7
7	170.2	25.6	357.5	62.0
8	164.8	36.0	357.9	31.5
9	163.6	36.5	374.5	34.5
10	173.9	33.7	357.7	31.9
Overall	170.7	5.05	363.5	13.23

TABLE III
DCCP UPLOAD AND DOWNLOAD OVER IPSTAR

in the downlink experiments. This is an expected result, as the poor performance of TFRC protocol with long delays has been identified in [22]. It is also observed that the received rate significantly varies which considering the intended use of DCCP/CCID3 for real time applications like IP telephony is a cause for concern and must be addressed if this protocol is to be used on satellite links. In Section VI we consider initial improvement and propose a modification to the DCCP feedback which improves the received rate.

B. Performance over the Unwired wireless network

The following tables and diagrams show a comparison of the received data rate for TCP and DCCP over the Unwired WiMax wireless network. Figures 5 and 6 illustrate the results of a randomly chosen sample DCCP upload and download run from Table IV.

Ten TCP upload and download measurements are summarised in Table V.

The average performance of the ten experiments over Unwired is given in the last row of Table IV.

It can be observed that, the DCCP performance is again inferior compared to TCP and that there is room for improvement in both uplink and downlink directions. Similarly to

Run #	Avg. Rate upload (kbit/s)	Std.Dev. upload (kbit/s)	Avg. Rate download (kbit/s)	Std.Dev. download (kbit/s)
1	93.6	13.4	321.4	54.4
2	95.6	11.3	322.6	58.2
3	118.3	12.2	364.7	79.9
4	121.3	11.7	298.5	52.1
5	111.1	16.2	320.5	59.1
6	115.5	15.3	301.5	56.2
7	111.5	10.5	276.4	54.8
8	112.5	13.4	318.6	48.3
9	116.2	15.9	325.5	74.9
10	113.3	11.7	337.5	53.7
Overall	110.9	9.16	318.7	23.25

TABLE IV
DCCP UPLOAD AND DOWNLOAD OVER UNWIRED

	mean (kbit/s)	std.dev. (kbit/s)
TCP upload rate	199.6	5.1
TCP download rate	723.5	23.0

TABLE V
TCP UPLOAD AND DOWNLOAD VALUES OVER UNWIRED

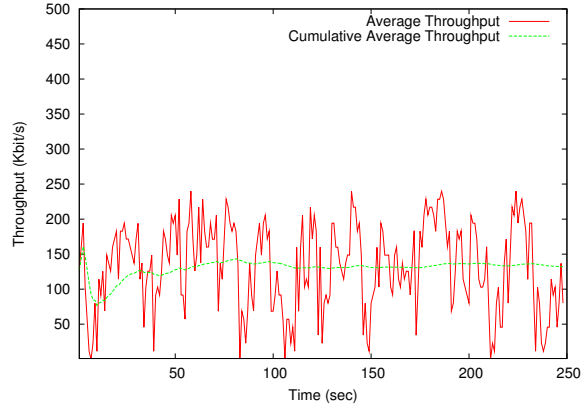


Fig. 5. DCCP upload over Unwired

the performance over the satellite link, there is a significant variation of the received rate in measured time intervals.

C. Discussion on the performance results

It is clear from the results presented in this section that the rate achievable over either a satellite or a wireless link does not match, particularly in the case of IPSTAR satellite network, the performance of a TCP connection with similar network loading and application parameters. This observation follows what we have expected based on Figure 2 as Unwired has a higher PER and IPSTAR a much higher RTT. So, the resulting performance for both networks is expected to be in the similar range. The uplink results are limited by the available nominal bandwidth of 256kbit/s. As a result, the formula is more applicable to the downlink results.

The causes of this mismatch are logically likely to be either packet losses in the control (return) DCCP path or the lateness

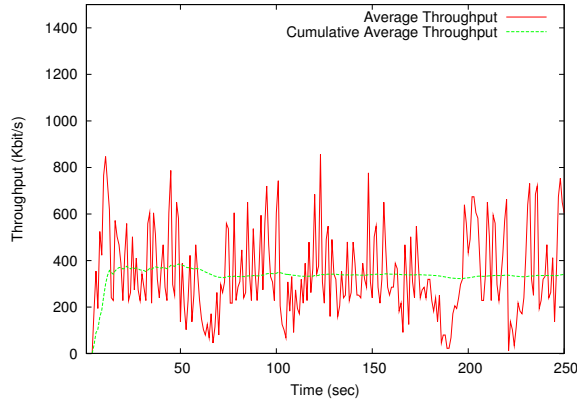


Fig. 6. DCCP download over Unwired

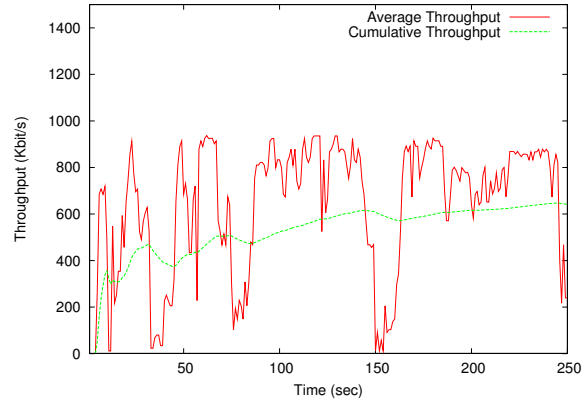


Fig. 7. DCCP download over IPSTAR, 10 feedback messages per RTT

# feedback messages	Average rate download (kbit/s)	Std.Dev. download (kbit/s)
1	363.5	13.23
10	552.3	121.6
100	572.8	134.8

TABLE VI
DCCP DOWNLOAD OVER IPSTAR

and therefore inaccuracy at the time of arrival of the feedback for calculating DCCP rate in long delay links.

Finally, these measurements allow to confirm the trend of the results obtained by simulation in [13], [14], [7].

VI. DCCP IMPROVEMENT

In the TFRC algorithm, the transmitting rate estimated at the sender side depends on the last feedback received. Indeed, this feedback reports keys information measured at the receiver side about the network state in order to allow the sender to transmit at an optimal rate. In a logical way, on such long delay networks, the feedback loop is linked to the delay of the connection and thus, cannot update the transmitted rate as efficiently as over a low delay network. Unlike TCP, DCCP is not a self-clocked protocol (i.e. data emission is not acknowledged). Thus, it could be interesting to study the increase of the number of feedback messages over these long delay networks in order to improve the transmitted throughput. To explore the impact of having more accurate control feedback, we have increased the number of DCCP control messages, which is by default one per RTT.

The following figures present a representative result from the download experiments on IPSTAR network, with increasing the number of DCCP feedback messages from 1 to 10 (Figure 7) and to 100 (Figure 8). The statistical analysis of these results is presented in Table VI, with the average rate of the standard DCCP download also presented for comparison purposes (overall results row from Table III).

It can clearly be observed that increasing the number of feedback messages results in a significant improvement in the

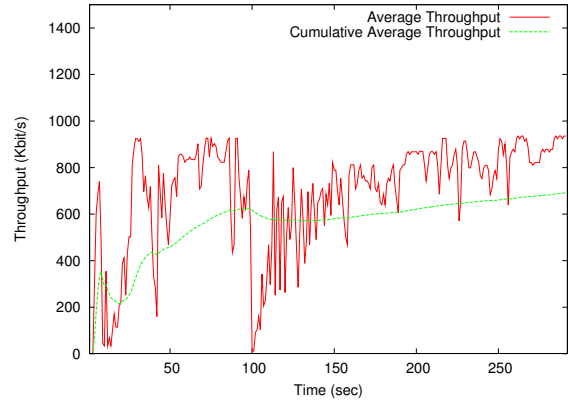


Fig. 8. DCCP download over IPSTAR, 100 feedback messages per RTT

achievable data rate. This is not matched by the improvement in the stability of the rate, however the results presented here are an initial evaluation and we plan to further explore both the dynamic changes to the feedback according to measured link parameters and potential DCCP modification in terms of algorithm, performance and friendliness.

It is recognised that increasing the frequency of feedback packets may decrease the efficiency of DCCP as it generates additional control traffic. However the real amount of feedback generated on any network is variable, as it is directly proportional to the RTT value, so a higher RTT allows for a higher number of feedback messages without impacting the network capacity. E.g. if we compare the amount of traffic generated by DCCP with 12 control messages per RTT on IPSTAR which has an RTT of 1200msec to the amount of control messages on a wired network with RTT of 100msec and DCCP with one control message per RTT, the traffic from feedback messages is identical.

Hence, we would like to approach the calculation of the “equivalent” number of feedback messages by having a con-

stant number of messages in a determined minimum time period, rather than per RTT which is how DCCP currently defines this. From the results presented in previous sections we can expect an improved DCCP rate for other long delay links, while not impacting the network loading any higher than is the case for lower delay links. We plan to further investigate the performance of satellite and wireless links and consider a dynamic algorithm for adjusting the amount of DCCP feedback based on RTT.

VII. CONCLUSIONS AND FUTURE WORK

We have presented results of experimental measurements which address the performance of DCCP over satellite and WiMax wireless links, with our experiments conducted on live networks operational in Sydney, Australia. To assist with analysing the DCCP performance we have also presented the characteristics of both networks. We have compared DCCP performance to the performance of standard Linux TCP on those networks and concluded that TCP significantly outperforms DCCP in regards to received data rates achieved under equivalent network loading. With our measurements, we have confirmed results previously obtained by simulation. We have proposed an improvement in the feedback control which results in a higher achievable rate when using DCCP on long delay networks.

We plan to further investigate DCCP over satellite and wireless links and evaluate performance of applications like VoIP. Further study will also address the tradeoff between minimising the amount of control feedback and maximising the achievable rate in DCCP.

VIII. ACKNOWLEDGEMENT

This research work has been supported by funding from National ICT Australia (NICTA). The authors would like to thank Sebastien Ardon and Guillaume Jourjon for the discussion about the DCCP improvements and Ian McDonald for use of his DCCP implementation.

REFERENCES

- [1] Inmarsat bgan. <http://bgan.inmarsat.com/>.
- [2] Ipstar. <http://ipstar.com/>.
- [3] Multitcp. <http://sourceforge.net/projects/multitcp/>.
- [4] ns-2 network simulator. <http://www.isi.edu/nsnam/ns/>.
- [5] Openvpn project. <http://openvpn.net/>.
- [6] Unwired. <http://unwired.com.au/>.
- [7] Waikato university, dccp site. <http://wand.net.nz/iam4/dccp/index.htm>.
- [8] Wildblue. <http://wildblue.com/>.
- [9] IEEE Standard 802.16, Air Interface for Fixed Broadband Wireless Access Systems, 2004.
- [10] H. Balakrishnan, V.N. Padmanabhan, S. Seshan, and R.H. Katz. A Comparison of Mechanisms for Improving TCP Performance over Wireless Links. In *Proc. of ACM SIGCOMM*, pages 256–269, August 1996.
- [11] Bob Briscoe. Flow Rate Fairness: Dismantling a Religion. *ACM Computer Communications Review*, 37(2):63–74, April 2007.
- [12] Marcondes C., Persson A., Sanadidi M.Y., Gerla M., Shimonishi H., Hama T., and Murase T. Inline Path Characteristic Estimation to Improve TCP Performance in High Bandwidth-Delay Networks. In *International Workshop on Protocols for Fast Long-Distance Networks (PFLDNET 2006)*, Nara, Japan, February 2006.
- [13] Arjuna Sathiseelan Gorry Fairhurst. Use of Quickstart for Improving the Performance of TFRC-SP Over Satellite Networks. Leganes, Spain, September 2006.
- [14] Arjuna Sathiseelan Gorry Fairhurst. Performance of VoIP using DCCP over a DVB-RCS Satellite Network. In *IEEE International Conference on Communications (ICC)*, Glasgow, UK, June 2007.
- [15] Sally Floyd and Eddie Kohler. Faster Restart for TCP Friendly Rate Control (TFRC) - draft-ietf-dccp-tfrc-faster-restart. Internet draft, IETF, 2006.
- [16] Sally Floyd and Eddie Kohler. Profile for DCCP Congestion Control ID 4: the Small-Packet Variant of TFRC Congestion Control. Internet draft, IETF, 2006.
- [17] Sally Floyd, Eddie Kohler, and Jitendra Padhye. Profile for DCCP Congestion Control ID 3: TFRC Congestion Control. Request For Comments 4342, IETF, March 2006.
- [18] G. Giambene, C. Caini, Nico C. Liberato, and R. Firrincieli. TCP Hybla Performance in GEO Satellite Networks: Simulations and Testbed. In *International Workshop on Satellite and Space Communications (IWSSC)*, Leganes, Spain, September 2006.
- [19] M. Handley, S. Floyd, J. Padhye, and J. Widmer. TCP-Friendly Rate Control (TFRC): Protocol Specification. Request For Comments 3448, IETF, January 2003.
- [20] Eddie Kohler, Mark Handley, and Sally Floyd. Datagram Congestion Control Protocol (DCCP). Request For Comments 4340, IETF, March 2006.
- [21] A. Jain S. Floyd, M. Allman. Quick-Start for TCP and IP - draft-ietf-tsvwg-quickstart. Internet draft, IETF, 2006.
- [22] Jorg Widmer. *Equation-Based Congestion Control*. Diploma thesis, University of Mannheim, Germany, February 2000.