

Application-level QoS : Improving Video Conferencing Quality through Sending the Best Packet Next

Ian McDonald and Richard Nelson

School of Computing and Mathematical Sciences, University of Waikato, Hamilton, New Zealand

Email: iam4@cs.waikato.ac.nz, richardn@cs.waikato.ac.nz

Abstract

In a traditional network stack, data from an application is transmitted in the order that it is received. An algorithm is proposed where information about the priority of packets and expiry times is used by the transport layer to reorder or discard packets at the time of transmission to optimise the use of available bandwidth. This can be used for video conferencing to prioritise important data. This scheme is implemented and compared to unmodified Datagram Congestion Control Protocol (DCCP). This algorithm is implemented as an interface to DCCP and tested using traffic modelled on video conferencing software. The results show improvement can be made to video conferencing during periods of congestion - substantially more audio packets arrive on time with the algorithm, which leads to higher quality video conferencing. In many cases video packet arrival rate also increases and adopting the algorithm gives improvements to video conferencing that are better than using unmodified queuing for DCCP. The algorithm proposed is implemented on the server only, so benefits can be obtained on the client without changes being required to the client.

1. Introduction

Typically an application will prepare data and transfer it to the transport layer where it is segmented into packets. A packet is queued and then sent when it reaches the head of the queue, after all other data before it has been transmitted. This however takes no account of time requirements on the delivery of data, or the possibility of different priority among different parts of the data.

This paper discusses a scheme where the transport layer, in cooperation with the application, determines which packets should be sent and in what order they should be sent in. The motivation for this is two-fold. It seems reasonable for the transport layer to discard packets which will no longer be useful for the receiver. It would also seem reasonable for audio packets to be sent before video packets, as Sasse [27]

shows that audio is more important than video to a user's experience in a video conference situation. This is particularly applicable when the number of layers of data being transmitted is being reduced, or is already at the lowest possible.

Our unique contribution is that we combine priority queues with tracking of expiry times which take into account the Round Trip Time (RTT). This is tested and implemented using the DCCP transport protocol [16] which is congestion controlled but unreliable.

This paper begins with a survey of previous work in Section 2. Section 3 introduces the Send Best Packet Next (SBPN) algorithm and discusses why DCCP was used. In Section 4 the methodology and the test environment are discussed. In Section 5 the findings from testing SBPN are discussed, and in the last section possibilities for further work are outlined.

2. Previous work

In this section previous work aimed at improving the transmission of video is examined. By nature video generates a considerable amount of network traffic that has strict requirements of the network. As such it is often constrained by the network conditions.

It is commonly proposed that video codecs should add and remove video layers as network conditions permit [26, 4]. While these codecs match the traffic generated to the available capacity, they do not ensure that the most important data is transmitted first. In a study of multimedia streaming [10] it was shown that the median time to change to a lower bit-rate stream was around 4 seconds. This indicates that there is scope to improve the user experience during the transition time as it is presumed that a lower bit-rate is being used because of loss. There are also cases where the bit-rate cannot be lowered further and loss occurs. Feng [5] builds upon changing which video layers are used and proposes a priority queue for delivery of pre-recorded video streams where lower priority video layers are

only sent after higher priority video layers have been sent. Krasic [18] investigates storing streaming video in multiple discrete layers which are tailored to the type of client that is requesting the data. The streaming server, which they call priority-progress, then decides the most appropriate layers for the client and maps these to a priority order within each time segment. At the end of each time segment, if the data has not been transmitted, then the data is discarded by a “progress regulator”.

In Tsaoussidis [28] Multimedia Transmission Protocol (MTP) is introduced, which is based on TCP Reno but without guaranteed reliability. Packet priority information is sent as a 2 bit field which determines whether packets are retransmitted or not. This enables MTP to not retransmit data which is of lower priority, as compared to TCP which retransmits all missing data. Time-lined TCP (TLTCP) [20] introduces the concept of tracking expiry time for data. TLTCP marks data with an expiry time as we also propose, and discards data if the expiry time is reached before the packet is sent. Unlike the approach in this paper, TLTCP does not take the RTT into consideration but proposes that it is worthwhile to do so. TLTCP is a partially reliable protocol - if the data has not expired and retransmission is requested by the receiver then the data is retransmitted. If the data has expired, a more recent packet is sent instead as the expired packet no longer has value.

In Papadimitriou Scalable Streaming Video Protocol (SSVP) [23] and Video Transport Protocol (VTP) [2] are introduced. These are layered on top of UDP. SSVP adjusts the send rate through altering the inter-packet gap, VTP smoothes the transmit rate and sends at a fixed rate. Both protocols adaptively alter the layers of video being transmitted. Amer et al [1] introduce a partial order transport service that allows reliable or unreliable partial order service. This allows marking of importance and time values by the application. This is used by the receiver to determine whether lost data is ignored or requested to be retransmitted depending on their temporal value.

Research by Lai and Kohler [17] investigated using late binding ring buffers with DCCP for transmission that were accessible by the application. The application could replace a packet in the buffer with a more recent packet if it had not been transmitted. The aim is for the most timely data to be sent and old data to be discarded.

The previous work shows that there is scope to improve the video transmission by considering which packets get sent when, and whether packets get sent at all. Our research draws together the range of different algorithms from the previous work to improve video conferencing and is described in the next section.

3. Send Best Packet Next

In this section a concept is introduced called Send Best Packet Next (SBPN), different variants of SBPN are also described, and the use of DCCP to implement SBPN is explained.

SBPN is a unique contribution of combining a priority queue, with expiry times for packets that take RTT into account. Discarding expired packets allows newer packets to be sent that are more likely to arrive on time. SBPN is completely sender based; it could be implemented on a server without requiring changes to the receiver.

3.1. DCCP

Datagram Congestion Control Protocol (DCCP) is a session-based unreliable transport protocol for datagrams. DCCP is a Standards Track RFC. Further discussion on the usage of DCCP is contained in the Problem Statement for DCCP [6].

DCCP was used for SBPN as it is a recent protocol designed to accommodate applications such as multimedia streaming and VoIP [25] that do not need reliability. DCCP is session based, tracks RTT and has congestion control in the base protocol.

UDP was not used because it does not track the RTT and does not have congestion control in the base protocol, and it was not desired to implement this at the application level. UDP also has difficulty traversing some firewalls due to its lack of session control and often applications fall back to TCP. It was also decided not to use TCP as it is a reliable protocol and TCP often combines packets, even when the Nagle algorithm [21] is turned off. With TCP being a reliable protocol if data is not received it is retransmitted when, for many applications, it would be better to discard the data.

DCCP has pluggable Congestion Control Identifiers (CCIDs) which allow different congestion control methods to be used within the protocol. CCID3 [8] implements TCP-Friendly Rate Control (TFRC) [11] for DCCP. This was used because it has a smoother response to loss than CCID2 [7] and is thus considered to be more suitable for multimedia [25]. CCID2 is TCP-like in using a congestion window and following many of the TCP semantics although it does not have retransmission.

3.2. SBPN1

The SBPN1 algorithm sorts the transmit queue as a priority queue. All higher priority packets get sent prior to any lower priority packets being sent. Audio packets are marked as higher priority than video packets. The packets in the queue are also marked with an expiry time. When a packet is due to be sent the expiry time is checked against

the current time added to half the RTT. If the expiry time is less than this calculation the packet is discarded rather than being sent.

3.3. SBPN2

One aspect of CCID3 is that it reduces the allowed transmit rate after short idle periods. The formula for allowed sending rates is $X = \max(\min(X_{calc}, 2 * X_{recv}), s/t_{mbi})$ where X_{calc} is the TCP throughput equation [22], X_{recv} is the receive rate since the previous feedback (at least once per RTT), s is the average packet size and t_{mbi} is 64.

This means that in practice the allowed transmit rate falls very quickly when no packets are being transmitted and thus not being received also. This has the impact of reducing the quality of the video conference, particularly during quiet periods. There is an Internet draft for TFRC faster restart [15] that partially addresses this issue.

To address the issue of reduced transmit rate the SBPN1 algorithm was modified to send an expired packet if it is the only packet in the queue to reduce the likelihood of the transmit rate being reduced.

4. Test Methodology

4.1. Video conference model

Packet traces were captured from Ekiga which is an open-source application which includes video transmission. The traces were captured on a 100 Mbps link and consisted of 10 seconds of talking / motion, followed by 5 seconds of silence / still, followed by 5 seconds of talking / motion. The traces were then analysed to see their protocol, timing, packet size, and whether each packet contained audio or video.

UDP was used by Ekiga to transmit audio and video. A minimal amount of TCP was used in the connection setup and for control purposes due to a requirement for reliability. The effects of TCP traffic were excluded from our experiments as they constituted less than 0.1% of the traffic. Ekiga had a fixed time period between audio packets of 20 milliseconds and employs silence suppression. All audio packets were exactly 214 bytes.

Ekiga had a simple transmission scheme with variable frame sizes that was dependent on the amount of motion in the frame. For high motion the average is 4450 bytes per video frame with a standard deviation of 1000 bytes, while for low motion the average is 1075 bytes with a standard deviation of 475 bytes. Frames are sent at the rate of 10 frames per second and are composed of packets with a maximum size of 1026 bytes which are sent as a bunch.

A model of the traffic was then built. Fixed rate audio was used with a payload size of 214 bytes every 20 milliseconds and silence suppression was not used. Video frames in the model were created with an average and a standard deviation for their type that matched the traffic Ekiga generated. With this model audio equates to 86 Kbps and video equates to 289 Kbps which gives a total of 374 Kbps. These rates include all packet overheads, including the MAC headers.

4.2. Test setup

All of the packets in an audio frame or video frame were queued immediately after each other by the application and the time to send was determined by the DCCP CCID3 congestion control. This differs from the applications which used UDP and determined the send time themselves as UDP does not have congestion control at the transport level. If audio and video frames were both due to be transmitted at the same time the application randomly chose which to send first to avoid bias.

An application was written to then replay this stream between two PCs with a Linux PC in the middle using the netem [12] software to add delay, and to rate limit the traffic or create loss. There was a separate application written to track the difference in time between the two PCs so that arrival time could be tracked at the receiver relative to the sender. The receiver recorded how long each packet took to arrive and whether it arrived at all.

Tests were run with symmetric one way delays of 15, 40, 75 milliseconds producing minimum RTTs of 30, 80, 150 milliseconds. For the purpose of simplicity symmetric links were used. It is recognised that not all links are symmetric [24] and the algorithm could easily be adapted if the one way delay could be measured for a link. Two priorities were used in the tests, one for audio and one for video. The number of priorities could be extended for multiple video layers.

Initial testing with the 20 second runs showed a high rate of variability due to the initial connection setup and the random nature of loss/congestion. The test was then done with 60 second runs with the data being repeated three times in each run. The first 3 seconds of data was discarded and the remaining data had significantly reduced variation due to the connection speed stabilising over a longer run. In all of the results in this paper data was generated from 30 runs and an average calculated with 95% confidence intervals from a T-test as recommended in Law [19].

The ITU recommend that delivery time for audio be less than 150 milliseconds and anything between 150 milliseconds and 400 milliseconds has lower quality [14]. With SBPN packets get marked with an expiry time which is 200 milliseconds from the time that the packets are created. The expiry time includes both time in the queue and time in tran-

sit. The choice of 200 milliseconds is arbitrary and it is envisaged that an application would allow the user to adjust this setting to meet their quality requirements.

4.3. Linux kernel implementation

The implementation was done on Linux 2.6.20 with a series of patches to ensure performance of CCID3 matched the TFRC [11] throughput equation.

The sendmsg system call in Linux allows control data to be sent as well as the data. The DCCP implementation of sendmsg was modified to transmit expiry time, priority and method in a structure passed into the msg_control field in the msg_hdr structure. This data is received by the dccp_sendmsg function, which uses the priority information to store the packet in the correct location in the transmit queue. The transmit queue remained as one queue and was changed to a queue sorted first by priority and secondly by expiry time. This was implemented as one queue for reasons of code simplicity, but for wider use it could be implemented as a set of queues. The dccp_write_xmit function starts transmission from the beginning of this priority queue. The function checks the expiry time just before a packet is transmitted and decides whether to discard the packet if the packet has expired.

The method field is used to select which algorithm will be used in testing. The expiry field contains the time that the data will be of no use to the receiver. In the tests this was set to be 200 milliseconds from the creation of the data by the application. These fields are used by all of the SBPN algorithms with the exception of unmodified DCCP.

5. Results

The desired outcome is that the number of on time audio packets increase and that the number of on time video packets decrease less than the audio increase. It should be noted that with the unmodified transmission scheme more audio packets arrived on time than video packets. The reason for this is that all packets for a single video frame get queued for transmission at the same time which means there is a higher probability the buffer will be full and a video packet discarded.

5.1. Testing with loss

SBPN1 was implemented and compared to unmodified DCCP using netem to discard packets between two nodes. A queue length of 5 packets was used which is similar to average UDP buffering. When run under these conditions the SBPN1 algorithm performed worse than unmodified DCCP as shown in Figure 1. Audio and video packet on time arrivals both decreased.

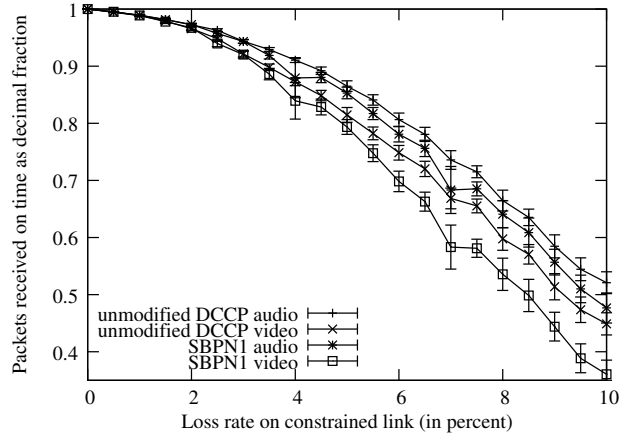


Figure 1. Testing with loss - 80 ms RTT, queue length 5

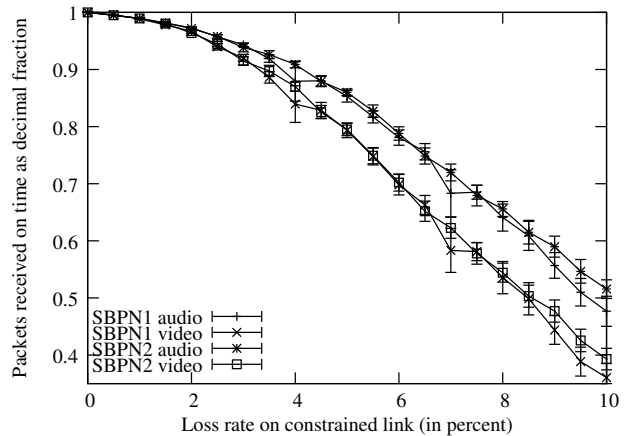


Figure 2. Testing with loss - SBPN1 vs SBPN2 - 80 ms RTT, queue length 5

Testing was carried out with the SBPN2 algorithm which only discards expired packets if there are other packets in the transmit queue. The results for the SBPN1 and SBPN2 algorithms with queue length of 5 packets can be seen in Figure 2. This did improve results compared to the SBPN1 algorithm but did not achieve results that were better than unmodified DCCP. This required some further investigation.

Figure 3 shows the difference between the data received and the data received on time. In the graph SBPN1 has almost no difference between the packets received and the packets received on time as it only transmits when the packet will reach the other end without expiring. SBPN2 has slightly more packets received than packets received on time as it sends data to keep the connection from going idle,

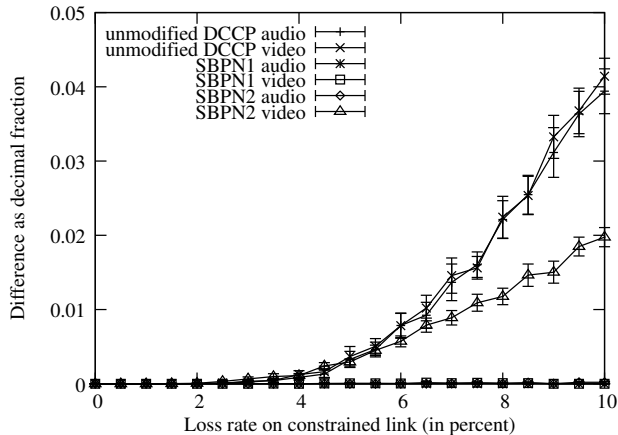


Figure 3. Difference in on time arrival - 80 ms RTT, queue length 5

and unmodified DCCP has even more packets received as it always tries to send. This indicates the improvement is not occurring because of running at a fixed rate of loss and if less packets are being sent then a lesser number of relevant packets will be received due to the smaller number of packets being transmitted. It is interesting to note on this graph that all the packets used to keep alive the connection for the SBPN2 algorithm are video packets - this shows that the audio packets are being transmitted and being received on time successfully. This can be seen from the graph as the difference between received and on time packets are zero (which is also the case for SBPN1).

This effect can also be seen in congestion control on mobile networks where loss is an indication of packet corruption, not congestion. Congestion control algorithms such as Vegas [3] and Veno [9] attempt to find the level of congestion in a network, as well as responding to packet loss - which is not always an indicator of loss. As such if these, or similar, algorithms were implemented for DCCP the results in this section may have been quite different. Implementing these algorithms for DCCP is outside the scope of this paper however.

5.2. Testing with congestion

As it was concluded that fixed loss does not give a realistic test tests were carried out with congestion. This was done by running tests with two competing TCP Reno flows, produced using iperf [13] that was modified to allow continuous running. Within the video conferencing flow there are a range of larger packets for video and smaller packets for audio resulting in an average packet size of 519 bytes. A fixed packet size was used with iperf of this size as DCCP effectively works on a rate of packets per second [29]. TCP




Figure 4. Unmodified DCCP vs SBPN2 - 80 ms RTT, queue length 5




Figure 5. Unmodified DCCP vs SBPN2 - 80 ms RTT, queue length 32

Reno was chosen as DCCP CCID3 uses the TCP throughput equation [22] modelled on Reno.

When the tests were run with competing TCP flows we obtained the results shown in Figure 4. This shows a clear difference for audio between unmodified DCCP and SBPN2. This result shows that more audio packets can be transmitted when there is congestion than under unmodified DCCP.

The algorithm was also run with a buffer size matching TCP. Often media applications will use TCP due to UDP having difficulties traversing NAT and firewall devices. Figure 5 shows the results from using a queue length of 32 which matches the default TCP buffer size of 16 Kbytes in Linux 2.6.20. These results show a larger improvement in audio and smaller reduction in the number of video packets




Figure 6. Unmodified DCCP vs SBPN2 - 30 ms RTT, queue length 5

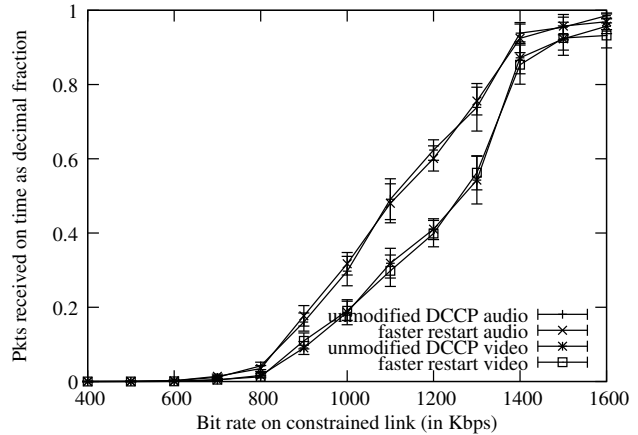


Figure 8. Unmodified DCCP vs Faster Restart - 80 ms RTT, queue length 5




Figure 7. Unmodified DCCP vs SBPN2 - 150 ms RTT, queue length 5

dropped.

To test if the algorithm is RTT dependent, the tests were also run with 30 millisecond and 150 millisecond RTT, shown respectively in Figures 6 and 7. These represent scenarios such as a relatively local network connection, or across greater distances.

TFRC faster restart [15] was also implemented to determine if this would counteract the effect of reduced transmit rates and ran a number of tests. A representative example of the results for this are shown in Figure 8 and this shows that TFRC does not produce significantly better results. (It should be noted that TFRC faster restart was implemented on Linux 2.6.23 so this graph cannot be compared directly to the other graphs.)

6. Conclusion and further work

This paper shows that it is possible to gain a real improvement to video conferencing applications when congestion is experienced. A new class of algorithms have been introduced that automatically bases packet delivery on "send the best packet next" principle. This has resulted in a clear improvement in terms of audio packets received on time when tested with competing TCP flows which will give a higher quality video conference experience. These results have been shown across a range of RTTs and with differing queue lengths. Using the best algorithm (SBPN2) shows audio on time improves by over 10% across most of the range while video on time does not decrease or decreases by a much smaller amount.

There are many areas for potential further research. A form of weighting could be considered for the queues which calculates a score based on packet type, expiry time and any other variables which would then be used to sort the transmit queue. To counter the effects of reduced transmit rate as shown in Section 5.1 it may be worthwhile to implement a version that only discards packets if the buffer is full and sends expired packets if there are not any unexpired packets.

As the algorithms proposed in this paper are made at the sender only, servers can implement these algorithms and give a significant benefit to clients without a change being required to the clients. The algorithms presented could also be used in a range of services such as real time IPTV.

7. Acknowledgements

We would like to thank Matthew Luckie, Sam Jansen and Gerrit Renker for their review and helpful suggestions

for this paper. We also thank the Linux kernel contributors who assisted with the implementation of DCCP.

References

- [1] P. D. Amer, C. Chassot, T. J. Connolly, M. Diaz, and P. Conrad. Partial-order Transport Service for Multimedia and Other Applications. *IEEE/ACM Transactions on Networking*, 2(5):440–456, 1994.
- [2] A. Balk, D. Maggiorini, M. Gerla, and M. Y. Sanadidi. Adaptive MPEG-4 Video Streaming with Bandwidth Estimation. In *Quality of Service in Multiservice IP Networks: Second International Workshop, QoS-IP 2003, Milano, Italy, February 24-26, 2003. Proceedings*, pages 525–538, 2003.
- [3] L. S. Brakmo and L. L. Peterson. TCP Vegas: End to End Congestion Avoidance on a Global Internet. *IEEE Journal on Selected Areas in Communications*, 13(8):1465–1480, 1995.
- [4] N. Feamster, D. Bansal, and H. Balakrishnan. On the Interactions Between Layered Quality Adaptation and Congestion Control for Streaming Video. *11th International Packet Video Workshop (PV2001), Kyongju, Korea*, April 2001.
- [5] W. Feng, M. Liu, B. Krishnaswami, and A. Prabhudev. A Priority-Based Technique for the Best-Effort Delivery of Stored Video. *Proc. of Multimedia Computing and Networking*, 1999.
- [6] S. Floyd, M. Handley, and E. Kohler. Problem Statement for the Datagram Congestion Control Protocol (DCCP). <ftp://ftp.rfc-editor.org/in-notes/rfc4336.txt>, March 2006.
- [7] S. Floyd and E. Kohler. Profile for DCCP Congestion Control ID 2: TCP-like Congestion Control. <ftp://ftp.rfc-editor.org/in-notes/rfc4341.txt>, March 2006.
- [8] S. Floyd, E. Kohler, and J. Padhye. Profile for DCCP Congestion Control ID 3: TFRC Congestion Control. <ftp://ftp.rfc-editor.org/in-notes/rfc4342.txt>, March 2006.
- [9] C. P. Fu and S. C. Liew. TCP VenO: TCP Enhancement for Transmission Over Wireless Access Networks. *IEEE Communications*, 21(2), 2003.
- [10] L. Guo, E. Tan, S. Chen, Z. Xiao, O. Spatscheck, and X. Zhang. Delving into Internet Streaming Media Delivery: A Quality and Resource Utilization Perspective. In *IMC '06: Proceedings of the 6th ACM SIGCOMM on Internet measurement*, pages 217–230, New York, NY, USA, 2006. ACM Press.
- [11] M. Handley, S. Floyd, J. Padhye, and J. Widmer. TCP Friendly Rate Control (TFRC): Protocol Specification. <ftp://ftp.rfc-editor.org/in-notes/rfc3448.txt>, January 2003.
- [12] S. Hemminger. Netem - Emulating Real Networks in the Lab. *Proc. Linux Conference Australia*, 2005.
- [13] NLANR/DAST : Iperf - The TCP/UDP Bandwidth Measurement Tool. <http://dast.nlanr.net/Projects/Iperf/>, Accessed 2007.
- [14] ITU. ITU-T Recommendation G.114. Technical report, International Telecommunication Union, 1993.
- [15] E. Kohler, S. Floyd, and A. Sathiseelan. Faster Restart for TCP Friendly Rate Control (TFRC). Work in progress. <http://www.ietf.org/internet-drafts/draft-ietf-dccp-tfrc-faster-restart-02.txt>, 2007.
- [16] E. Kohler, M. Handley, and S. Floyd. Datagram Congestion Control Protocol (DCCP). <ftp://ftp.rfc-editor.org/in-notes/rfc4340.txt>, March 2006.
- [17] E. Kohler and J. Lai. Efficiency and Late Data Choice in a User-kernel Interface for Congestion-Controlled Datagrams. In *Multimedia Computing and Networking*, 2004.
- [18] C. Krasic, J. Walpole, and W.-C. Feng. Quality-adaptive Media Streaming by Priority Drop. In *NOSSDAV '03: Proceedings of the 13th international workshop on Network and operating systems support for digital audio and video*, pages 112–121, New York, NY, USA, 2003. ACM Press.
- [19] A. M. Law. Practical Statistical Analysis of Simulation Output Data: The State of the Art. In R. G. Ingalls, M. D. Rossetti, J. S. Smith, and B. A. Peters, editors, *Winter Simulation Conference*, pages 67–72. Winter Simulation Conference, December 2004.
- [20] B. Mukherjee and T. Brecht. Time-lined TCP for the TCP-friendly Delivery of Streaming Media. *International Conference on Network Protocols (ICNP), Osaka Japan*, pages 165–176, 2000.
- [21] J. Nagle. Congestion Control in IP/TCP Internetworks. *SIGCOMM Comput. Commun. Rev.*, 14(4):11–17, October 1984.
- [22] J. Padhye, V. Firoiu, D. Towsley, and J. Kurose. Modeling TCP Throughput: A Simple Model and its Empirical Validation. In *SIGCOMM '98: Proceedings of the ACM SIGCOMM '98 conference on Applications, technologies, architectures, and protocols for computer communication*, pages 303–314, New York, NY, USA, 1998. ACM Press.
- [23] P. Papadimitriou and V. Tsaoussidis. SSVp: A congestion control scheme for real-time video streaming. *Computer Networks*, 51(15):4377–4395, October 2007.
- [24] V. Paxson. End-to-end Internet packet dynamics. *IEEE ACM Transactions on Networking*, 7(3):277–292, 1999.
- [25] T. Phelan. Datagram Congestion Control Protocol (DCCP) User Guide. Work in progress. <http://tools.ietf.org/id/draft-ietf-dccp-user-guide-03.txt>, 2005.
- [26] R. Rejaie, M. Handley, and D. Estrin. Quality Adaptation for Congestion Controlled Video Playback Over the Internet. In *SIGCOMM '99: Proceedings of the conference on Applications, technologies, architectures, and protocols for computer communication*, volume 29, pages 189–200, New York, NY, USA, October 1999. ACM Press.
- [27] A. Sasse, U. Bilting, C. D. Schulz, and T. Turletti. Remote Seminars through Multimedia Conferencing: Experiences from the MICE Project. In *Proc. INET'94/JENC5*, pages 251/1–251/8, Prague, June 1994.
- [28] V. Tsaoussidis and S. Wei. Reliability / Throughput / Jitter Tradeoffs for Real-time Transport Protocols. In *Proceedings of the 20th IEEE Real Time Systems Symposium*. IEEE Computer Society Press, December 1999.
- [29] J. Widmer, C. Boutremans, and J.-Y. Le Boudec. End-to-end Congestion Control for TCP-friendly Flows with Variable Packet Size. *SIGCOMM Comput. Commun. Rev.*, 34(2):137–151, April 2004.