

An Experimental Investigation of the End-to-End QoS of the Apple Darwin Streaming Server

Luca De Cicco, Saverio Mascolo and Vittorio Palmisano
{ldecicco, mascolo, vpalmisano}@poliba.it

Dipartimento di Elettrotecnica ed Elettronica, Politecnico di Bari, Via Re David 200,
Italy

Abstract. Video content distribution over the traditional best-effort, store-and-forward Internet Protocol is of ever increasing importance due to the great success of new web services such as personal video broadcast or television over IP (IPTV). In this paper we investigate the end-to-end quality of service (QoS) that is provided by the Apple Darwin Streaming Server and the Quick-Time client player in the presence of time-varying available bandwidth and multiple concurrent streaming sessions. The considered end-to-end QoS parameters are the loss rates and the friendliness experienced when the available bandwidth changes and when multiple QuickTime streaming sessions and/or TCP sessions compete in order to obtain a bandwidth share.

We found that the Darwin Streaming Server implements a TCP-like congestion control that is more aggressive than TCP; in particular, when more QuickTime flows share the same link with TCP flows, QuickTime gets more bandwidth than TCP. Moreover, when more QuickTime flows share the same link, they exhibit a high loss rate.

Keywords: End-to-End QoS, Multimedia Congestion Control, Reliable UDP, Apple Darwin Streaming Server, QuickTime Player

1 Introduction

Audio/Video content distribution is nowadays a potential killer application for the Internet as it is proved by the great success of YouTube¹ and by the introduction of new applications such as Joost² and Babelgum³, which aim at providing television distribution over IP. The most part of Internet traffic is still delivered using the TCP transport protocol, which has been the key factor of Internet stability so far. This is the reason for which many Web sites (such as YouTube) that host small length and low resolution videos use only pseudo-streaming technologies that are based on the simple TCP download. In this way the generated traffic is not harmful for the stability of the Internet because the TCP transport

¹ <http://www.youtube.com/>

² <http://www.joost.com/>

³ <http://www.babelgum.com/>

protocol implements an effective congestion control algorithm [1]. However, it is not clear if the perceived quality is satisfactory for the user. In fact the source of the great success obtained by YouTube is very much likely to be due to the richness of contents and its large user base rather than to the quality of the video delivering.

The TCP window-based congestion control guarantees congestion avoidance by using the additive increase/multiplicative decrease paradigm [1] and reliable delivery of the content through packet retransmissions but not content delivery within delay constraints. On the other hand, multimedia streaming services can tolerate some low packet loss percentage but require more tight quality of service (QoS) requirements in terms of end-to-end delays and jitter. For this reason the UDP protocol is the preferred transport protocol for multimedia streams, because, as matter of fact, it is a simple packet multiplexer/demultiplexer, where the packet sending rate can be managed at the application level. However, many multimedia applications which use UDP do not implement effective congestion control mechanisms, thus possibly leading to a network *congestion collapse* [2] due to the presence of unresponsive flows on the same bottleneck link. This circumstance can cause a high loss rate, which is an important factor that affects the perceived quality [3, 4].

Several efforts have been made to design multimedia congestion control protocols that are TCP friendly, where friendliness here means that the multimedia flows will share the network bandwidth with TCP flows fairly. The TCP Friendly Rate Control (TFRC) [5] protocol and the Datagram Congestion Control Protocol (DCCP) framework [6] are two IETF standards proposed as possible congestion control algorithms for the transport of multimedia flows. An interesting solution is the Reliable UDP proposed by Apple, which is a TCP-like congestion control protocol that aims at providing a set of QoS enhancements for RTP multimedia flows [7] (see Sec. 3 for more details).

The Darwin Streaming Server (DSS) is the open source version of the commercial Apple's QuickTime Streaming Server (QTSS) that allows the distribution of streamed multimedia contents over the Internet. The protocols employed by DSS are the standard RTP and RTCP⁴. DSS is based on the same code base of QTSS, but its source code is freely distributed under the Apple Public Source License. Both DSS and the official commercial QuickTime Player (QTP)⁵ implement the Reliable RTP congestion control. DSS uses well-known standards (such as RTP, RTCP, SDP and HTTP) for content distribution. Thus, every multimedia player that supports RTP can be used as client.

In this paper we have used the official DSS and QTP for investigating the effectiveness of Reliable UDP congestion control algorithm in the presence of changing available bandwidth and/or packet losses. The goal of these investigations is to evaluate how the congestion control algorithm implemented by Reliable UDP allows the sending rates be managed in order to match the available

⁴ <http://developer.apple.com/opensource/server/streaming>

⁵ <http://www.apple.com/it/quicktime/download/>

bandwidth when multiple streaming sessions and/or TCP connections share the same link, thus revealing intra-protocol and inter-protocol fairness behaviour.

The rest of the paper is organized as follows: Section 2 presents the previous work on streaming server performance evaluations; Section 4 describes the considered experimental testbed and the scenarios; Section 5 reports the experimental results and finally Section 6 draws the conclusions.

2 Related Work

Many recent investigations have focused on multimedia streaming client/server applications. In [8] an investigation on the Internet streaming quality and efficiency is performed by collecting connection data from thousands of broadband home users accessing both on-demand and live streaming media. Authors have found that input rate adaptation, even though implemented in media authoring, is poorly utilized, particularly when a pre-buffering phase is used. The pre-buffering phase (also called *Fast Streaming*) is widely used and much quality degradation is caused by re-buffering events.

Authors of [9] present an evaluation of RealVideo streaming over UDP and over TCP. They have found that RealVideo over UDP does not respond to Internet congestion by adapting the sending and/or the encoding rate. In particular, under very constrained bandwidth conditions RealVideo UDP streams do not share the bandwidth fairly with concurrent TCP connections. Moreover, authors report that only the 35% of RealServer implement some form of encoding scalability (called *Media Scaling*), and less than the 50% of the clips were using more than 4 encoding levels so that they can only adapt to the available bandwidth coarsely.

In [10] an investigation of the Windows Streaming Media (WSM) is performed in order to analyze content multiple encoding. They found that, if the network capacity is lower than the minimum available encoding level, WSM produces high packet loss rates and it is unfair with concurrent TCP flows.

In [11] a comparative analysis of RealPlayer, Windows Media Player and Quicktime is performed. They used a UDP cross traffic generator concurrent with each multimedia stream flow in order to emulate a network congestion condition. Authors have found that Quicktime provided the lowest packet loss rate among the applications, thus indicating that DSS performs an effective congestion control algorithm as compared to the other media streaming solutions.

In this work we aim at performing an extensive investigation of DSS by testing it on different scenarios in order to:

1. evaluate how Reliable UDP reacts to congestion episodes;
2. estimate the friendliness between more concurrent QuickTime multimedia flows and between QuickTime flows and TCP flows.

3 Apple's Reliable UDP

Reliable UDP is a set of extensions to the RTP protocol designed in order to provide retransmission and congestion control mechanism to the unreliable UDP protocol. These extensions allow multimedia streams to behave like TCP flows, while providing soft real-time features. Apple's version of Reliable UDP implements a congestion control based on the Additive Increase/Multiplicative Decrease approach: the sender maintains a *congestion window* (CWND) such as the one used by the TCP congestion control; during the *slow-start* phase, for each ACKed packet, CWND increases by 1 segment, whereas during the *congestion avoidance* phase CWND increases by 1 segment every *round trip time* (RTT). When a timeout expires (which in the DSS implementation is always equal to 250 ms) or 3 duplicate ACKs (3DUPACKs) are received the slow-start threshold is set to $3/4$ of CWND and CWND is halved. This behaviour makes Reliable UDP more aggressive than TCP when a loss event is detected, because when 3DUPACKs are received TCP halves the slow-start threshold and enters the congestion avoidance phase, whereas Reliable UDP enters the slow-start phase and then the congestion avoidance phase.

The tests we have carried on confirm this behaviour and show that Reliable UDP tends to use more bandwidth with respect to TCP concurrent flows.

4 Experimental testbed

The testbed we have set up is made of a Linux machine, hosting the DSS, and a Windows machine where the QuickTime players have been installed. The Linux machine has been equipped with a tool developed by us (*ipqshaper*) that performs bandwidth and delays shaping, and introduces packets losses. This tool uses the Application Programming Interface (API) provided by *Netfilter* [12] in order to redirect the packets sent or generated by the applications to a user-space drop-tail queue, where traffic shaping and measurement are performed.

As it is shown in Figure 1, the packets generated by *DSS* are redirected to *ipqshaper*'s queue, where traffic shaping is performed. The outbound packets from *DSS* are sent to the QuickTime players (P_1, \dots, P_n) that are installed on the Windows machine. *iperf* (T_1, \dots, T_n) were installed in order to generate TCP concurrent flows. The queue size was set equal to 10 *KB*, and a 10 *ms* delay was applied to all outgoing packets.

It should be noticed that the experimental testbed is strictly equivalent to the one that could be obtained using a Dummynet-like router [13], the only difference being that in our case we are able to use only two hosts instead of three.

The video flows are generated by using the benchmark video sequence *Foreman*, which is encoded into an MPEG4 format, at a resulting average bitrate

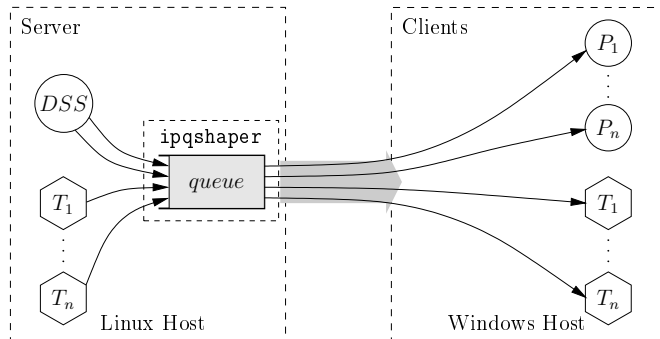


Fig. 1. Experimental testbed

B_{avg} of 242 kb/s. We used the commercial version of QuickTime Player in order to produce the correct *hinted* .mov files, required by the DSS for the streaming.

The experimental scenarios we have tested are made of:

1. Multiple QuickTime concurrent flows with a super imposed constant bandwidth limitation;
2. Multiple QuickTime concurrent flows with a variable bandwidth limitation;
3. Single and Multiple Quicktime flows with a super imposed loss rate;
4. Multiple QuickTime and TCP concurrent flows with a constant bandwidth limitation;
5. Multiple QuickTime and TCP concurrent flows with a variable bandwidth limitation.

The scenario (3) has been realised using the Gilbert model in order to emulate loss events that affect noisy channels (e.g. IEEE 802.11a/b/g connections [14]). This model uses a two-states Markov chain: a *good* and a *bad* state. When the process is into the *good* state, no loss events are generated. When in the *bad* state, the arriving packets are dropped with probability 1. If the transition probabilities are p_{good} and p_{bad} , it can be proved that the expected values are: $l_{good} = \frac{1}{1-p_{good}}$ and $l_{bad} = \frac{1}{1-p_{bad}}$. The values $l_{good} = 11.63$ and $l_{bad} = 1.78$ used here are the ones reported in [14] for IEEE 802.11 connections.

Throughput $r(t)$, loss rate $l(t)$ and goodput $g(t)$ shown in test results are defined as follows:

$$r(t) = \frac{S(t) - S(t - \Delta T)}{\Delta T}, l(t) = \frac{L(t) - L(t - \Delta T)}{\Delta T}, g(t) = r(t) - l(t)$$

where $S(t)$ and $L(t)$ are the number of sent bits and lost bits at time t respectively. We have considered $\Delta T = 0.5$ s in our measurements.

We evaluate the link utilization (LU) as follows:

$$LU = \frac{\sum_{i=1}^N R_i}{C} \cdot 100$$

where R_i is the average throughput of the i -th flow, N is the number of concurrent connections accessing the bottleneck and C is the link capacity.

In order to evaluate the fairness, we employ the Jain Fairness Index (JFI) [15]. Moreover, in order to evaluate the temporal evolution of the fairness index, we define the instantaneous JFI as follows:

$$JFI(t) = \frac{\left(\sum_{i=1}^N g_i(t)\right)^2}{N \cdot \sum_{i=1}^N g_i^2(t)}$$

Where $g_i(t)$ is the goodput of the i -th flow and N is the number of concurrent connections accessing the bottleneck.

5 Results

In this section we present the more significative experimental results we have obtained.

5.1 Darwin Streaming Server evaluation

Two and four QuickTime concurrent flows In this scenario we tested DSS in presence of 2 or 4 concurrent streaming sessions over the same link with super-imposed bandwidth reductions in order to evaluate the fairness of Reliable UDP.

In the case of 2 concurrent QTP streaming sessions we imposed a bandwidth equal to $0.75 \cdot B_{avg} \cdot 2 = 480 \text{ kb/s}$, whereas in the case of 4 concurrent sessions we imposed a bandwidth equal to $0.75 \cdot B_{avg} \cdot 4 = 960 \text{ kb/s}$. Each connection starts after a 10s delay from the other in order to avoid overlaps between the pre-buffering phases.

As it is shown in Figure 2, during the first 10 s of each connection, each flow tends to take up all the available bandwidth in order to fill the playout buffer (over-buffering phase). After this phase, DSS tries to allocate for each flow a bandwidth equal to the average streaming rate B_{avg} , but the first started flows tends to occupy more resources, as shown on Table 1. Moreover, the last started flows experience a higher loss rate, i.e. the adopted congestion control is aggressive and tries to reallocate bandwidth continuously. The link utilization in the case of two QT flows results equal to 79% whereas in the case of four QT flows results equal to 76%, that is, QT flows are unable to fully utilize the link capacity.

As it is shown in Figure 3, the JFIs oscillate in the range [0.5, 0.98] around their average values, meaning that the DSS congestion control becomes unfair with respect to other QT flows when the number of concurrent connections accessing the same bottleneck raises.




Fig. 2. Two and four QuickTime concurrent flows.

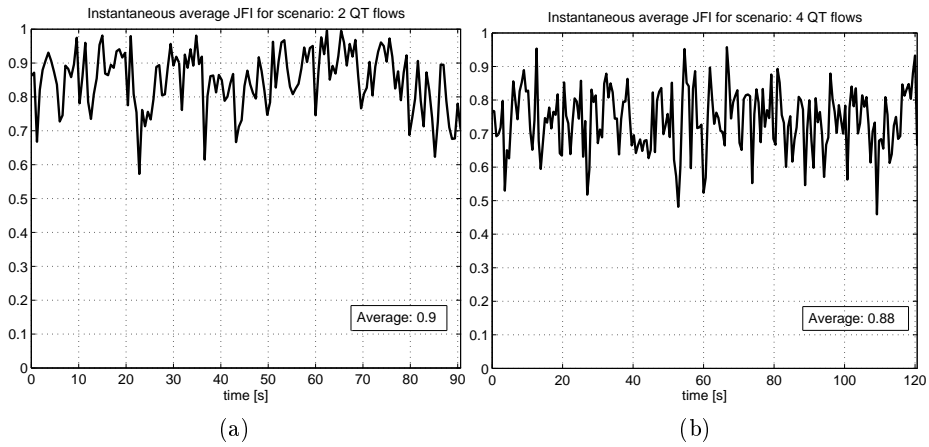


Fig. 3. JFI index for two and four QuickTime concurrent flows.

One and four QuickTime concurrent flows sharing a square wave available bandwidth In this scenario we used an available bandwidth that varies as a square wave with maximum value $A_M = 1600 \cdot N$ kb/s and minimum value $A_m = 100 \cdot N$ kb/s (where $N = 2, 4$ is the number of concurrent flows) with period equal to 40 s in order to investigate how DSS adapts his sending rate.

As it is shown in Figure 4(a), DSS adapts its sending rate in order to match the available bandwidth, using a pre-buffering phase every time a larger bandwidth is available. After each pre-buffering phase (that lasts roughly 10 s), the throughput is back to B_{avg} .

In the case of four QuickTime concurrent flows, we started all flows at the same time. As it is shown in Figure 4(b), when the available bandwidth increases from 100 to 1600 kb/s, each client starts an over-buffering phase. Therefore, in this case all clients equally compete for a bandwidth share. In fact Table 1 shows that the four clients loss rate levels are roughly the same and the JFI is close to 1 (fair share). We can infer that in this case DSS congestion control is more fair than the case described in the previous section, likely because the over-buffering phases start at the same time for each client.

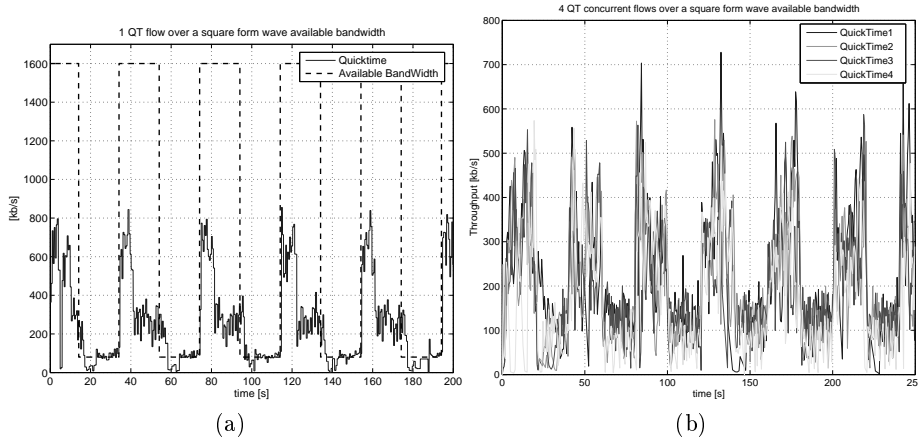


Fig. 4. One and four QuickTime concurrent flows over a square wave available bandwidth.

One QuickTime flow over a lossy link In this scenario we have evaluated the behaviour of DSS and QTP in the presence of a lossy link, emulated using a Gilbert model. Figure 5 shows the throughput and the loss rate imposed by the Gilbert model. As it is shown in Table 1, in this case the throughput is roughly equal to B_{avg} , i.e. the congestion control implemented by DSS exhibits

a fast bandwidth reallocation, likely because Reliable UDP uses a slow-start phase after each loss event.




Fig. 5. One QuickTime flow over a lossy link.

5.2 Darwin Streaming Server vs TCP

Many QuickTime flows with many concurrent TCP flows In this scenario we evaluated the TCP-friendliness of QuickTime streams using more concurrent connections: 1 QT and 3 TCP streams, 2 QT and 2 TCP streams and finally 3 QT and 1 TCP stream. We imposed an available bandwidth equal to $0.75 \cdot B_{avg} \cdot 4 = 960 \text{ kb/s}$. As it is shown from figure 6, QT flows tend to be unfair with respect to TCP flows. By looking at Table 2, in the case of 2 QT and 2 TCP flows, the calculated JFI is low (0.81), instead in the case of 3 QT and 1 TCP flows the JFI is 0.87.

The evaluated link utilization in the case of 1 QT and 3 TCP is equal to 81%, whereas in the other cases results equal to 73%, i.e. TCP flows are able to utilize more bandwidth than the QT flows.

Experiment	Flow	TP (kb/s)	LR (kb/s)	LR%	JFI
2 QT, constant bw	QT1	208.34	18.36	9	0.98
	QT2	171.97	26.83	16	
4 QT, constant bw	QT1	217.81	30.56	14	0.96
	QT2	183.37	37.15	20	
	QT3	178.14	48.21	27	
	QT4	154.34	51.54	33	
1 QT, variable bw	QT	257.67	10.91	4	-
4 QT, variable bw	QT1	212.52	24.96	12	0.99
	QT2	202.49	31.35	15	
	QT3	195.53	36.32	19	
	QT4	163.05	28.73	18	
1 QT, lossy link	QT	203.45	25.70	13	-

Table 1. DSS test results showing average throughputs (TP), lossrates (LR) and Jain Fairness Indexes (JFI).




Fig. 6. Quicktime flows with concurrent TCP flows.

Many QuickTime flow with many concurrent TCP flows over a square wave available bandwidth In this scenario we have evaluated the TCP-friendliness of QuickTime streams in presence of time-varying available bandwidth, that varies as a square wave with maximum value $A_M = 1600 \cdot N$ kb/s and minimum value $A_m = 100 \cdot N$ kb/s (where N is the number of total concurrent flows). Figure 7 shows the test results. In the case of 1 QT and 1 TCP flow, after each over-buffering phase, the QT flow takes a bandwidth roughly equal to B_{avg} , so that the TCP flow (which is a greedy source) can take a larger fraction of bandwidth. On the other hand, in the case of 2 QT and 2 TCP flows, the QT flows always take more bandwidth than TCP ones.




Fig. 7. QuickTime flows with concurrent TCP flows over a square wave available bandwidth.

6 Conclusions

We have carried out an investigation of the Darwin Streaming Server (DSS) in order to evaluate the behaviour of Reliable UDP in the case of time varying network bandwidth and in the presence of multiple concurrent QuickTime and TCP flows.

Main results are: i) when a Reliable UDP-capable client is used, there is an effective adaptation to the network bandwidth; ii) when more concurrent QuickTime flows share the same link, the available bandwidth is not shared fairly, depending on the over-buffering phase adoption by each flow; iii) QuickTime flows are unfriendly with respect to concurrent TCP flows, since they experience a goodput much higher than TCP ones.

Experiment	Flow	TP (kb/s)	LR (kb/s)	LR%	JFI
1 QT + 3 TCP, const bw	QT	250.44	27.09	11	0.96
	TCP1	198.18	24.45	12	
	TCP2	150.50	23.94	16	
	TCP3	179.20	23.61	13	
2 QT + 2 TCP, const bw	QT1	278.82	22.52	8	0.81
	QT2	215.62	31.61	15	
	TCP1	87.70	18.49	21	
	TCP2	121.14	21.92	18	
3 QT + 1 TCP, const bw	QT1	249.55	27.32	11	0.87
	QT2	216.97	36.67	17	
	QT3	144.67	33.40	23	
	TCP	93.76	18.46	20	
1 QT + 1 TCP, variable bw	QT	262.96	15.65	6	0.93
	TCP	453.86	24.17	5	
2 QT + 2 TCP, variable bw	QT1	253.24	20.76	8	0.97
	QT2	252.64	20.28	8	
	TCP1	195.46	24.23	12	
	TCP2	172.51	23.55	14	

Table 2. DSS vs TCP test results showing average throughputs (TP), loss rates (LR) and Jain Fairness Indexes (JFI).

References

1. Van Jacobson and Michael J. Karels, "Congestion avoidance and control," *ACM SIGCOMM Computer Communication Review*, Jan. 1988.
2. Sally Floyd and Kevin Fall, "Promoting the Use of End-to-End Congestion Control in the Internet," *IEEE/ACM Transactions on Networking*, 3 May 1999.
3. L. Ding and R. Goubran, "Assessment of effects of packet loss on speech quality in VoIP," *IEEE, In Proc. HAVE 2003*, pp. 49–54, 2003.
4. T. Hayashi, S. Yamasaki, N. Morita, H. Aida, M. Takeichi, and N. Doi, "Effects of IP packet loss and picture frame reduction on MPEG1 subjective quality," *Multi-media Signal Processing, 1999 IEEE 3rd Workshop on*, pp. 515–520, 1999.
5. M. Handley, S. Floyd, J. Padhye, and J. Widmer, "TCP Friendly Rate Control (TFRC): Protocol Specification," *Proposed standard*, Jan. 2003.
6. Eddie Kohler, Mark Handley, and Sally Floyd, "Designing DCCP: Congestion Control Without Reliability," *Proposed standard*, May 2003.
7. *QuickTime Streaming Server Modules Programming Guide*.
8. 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 Proc. of ACM SIGCOMM IMC '06*, pp. 217–230, 2006.
9. J. Chung, M. Claypool, and Y. Zhu, "Measurement of the Congestion Responsiveness of RealPlayer Streaming Video Over UDP," *In Proc. of the Packet Video Workshop (PV)*, 2003.
10. J. Nichols, M. Claypool, R. Kinicki, and M. Li, "Measurements of the congestion responsiveness of windows streaming media," *In Proc. ACM NOSSDAV '04*, pp. 94–99, 2004.

11. S. Hessler and M. Welzl, "An Empirical Study of the Congestion Response of RealPlayer, Windows MediaPlayer and Quicktime," *In Proc. of IEEE ISCC '05*, pp. 591–596, 2005.
12. Oskar Andreasson, *Iptables Tutorial 1.2.0*. World Wide Web, <http://iptables-tutorial.frozentux.net/iptables-tutorial.html>, 2005.
13. L. Rizzo, "Dummynet: a simple approach to the evaluation of network protocols," *ACM SIGCOMM Computer Communication Review*, vol. 27, no. 1, pp. 31–41, 1997.
14. J. Hartwell and A. Fapojuwo, "Modeling and characterization of frame loss process in IEEE 802.11 wireless local area networks," *In Proc. IEEE VTC '04*, vol. 6, 2004.
15. D. Chiu and R. Jain, "Analysis of the increase and decrease algorithms for congestion avoidance in computer networks," *Computer Networks and ISDN Systems*, vol. 17, no. 1, pp. 1–14, 1989.