Objective and Perceptual Impact of TCP Dropping Policies on VoIP Flows

Erika P. Alvarez-Flores
Centro de Estudios Superiores del Estado de Sonora
Hermosillo, México
erika.alvarez@cesues.edu.mx

Juan J. Ramos-Munoz, Jose M. Lopez-Vega and Juan M. Lopez-Soler
Research Center on Information and Communications Technologies (CITIC), University of Granada
Granada, Spain
jjramos@ugr.es, jmlvega@ugr.es and juanma@ugr.es

Abstract— **The integration of voice and data has allowed the development of new services and applications on the network. Despite of the growth of voice over IP applications, main active queue management (AQM) algorithms primarily focus on TCP traffic. An AQM scheme implicitly adopts a policy for selecting the packet to be marked or dropped. However, the traditional packet selection policies do not take into account the effect on the Quality of Service for different traffic types. In this case, is not always satisfied the quality of service of applications such as Voice over IP (VoIP), which require limited packet delay and loss rate. In this work we study the effect of TCP dropping policy over VoIP for scenarios with mixed traffic conditions. We evaluate their objective and perceptual impact on VoIP flows. As main result, we show that the adoption of one of the proposed AQM dropping procedures improves the user's perceptual score for VoIP, with no penalty on the TCP throughput and loss rate.**

Keywords - Active Queue Management; QoS; VoIP

I. INTRODUCTION

The adoption of Voice over IP (VoIP) technology has allowed the use of the same infrastructures for voice and data, saving installing and managing costs, thus enabling new services. Customer relationship systems (CRM), private branch exchange (PBX), or simply affordable international calls services profit with VoIP. However, IP networks were not designed for carrying multimedia data with real-time requirements. Therefore, the best-effort service that IP networks provides does not fit the VoIP needs.

In particular, one of the most impacting problems of IP for this type of traffic are the network congestion episodes: if routers become congested, the router may start dropping packets. While for TCP traffic it means that packets have to be resent and the transmission rate decreased, for real-time, UDP-based traffic the loss of packets may degrade the quality of the quality perceived by the final users. In addition, congestion may increase the end-to-end packet delay and the average packet jitter, impairments which also affects the VoIP flows quality.

To overcome this limitation and preventing congestion, Active Queue Management (AQM) schemes such as the well known Random Early Detection (RED) [1] have been proposed. AQM schemes monitor the router queue, triggering countermeasures to alleviate the congestion by marking or dropping packets. A TCP packet drop results in the decreasing of the TCP flow throughput, since the TCP congestion algorithm will be initiated. However, UDP flows such as VoIP do not react to these packet losses.

Generally, AQM schemes do not differentiate TCP and UDP traffic. In addition, in shared AQM queues, an inherent problem is that all packets will be exposed to the same drop probability regardless its source pattern. We consider that identifying the responsible source which is causing congestion may alleviate this issue. Unfortunately, up to the author's knowledge, no scalable solution for doing this has been found.

After a number of simulations, we have checked that a good approach for selecting the packet is to drop among both reactive (TCP) and non-reactive (UDP) packets. Otherwise, TCP traffic would be unfairly punished, and consequently non-reactive (UDP) traffic would be incorrectly favoured.

Our approach intends to reach a trade-off that avoids that some flows monopolize the available bandwidth and consequently penalize other flows in active queue management (AQM).

For this goal, in this work we study different dropping strategies for shared AQM queues and evaluate their impact on the VoIP quality of service (QoS). We will experimentally show that if we appropriately select the packet to be dropped, the network level QoS and the VoIP subjective end-user quality will be enhanced. In addition, we show that our approach is TCP friendly. That is, we prove that the VoIP traffic improvements have little impact on the TCP traffic performance.

The remainder of the paper has been organized as follows. In Section II, we briefly describe the AQM RED scheme. Different AQM packet dropping procedures are explained in Section III. The VoIP traffic evaluation framework is detailed in Section IV. Next, in Section V, we report the performance of the studied dropping strategies after some simulations. Finally, the paper is concluded in Section VI.

II. RANDOM EARLY DETECTION AND RELATED WORKS

The Random Early Detection queue management was first described in the seminar paper [1] by Floyd and Jacobson. RED gateways drop or mark each arriving packet with a certain probability, where the exact probability is a function of the average queue size. Its effectiveness is heavily dependent on the setting of its parameters.

Let Avg, Min_{th}, Max_{th} and Max_p be defined respectively as the RED average queue size, minimum threshold, maximum threshold and maximum packet drop probability. As detailed in [1], the packet drop probability is given by

$$P = Max_p \frac{Avg - Min_{th}}{Max_{th} - Min_{th}} \quad (1)$$

RED estimates Avg as the exponentially weighted moving average, expressed as

$$Avg(t+1) = (1 - w_q) Avg(t) + w_q Q \quad (2)$$

in which, $w_q \in [0,1]$ weights the current queue size (denoted by Q), and $(1\text{-}w_q)$ weights the previous long-term average value ($Avg\ (t)$).

RED scheme has motivated a significant number of interesting works. Different aspects have been widely studied in the literature since it was proposed. For instance, [2] and [3] have addressed the stability issue, giving as results the Stabilized RED algorithm (SRED) and improvements in Additive Increase and Multiplicative Decrease/RED systems, respectively. Some AQM schemes such as Balanced RED (BRED) [4] and Fairness-Improvement for RED (FI-RED) [5] have been also proposed to deal with fairness. The settings of RED's parameters have been also addressed, giving as results adaptive schemes such as the Self-Configuring RED [6] or frameworks to find the optimal value of Max_p in RED gateways [7]. Moreover, in [8] alternative packet dropping strategies, as the Drop Front Strategy, have been investigated. Finally, the provision of a better control over the burstiness traffic level has been considered in [9], providing as result the Modified RED (MRED) scheme.

However, the majority of the aforementioned schemes do not take into account the nature of the processed traffic when discarding packets, degrading thus the performance for certain types of traffic.

III. DROPPING STRATEGIES

For victim selection, we use an AQM variant referred to as the Drop-Sel algorithm [10]. Interestingly, Drop-Sel can be integrated into any AQM scheme. Drop-Sel defines three classes of traffic: real-time flows (VoIP), other UDP (O-UDP) and elastic flows (TCP). Drop-Sel records the queue occupancy of each class. When a new packet arrives to the router, and the AQM scheme decides to drop a packet, the Drop-Sel algorithm chooses a packet from the traffic with the highest consumption of memory space in the queue.

The basic idea behind Drop-Sel is to discard the packet belonging to the class of traffic which most contributes to congestion.

Additionally, to improve the QoS perceived by the end user, especially over connections with long propagation delays, it is important to adopt the packet dropping policy with improves the user´s perception.

In this occasion, we evaluate five procedures for TCP packet victim selection. Three of them consider the overall traffic pattern, and the other two are flow based. The policy for the UDP class is always to discard the UDP packet nearest to the front of the queue, whereas for TCP class, the following dropping schemes were considered:

- First-TCP: selects the TCP packet nearest to the front of the queue.
- Last-TCP: selects the TCP packet nearest to the tail of the queue.
- Random-TCP: randomly selects a TCP packet.
- Flow-TCP: selects the packet nearest to the front of the queue belonging to the most populated TCP flow.
- First-Flow-TCP: selects 2 packets. The TCP packet nearest to the front of the queue, and the packet nearest to the front of the queue belonging to the most populated TCP flow. In other words, we combined First-TCP and Flow-TCP. If the algorithm confirms that the selected packets are the same, it considers only one drop.

Fig. 1 shows a packet discarding example for each victim selection procedure.

IV. VOIP QOS ASSESSMENT

For a given voice communication system, the evaluation of the perceptual quality is a costly process, which even could be hardly reproducible. However, for the multimedia communication general case, and for VoIP applications in particular, QoS should somehow definitively include the final perceived user quality.

For quantifying the effect of the aforementioned dropping strategies on the transmission quality, we adopt the E-model and ITU-T Recommendation G.107 [11]. The E-model was initially conceived for network planning design purposes; it predicts the subjective effect of combinations of impairments using stored information on the effects of individual impairments. However, it has also been adopted to estimate the subjective QoS perceived by the user in many voice transmission systems.

For this purpose, the model is usually simplified for the sake of practicality. Henceforth, we adopt the E-model setup

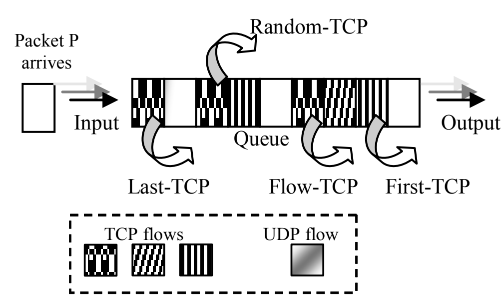

Figure 1. Dropping Strategies Example.

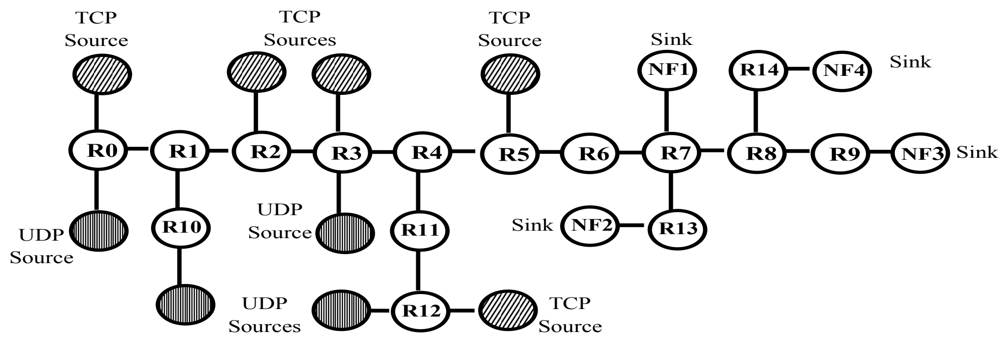

Figure 2. Simulated Topology.

proposed in [12], and obtain the R factor using (3), defined as:

$$R=94.2-0.11(d-177.3)\,H - 0.024d - 30\log(1+15p) \quad (3)$$

where d -expressed in milliseconds- is the end-to-end average delay of the VoIP packets, p is the loss packet probability, and H shapes the delay contribution according to the following equations,

$$\begin{aligned} H &= 0 \; if\,(d - 177.3) < 0 \\ H &= 1 \; if\,(d - 177.3) \geq 0 \end{aligned} \quad (4)$$

To provide more readable subjective evaluations, the R factor can be mapped to MOS score [11]. Like [13], we show the impact of the dropping strategies on VoIP traffic in terms of the MOS scale.

V. Experimental Evaluation

A. Experimental Setup

In order to properly assess both the network-level and the user-level impact of the proposed dropping strategies, we conducted a number of simulations with the ns-2 simulator [14].

Althought Drop-Sel algorithm can be used in conjunction with any AQM scheme, we adopted the RED scheme in our simulations. More precisely, we evaluated the selected AQM RED scheme implementing the five alternative victim selection algorithms described in Section III.

We simulated the reference topology shown in Fig. 2. Links bandwidths are all set to 10 Mbps, except for link R6-R7. To cause network congestion at the AQM node, the R6-R7 link bandwidth is restricted to a narrower bandwidth of 4.5 Mbps. Therefore, R6-R7 is the "bottleneck" link.

The simulated topology represents a complex reference scenario in which a number of elastic TCP and voice over UDP flows compete for the resources of the AQM router.

Since sources are located at different distances from the AQM node, they will undergo a range of different round-trip time (RTT) delays for the TCP flows, and end-to-end delays for the different VoIP flows. The resulting end-to-end delays are close to the maximum allowed delay, 300 ms, so the delay introduced at the router may cause useless expired VoIP packets [15].

TABLE I UDP FLOWS SPECIFICATION

Flow	Source connected	Sink	Interpacket period	Packet size	End-to-end delay
A	R0	NF1	30ms	292 bytes	234 ms
B	R10	NF2	10ms	132 bytes	225 ms
C	R10	NF2	10ms	132 bytes	225 ms
D	R3	NF3	60ms	532 bytes	218 ms
E	R12	NF4	30ms	292 bytes	231 ms

The audio sources generate RTP packets encapsulated into UDP datagrams during random periods, in accordance with the configuration detailed below. The FTP sources also generate the TCP traffic at random variable periods.

The VoIP and FTP applications are modeled as ON/OFF traffic sources. The ON period for the VoIP application lasts 180 seconds, and the OFF period lasts 100 seconds. The FTP traffic follows a Pareto distribution with a shape parameter of k=1.4, an average ON period equal to 2 seconds, and an OFF period that follows an exponential distribution with an average duration of 1 second.

During their ON period, the FTP sources generate TCP segments with a length of 1500 bytes at 1 Mbps. Since different TCP flavours may lead to different results [16], we will evaluate two different implementations. They differ in the ACK algorithms employed by the TCP receiver; the classical ACK (noted hereafter as SACK(1)) and delayed ACK algorithm (called hereafter as SACK(2)) will be tested.

In our setup VoIP sources generate constant bit rate flows that represent voice streams encoded with a G.711 vocodec [17]. To provide a scenario with heterogeneous VoIP applications, the VoIP flows may generate several G.711 frames per packet. This configuration is described in more detail in Table I.

The generated flows go from the node connected to a VoIP or a FTP traffic generator to one of the sink nodes. Specifically, from node R0 to the NF1 sink node, from R3 to NF3, from R12 to NF4, and from nodes R2, R5 and R10 to sink node NF2. Each simulation lasts 500 seconds.

 ISBN: 978-1-61208-215-8

TABLE II LOSS RATE OF AUDIO PACKET BY FLOW

Flow	Metric	SACK(1) Last TCP	SACK(1) Random TCP	SACK(1) First TCP	SACK(1) Flow TCP	SACK(1) First Flow TCP	SACK(2) Last TCP	SACK(2) Random TCP	SACK(2) First TCP	SACK(2) Flow TCP	SACK(2) First Flow TCP
	Drop rate at AQM (%)	1.10	1.08	0.98	1.04	0.66	0.43	0.52	0.56	0.49	0.32
A	Useless packets (%)	24.24	20.21	17.37	16.15	10.43	19.88	16.56	15.33	15.51	11.02
	Total	25.34	21.29	18.35	17.19	11.09	20.31	17.08	15.89	16.00	11.34
	Drop rate at AQM (%)	0.93	1.04	0.92	0.93	0.66	0.50	0.48	0.48	0.55	0.31
B	Useless packets (%)	12.51	7.83	4.72	4.87	2.25	9.65	6.26	4.79	5.29	2.81
	Total	13.44	8.87	5.64	4.80	2.91	10.15	6.74	5.27	5.84	3.12
	Drop rate at AQM (%)	1.54	1.43	1.46	1.38	0.99	0.86	0.68	0.70	0.79	0.47
C	Useless packets (%)	12.28	7.66	4.60	4.77	2.19	9.44	6.23	4.64	5.24	2.77
	Total	13.82	9.09	6.06	6.15	3.18	10.30	6.91	5.34	6.03	3.24
	Drop rate at AQM (%)	1.01	0.88	0.70	0.63	0.50	0.42	0.48	0.35	0.40	0.22
D	Useless packets (%)	4.71	1.20	0.52	0.80	0.17	3.35	1.18	0.72	0.88	0.23
	Total	5.72	2.08	1.22	1.43	0.67	3.77	1.66	1.07	1.28	0.45
	Drop rate at AQM (%)	0.90	0.68	0.67	0.84	0.70	0.42	0.38	0.38	0.38	0.28
E	Useless packets (%)	19.05	14.39	11.39	10.43	6.01	15.52	11.98	10.13	10.35	6.78
	Total	19.95	15.07	12.06	11.27	6.71	15.94	12.36	10.51	10.73	7.06
Average	Drop rate at AQM (%)	1.17	1.13	1.08	1.07	0.78	0.60	0.54	0.54	0.59	0.36
results	Useless packets (%)	14.12	9.61	6.70	6.58	3.51	11.10	7.84	6.37	6.81	4.08
all flows	Total	15.29	10.74	7.78	7.65	4.29	11.70	8.38	6.91	7.40	4.44

B. *Experimental Results*

Given that VoIP quality depends on the loss rate, we evaluated in particular this factor. In our scenario there are two sources of packet losses: firstly, packets are dropped at the AQM router for notifying or preventing congestion; and secondly, useless late packets – those which accumulate an end-to-end delay exceeding 300 ms- are also dropped at the final user node.

To show the effect of the packet victim selection procedures on the packet loss probability, Table II gives the results obtained for the different VoIP flows of the simulated scenario.

It is experimentally shown that in both the drop rate at the AQM router and the drop rate at the final user due to the useless packets, Last-TCP causes the highest packet loss probability at the RED queue. For instance with SACK(1), First-TCP produces an average total loss rate of 7.78% audio packets, while Last-TCP generates 15.29%. Although this degraded performance also occurs with SACK(2), this is less severe. In that case, First-TCP produces an average total loss rate of 6.91% while Last-TCP generates 11.70%

On the other hand, First-TCP, Flow-TCP, and First-Flow-TCP algorithms (with both SACK TCP variants) achieve a significant decrease in the number of useless packets at final user's site, compared to Random-TCP. For example, it is shown for flows A and E, the ones with largest end-to-end delay, that loss rate fluctuates from 11.98% to 20.21% with Random-TCP while it varies from 6.01% to 17.37% with the other three procedures. This result can be explained because of dropping packets from the front generates an empty slot at the head of the queue, and it reduces the delay of all queued packets behind the dropped packet. Therefore, it provides a reduction in the overall end-to-end VoIP delay, making it possible to diminish the number of useless packet.

However, the most significant reduction is obtained with First-Flow-TCP. This algorithm applies a more aggressive dropping policy that causes double packet drops, providing an earlier congestion detection and notification simultaneously to several flows. An interesting point is that First-Flow-TCP controls more effectively the sending rate of the TCP connections with the shortest RTT, generating additional empty slots in the queue at the same time. We see in Table II that this dropping strategy outperforms the other procedures, regardless of how far the sources are from the AQM router. Note that it improves both the drop rate at the router and the number of useless packets percentage. Thus, it causes the lowest loss probability between the evaluated dropping strategies with both SACK TCP variants.

Note that even though sometimes two different TCP flow simultaneously reduce their sending rate, the global TCP performance is not severely degraded. Table III and IV show the TCP loss rate at the AQM router and the TCP arrival rate at the final user, respectively. Note that there is no significant difference between the results obtained by the discarding procedures for the same TCP flavour.

If we reduce the packet loss at the router and at the final user side, we could expect a significant improvement in the subjective end-user perceived quality of VoIP. Fig. 3(a) and 3(b) show the MOS values obtained for each flow with different packet victim selection procedures.

As it can be observed, First-Flow-TCP achieves the best results for VoIP traffic. For all flows, the highest MOS values are obtained using the First-Flow-TCP procedure. This means that using a more aggressive dropping policy with TCP traffic, better QoS is provided for VoIP traffics.

On the other hand, without imposing an aggressive dropping rate, the experiments carried out show that the discarding of packets nearest to the front of the queue achieves significantly improvement in MOS values.

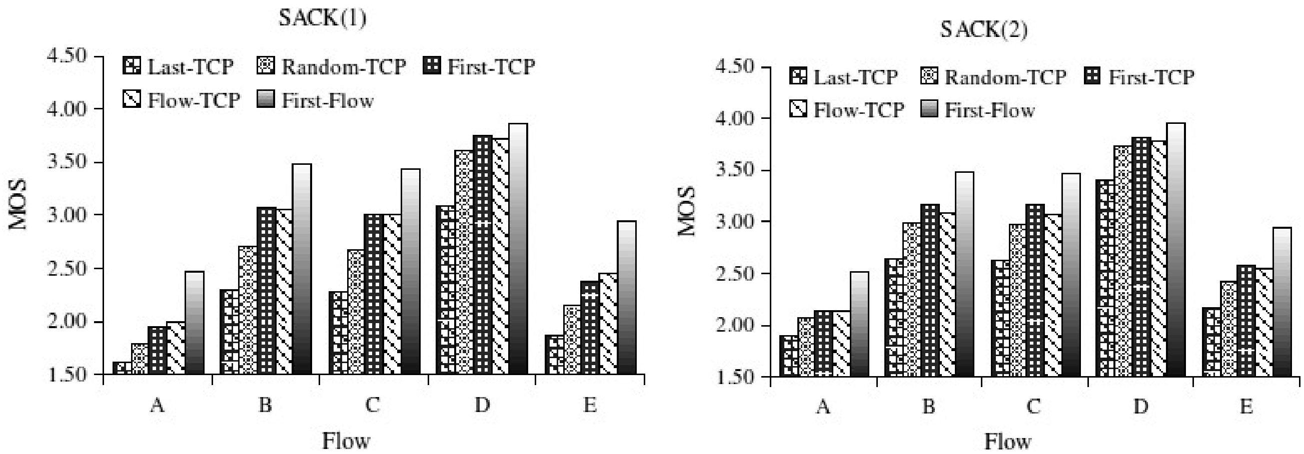

(a) (b)

Figure 3. E-MODEL based MOS evaluation of dropping policies.

TABLE III LOSS RATE OF TCP PACKET (%)

	Last TCP	Random TCP	First TCP	Flow TCP	First Flow
SACK(1)	4.00	3.82	3.84	3.36	3.64
SACK(2)	2.06	1.95	1.94	1.84	1.94

TABLE IV ARRIVAL RATE OF TCP PACKET (MBPS)

	Last TCP	Random TCP	First TCP	Flow TCP	First Flow
SACK(1)	4.098	4.110	4.117	4.132	4.099
SACK(2)	3.879	3.891	3.924	3.985	3.861

As expected, MOS results obtained for the different procedures have a rational correlation with network-level results.

VI. CONCLUSION AND FUTURE WORK

In this work, we have studied different TCP packets dropping policies for the Drop-Sel victim selection scheme. We have assessed their impact on the quality of service perceived by end users. We have evaluated their performance by means of simulation, obtaining network parameters and perceptual scores which experimentally demonstrate the benefits of using specific policies for TCP packets. Such benefit results in an enhanced perceived quality of the VoIP flows, without degrading the TCP flows performance.

The experiments also show that despite of using an aggressive discard policy for TCP, as the First-Flow-TCP procedure does, discarding up to two packets at once, the TCP performance is not significantly penalized. Those results have been validated for two SACK TCP variants.

As future work, we plan to evaluate the impact of AQM schemes on VoIP and other real-time media flows such as video streaming. For instance, the use of TCP and HTTP as transport protocols for streaming media with time constraints [18] will be also considered. Additionally, we will study Drop-Sel based algorithms to cope with that class of TCP traffic.

ACKNOWLEDGMENT

This work was supported by the "Ministerio de Ciencia e Innovación" of Spain under research project TIN2010-20323.

REFERENCES

[1] S. Floyd and V. Jacobson, “Random early detection gateways for congestion avoindance,” IEEE/ACM Transaction on Networking, August, 1993. doi: 10.1109/90.251892.

[2] T. J. Ott, T. V. Lakshman, and L. Wong, “SRED: Stabilized RED,” IEEE, 1999, pp. 1346-1355. doi: 10.1109/INFCOM.1999.752153.

[3] L. Wang , L. Cai , X. Liu, and X. Shen, “Stability and TCP-friendliness of AIMD/RED system with feeback delays,“ Computer Networks 51, 2007, pp. 4475-4491. doi: 10.1016/j.comnet.2007.06.022.

[4] F. M. Anjum and L. Tassiulas, “Fair bandwidth sharing among adaptive and non-adaptive flows in the internet,” IEEE, 1999, pp. 1412-1420. doi: 10.1109/INFCOM.1999.752161.

[5] H. Ohsaki, T. Eguchi, and M. Murata, “FI-RED: AQM Mechanism for improving fairness among TCP connections in tandem networks,” Proc. IEEE SAINT ’05, 2005. doi: 10.1109/SAINT.2006.33.

[6] W. Feng, D. Kandlur, D. Saha, and G. Kang, “Self-configuring RED gateway,” Proc. IEEE INFOCOM, 1999. pp. 1320-1328. doi: 10.1109/INFCOM.1999.752150.

[7] B. Zheng and M. Atiquzzaman, “A framework to determine bounds of maximum loss rate parameter of RED queue for next generation routers,” J. Network and Computer Applications, 2008. doi: 10.1016/j.jnca.2008.02.003.

[8] T. V. Lakshman, A. Neidhardt, and T. J. Ott, “The drop front strategy in TCP and in TCP over ATM,” IEEE, 1996. doi: 10.1109/INFCOM.1996.493070.

[9] G. Feng, A. Agarwal, A. Jayaraman, and C. Siew, “ Modified RED gateways under bursty traffic,” IEEE Commun Lett, 2004, pp. 323-325. doi: 10.1109/LCOMM.2004.827427.

[10] E. P. Alvarez-Flores, J. J. Ramos-Munoz, P. Ameigeiras, and J. M. López-Soler, “Selective packet dropping for VoIP and TCP flows”, Telecommunication System Journal, Vol. 46 No.1, 2011, pp. 1-16. doi: 10.1007/s11235-009-9252-z

[11] ITU-T, Recommendation G-107, “The E-model, a computational model for use in transmission planning,” March 2005.

[12] R. G. Cole and J. H. Rosenbluth, “Voice over IP performance monitoring,” Proc. SIGCOMM Comput. Commun. Rev. 31, 2001, pp. 269-275. doi: 10.1145/505666.505669.

[13] V. A. Reguera, F. F. Alvarez Paliza, W. Godoy Jr., and E. M. García Fernández, “On the impact of active queue management on VoIP quality of service,” Computer Communications 31, 2008 pp. 73-87. doi: 10.1016/j.comcom.2007.10.016.

[14] Network Simulator ns2. [Online]. Available from: http://www.isi.edu/nsnam/ns/.

[15] C. Perkins, O. Hodson, and V. Hardman, “A survey of packet-loss recovery techniques for streaming audio,” Proc. IEEE Network, 1998, pp. 40-48. doi: 10.1109/65.730750 .

[16] M. Allman, V. Paxson, and W. Stevens, “TCP Congestion Control”, IETF RFC 2581, April, 1999.

[17] ITU-T, Recommendation G.711, “Pulse code modulation (PCM) of voice frequencies”.

[18] R. Pantos, "HTTP Live Streaming", IETF Internet-Draft, September 30, 2011.