

Selective packet dropping for VoIP and TCP flows

Erika P. Alvarez-Flores · Juan J. Ramos-Munoz ·
Pablo Ameigeiras · Juan M. Lopez-Soler

Published online: 19 December 2009
© Springer Science+Business Media, LLC 2009

Abstract The adoption of the IP protocol for serving diverse applications arises the need for mechanisms to prevent network congestion in scenarios with different traffic types (responsive and unresponsive) sharing limited network resources. To deal with this issue, a number of algorithms for active queue management (AQM) have been proposed. However, most of them do not observe the traffic type and usually disregard this knowledge. In this way, the provided service could not comply with the distinctive requirements of the different type of traffic, such as VoIP services, which demand bounded packet latency and loss rate.

This paper proposes a new approach to be applied for preventing network congestion in AQM routers. Our scheme includes a procedure for selecting the packet to be dropped that improves the fairness among different classes of flows. We evaluate the use of this approach on distinct AQM schemes in scenarios with different degrees of UDP and TCP traffic mix. Objective and subjective performance measurements are reported. The experimental evaluation indicates that our approach improves the fairness among different traffic classes without using any packet scheduler. In fact, it also improves the VoIP traffic performance in terms of packet dropping probability, MOS (Mean Opinion Score) and intelligibility. We also show that our approach has no negative impact on the packet delay. Moreover, it is not achieved at the expense of TCP responsive traffic.

Keywords Active queue management · RED · VoIP · TCP-friendly · Fairness · Traffic mix

1 Introduction

Almost universally accepted for short and medium term, the IP protocol is considered to be the convergence technology for interconnecting networks. All kinds of applications, ranging from pervasive computing and multimedia to more traditional ones, are envisaged as IP users. In consequence, IP-based convergence technology supports a wide variety of service types, with very different traffic characteristics and quality requirements. Most of these applications are based either on the TCP protocol, which reacts to packet losses to avoid network congestion, or on the UDP protocol, which does not react to packet losses at all. This large variety of traffic flows contending for the same network resources imposes interesting challenges on the conveying networks.

One of the critical challenges facing IP-based technologies is the question of preventing network congestion. To deal with this problem, a number of algorithms have been described, most of them relying on the Random Early Detection (RED) [1] queue management approach. RED gateways drop or mark each arriving packet with a certain probability, where the exact probability is a function of the average queue size.

In general terms, AQM schemes prevent congestion by dropping or marking packets at the output queue in the designated router. However, these schemes typically disregard the traffic type information of the packets in the output queue. The present paper proposes a new approach to be applied in AQM schemes that considers the traffic type in selecting the packet to be dropped. We assume that a single

This work has been partially supported by the Spanish Science and Innovations Ministry (MICINN) Project TIN2009-13992-C02-02, and by the University of Granada (SPAIN)—Plan Propio 2009.

E.P. Alvarez-Flores · J.J. Ramos-Munoz · P. Ameigeiras ·
J.M. Lopez-Soler (✉)
ETSI Informática y de Telecomunicación, 18071 Granada, Spain
e-mail: juanma@ugr.es

output AQM queue is used for both responsive and unresponsive traffic. As explained below, the scheme described observes the network traffic in the queue and, in accordance, selects the packet to be dropped from the traffic class with the highest router memory consumption.

The designed algorithms have been evaluated by means of network simulations in various scenarios with different degrees of UDP and TCP traffic mix.

For VoIP traffic we assess the performance in terms of Quality of Service (QoS) as well as the associated Quality of Experience (QoE). The latter performance metric is evaluated bearing in mind that merely providing excellent low-level QoS is useless if this is not correlated with the (high-level) end user's satisfaction. Like [2], we report the impact of the proposed scheme for VoIP traffic on the Mean Opinion Score (MOS) scale. To complete the evaluation, in some illustrative experiments, we estimate the QoE in terms of intelligibility. For this purpose, we additionally provide scores of an automatic speech recognizer located at the receiver's premises.

The experimental evaluation indicates that the proposed algorithm significantly improves the fairness among the distinct traffic classes without needing any packet scheduler. Moreover, it improves the VoIP traffic performance in terms of packet dropping probability, MOS and intelligibility without penalizing responsive traffic.

The remainder of this paper is structured as follows: Sect. 2 provides a concise overview of the most relevant AQM approaches. Section 3 describes the proposed selective packet dropping algorithm. In the light of the simulations made, Sect. 4 assesses the performance of the proposed approach. We provide both network-level QoS measurements (in terms of fairness, packet delay and loss probability) and user-level VoIP QoE evaluation (MOS scores). In this section, intelligibility test based on Automatic Speech Recognition (ASR) scores are provided for comparison with the reference system. Finally, the main conclusions are summarized in Sect. 5.

2 Active queue management

Active Queue Management is a class of packet queuing algorithms that is intended to prevent network congestion. It uses a probabilistic approach for reacting to congestion conditions. Previous works on congestion avoidance gateways like the Early Random Drop (ERD) [3] and DECbit schemes [4] preceded the seminal paper of Random Early Detection algorithm (RED) [1]. RED gateways drop or mark each arriving packet with a certain probability, whereby the exact probability is a monotonically increasing function of the average queue size. The main drawbacks of this scheme are that under some circumstances RED could react slowly,

and that its effectiveness is heavily dependent on the setting of its parameters. Adaptive approaches have been proposed to cope with these issues, such as the Adaptive RED (ARED) [5] and the Proportional-Differential RED (PD-RED) [6]. ARED increases or decreases the RED dropping probability in accordance with the traffic workload, to keep the average queue occupancy within an objective range. PD-RED uses a Proportional-Differential controller to adjust the dropping probability according to the average error and the target queue sizes. Nevertheless, since these methods only consider queue occupancy, they do not maximize the use of the output link capacity.

Besides RED-based approaches, there exist a number of more recent systems that also calculate the dropping probabilities as a function of the queue length, like the Fast and Autonomic Fuzzy AQM Controller [7]. This scheme includes a robust fuzzy logic algorithm and a self-configuring mechanism for lowering the queue occupancy and, accordingly for reducing the packets end-to-end delay.

Two other schemes that detect the congestion by monitoring the incoming traffic or the loss rate are the Adaptive Virtual Queue (AVQ) [8] and GREEN [9]. The operation of the AVQ scheme is based on a virtual queue whose capacity depends on the arriving packets rate. Alternatively, GREEN adjusts the congestion notification rate based on a estimation of the incoming rate and the average output capacity.

Other AQM approaches calculate the packet dropping probability as a function of not only the queue size, but the arriving traffic rate or the loss ratio metrics. For example, Random Exponential Marking (REM) [10] and RaQ [11]. The REM system makes use of an exponential probability to adjust the queue length. It uses a price function which is calculated as function of the queue size and the incoming packet rate. It stabilizes the incoming rate to the link capacity, and thus the queue size is consequently reduced. RaQ stabilizes the queue length by using control theory concepts. Basically, it uses a dual loop feedback control system that uses Proportional rate and Proportional-Integral queue length controls.

In general, the aforementioned AQM schemes do not take into account the responsive behavior of the involved traffic sources. If responsive and unresponsive flows (e.g., TCP and UDP sources) are mixed in a single queue, the mentioned AQM schemes undistinguishing react to congestion, regardless the responsive nature of the involved traffic. It is expected that the problem would be alleviated if the algorithm considers the different traffic classes. In accordance with this idea, low complexity algorithms that use per flow-state information have been proposed. For instance, Flow Random Early Drop (FRED, [12]) and Approximate Fair Dropping (AFD, [13]). Additionally, the Dynamic Class Based Thresholds (D-CBT, [14]) method prevents congestion by using flow classification and applying distinct policies according to the given flow class.

There do exist, moreover, algorithms that solve the monopolization of bandwidth without considering per-flow-state information. Core-Stateless Fair Queuing (CSFQ, [15]) and Rainbow Fair Queuing (RFQ, [16]) are representative algorithms in this group. They both provide a high level of fairness, but they need the collaboration of core and edge routers. To this end, the edge routers label each packet with information that is used at the core routers to calculate the dropping probability. Despite of both provide flow level fairness, they cannot satisfy the real-time requirements of some flows, as demanded by interactive VoIP streams. To overcome this drawback, the Jitter Detection algorithm (JD, [17]) categorizes the packets into two classes, TCP and UDP. The TCP class is handled by the RED algorithm and the UDP class is managed under the JD scheme. The latter scheme calculates the UDP dropping probability as an exponential smoothing function of the packet delay. The RED-Worcester scheme [18] uses an exponential smoothing of the instantaneous values of the suggested delays to adjust the dropping probability. In this way, it adapts the average queue occupancy to a target range, taking into account each packet's requirements. Alternatively, a novel modification to CHOKe (CHOOse and keep for responsive flows, CHOOse and kill for unresponsive flows, [19]), called PUNSI has been proposed [20]. PUNSI computes the dropping probability of TCP packets similarly to RED, and uses additional policies when the incoming packet belongs to an UDP flow. In spite of this differentiated treatment, the policy applied to benefit TCP streams strongly penalizes the UDP flows.

3 The Drop-Sel victim selection algorithm

Whenever an AQM scheme determines that a packet has to be marked or dropped, it executes a victim selection algorithm to choose the packet to mark or discard. For the sake of simplicity, in this work we will assume that the adopted policy is to drop packets instead of marking.

The selection algorithms are as simple as selecting the last, the first, or a random packet from the queue. Those algorithms will be referred hereafter to as Drop-Tail, Drop-Front and Drop-Random, respectively. Most AQM schemes use this kind of mechanisms to select the packet to be dropped. However, this decision may lead to an unfair traffic treatment-providing unbalanced loss rates- and possibly, it may underutilize the network resources.

Additionally, another matter of concern for VoIP flows is the average packet queuing delay. If the end-to-end delay of a VoIP packet is greater than 300 ms, it is considered a loss. Moreover, the quality of a VoIP flow depends heavily on the average packet delay [21], and therefore it is crucial to minimize it. In this respect, previous research has shown that the Drop-Front scheme is better at reducing the average queuing delay [22, 23]. Despite this benefit, however,

the selection scheme does not discriminate between flows, and as a result, a class of traffic may have a higher chance to be selected and be forced to have less than its fair share of bandwidth. A more effective victim selection algorithm can be envisaged if the selection is made regarding the class of traffic to which the packet belongs. In doing so, better performance for all the involved traffic classes can be hopefully obtained without needing the use of any extra scheduler.

We propose a victim selection algorithm, referred to hereafter as Drop-Sel. Our aim is to improve the fairness and to provide, if possible, an overall better network utilization.

Drop-Sel avoids a packet scheduler that uses multiple output queues, using a same single queue for every type of traffic. The main purpose behind Drop-Sel is to provide a fair service for responsive and unresponsive sources that whenever it would be feasible will improve the QoS and QoE of VoIP applications.

Given that the packet flows with the highest workload are the main contributors to the congestion episode, our scheme simply identifies packets belonging to such traffic class, and singles them out in order to protect the other flows. Thus, this algorithm penalizes the sources with the highest queue occupancy, assigning them a higher probability of being discarded. As a result, a differentiated treatment for each packet based on their traffic class is provided.

In this work, similarly to [14], we will consider three classes of traffic: interactive non-responsive UDP real-time class (VoIP flows), other non real-time UDP flows, and finally the TCP class.

To provide a fair dispatching service, Drop-Sel maintains, for each traffic class, a count of the enqueued packets specified in bytes. Drop-Sel will select the packet nearest to the front of the queue from the traffic class with the highest count. Therefore, when a packet arrives at the queue, it is classified as real-time (VoIP), other UDP (O-UDP hereafter), or elastic (TCP), and the counter is updated accordingly. Concurrently, the packet is enqueued or dropped with a probability that depends on the AQM discipline. If the AQM algorithm decides to drop the packet, it executes the victim selection algorithm. To do this, the real-time, other UDP and TCP counters are compared in order to determine the class with the highest queue occupancy. Then, the algorithm looks for the packet within the selected class that is nearest to the front of the queue, and drops it. In Fig. 1, the Drop-Sel procedure is summarized in a simple flow-chart.

Drop-Sel preferentially accepts packets from the traffic classes with the lowest consumption of memory space in the queue during congestion periods. Consequently, as it will be shown, the algorithm provides better chance to the traffic classes which consume less resources.

Specifically, the advantages of Drop-Sel are the following. First, it only considers the instantaneous queue occupancy in a per-class based approach instead of per-flow

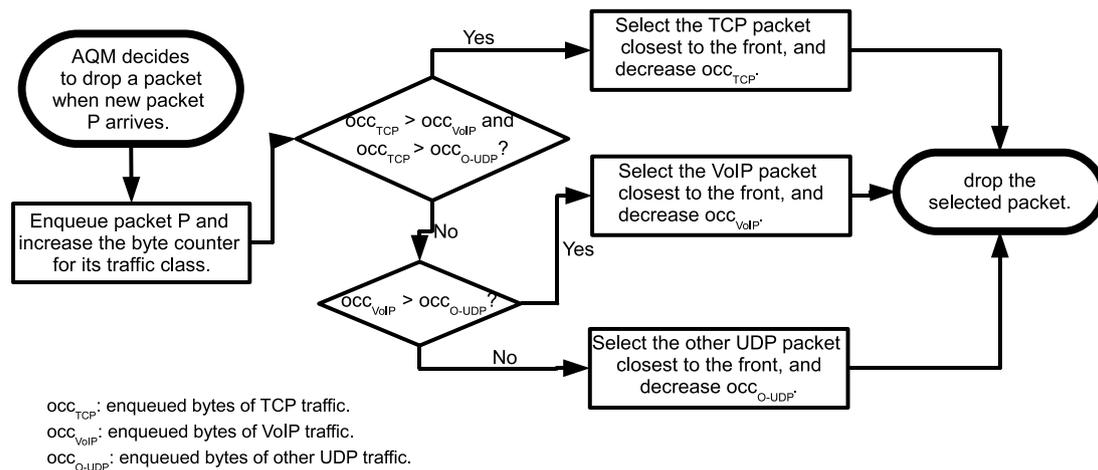


Fig. 1 Drop-sel flow diagram

based. In this sense, Drop-Sel is scalable. Second, for distinct traffic classes, it uses a single queue, thus it does not require an scheduler to provide fairness. Additionally, even if some classes do not require to use their fair share of the networks resources (memory and bandwidth), Drop-Sel will fairly give the unused extra resources to the other traffic classes. Third, the sources do not need to insert extra signaling information into the packet headers.

4 Drop-Sel experimental evaluation

To compare the performance of the victim selection algorithms described in Sect. 3, a number of simulations were conducted with the ns-2 simulator [24]. A set of selected AQM schemes were evaluated under a variety of network topologies, traffic sources, and different congestion levels as well.

To properly assess the benefits provided by the proposed scheme to VoIP flows, for this traffic type we measured both network-level QoS parameters (the average packet delay and loss rate), and user-level QoS perceptual scores.

4.1 Experimental setup

To carry out a comprehensive experimental evaluation, we have simulated a total of three scenarios. Each scenario (namely S1, S2, and S3) differ from another in the topology adopted and the traffic workload assigned for each traffic class.

For scenarios S1 and S2, we consider the standard single-bottleneck dumbbell topology shown in Fig. 2. In this scheme, a number of TCP, VoIP, and other non-voice UDP flows compete for the shared resources of the AQM router (R0). A number of FTP sources generate the TCP segments

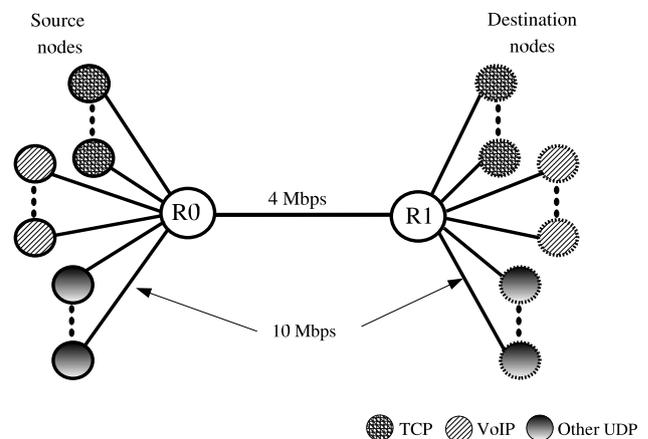


Fig. 2 Dumbbell topology for scenarios S1 and S2

with a length equal to 1500 bytes, and the VoIP sources generate RTP packets [25] (encapsulated into UDP datagrams) that stand for G.711 encoded voice [26]. The size of non-voice UDP datagrams is also equal to 1500 bytes.

The simulated VoIP flows are generated using three VoIP applications (A, B, C) with different interpacket periods and packet sizes (see Table 1). To consider the impact of the end-to-end latency, in scenario S1, all the VoIP applications have equal end-to-end latency (39 ms), while in scenario S2, the end-to-end delays range from 39 to 269 ms.

In scenario S1, the router R0 is stressed with four different workloads, summarized in Table 2 (cases CA1, CA2, CA3, and CA4). For scenario S2, in which different delays were considered, we simulate four extra cases (labeled as CA5, CA6, CA7, and CA8 in Table 3). In both scenarios (S1 and S2), R0 is supposed to have a Drop-Tail RED AQM scheme.

In the experiments of scenario S1, all the traffic sources start at instant 0 s, and are active up to the end of the sim-

Table 1 VoIP flows specification for dumbbell and complex topologies

Application	Interpacket period	Packet size
A	10 ms	92 bytes
B	30 ms	252 bytes
C	60 ms	492 bytes

Table 2 Workloads of UDP and TCP traffic flows for scenario S1

Case	VoIP	O-UDP	Total UDP	TCP
CA1	25%	50%	75%	25%
CA2	50%	25%	75%	25%
CA3	25%	25%	50%	50%
CA4	12%	12%	24%	76%

Table 3 Workloads of UDP and TCP traffic flows for scenario S2

Case	VoIP	O-UDP	Total UDP	TCP
CA5	47%	36%	83%	17%
CA6	54%	19%	73%	27%
CA7	13%	52%	65%	35%
CA8	29%	6%	35%	65%

ulation (500 s). On the contrary, in S2, all the sources were modeled using ON/OFF traffic patterns. The ON period for the VoIP application lasts 180 seconds, while the OFF period lasts 100 seconds [27]. 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 [28]. Finally, for the other UDP applications the ON period lasts 300 seconds, and the OFF period lasts 200 seconds. Each traffic generator starts sending packets with a uniform random probability during the first 15 seconds of the simulation. In this way, assuming that the different flows are not synchronized, a number of congestion periods are randomly generated. To complete the evaluation, we define a third scenario (S3), in which we also consider a more realistic topology with multiple links. The simulated topology, shown in Fig. 3, consists of a total of 7 routers (labeled as R0 through R6) and 4 sink nodes (labeled as NF1 through NF4). The links bandwidths are all set to 10 Mbps, with the exception of links R2-R3 and R5-R6: to cause network congestion at the AQM nodes (R2 and R5), the R2-R3 link bandwidth is restricted to 3.5 Mbps and the bandwidth of the R5-R6 link is set equal to 4.5 Mbps. Therefore, R2-R3 and R5-R6 are the “bottleneck” links (L1 and L2 in Fig. 3). Since the remaining links have enough capacity to cope with the generated traffic, no extra queuing delays will be introduced.

Table 4 Workloads of UDP and TCP traffic flows for scenario S3

Case	VoIP	O-UDP	Total UDP	TCP
CA9	12%	39%	51%	49%
CA10	12%	21%	33%	67%

A set of TCP and UDP sources were also arranged for this topology. In particular, for S3 we adopt the workloads summarized in Table 4. The VoIP, FTP, and other UDP applications are modeled as ON/OFF traffic sources (as in scenario S2). The number of sources is set according to the relative distribution analyzed in [29].

The simulated flows go from the sources connected to a router towards one of the sink nodes. More specifically, from router R0 to NF1, from R1, R2, and R5 to NF2, from R3 to NF3, and from R4 to NF4.

4.2 Drop-Sel fairness evaluation

In this section we evaluate and discuss the impact of including the Drop-Sel victim selection procedure in RED, AVQ, REM. In particular, we examine its impact on the average throughput for the considered traffic classes (TCP, VoIP and other competing UDP) under different workload conditions. As main result we anticipate that Drop-Sel improves the overall fairness among the traffic classes even when the traffic load is mostly generated by unresponsive sources (VoIP and UDP) without punishing TCP responsive sources.

First, we show in Fig. 4 the average input and output throughput for each of the simulated victim selection algorithms at R0, for the RED scheme and scenario S1 (cases CA1, CA2, CA3 and CA4 of Table 2). We can observe clearly that, in case CA1 (Fig. 4a), it is shown that Drop-Tail and Drop-Front schemes are not fair. On the contrary, Drop-Sel tends to make equal the output throughput for the three traffic classes regardless of the responsive nature of the contending sources. Similarly, for case CA2 in scenario S1 (Table 2) in which TCP and O-UDP flows compete with predominant VoIP flows, we observe that Drop-Sel equates the VoIP bandwidth consumption (34%) to TCP (35%) and O-UDP (31%) traffic (Fig. 4b). In addition, for cases CA3 and CA4 (Fig. 4c and 4d) in which TCP sources respectively generate the 50% and 76% of the traffic load, Drop-Sel does not significantly penalize the TCP sources. It can be explained because of the responsive nature of the TCP sources. For TCP sources, when Drop-Sel drops a packet, the sending rate adapts to the available bandwidth. This way, the presence of TCP packets in the queue further decreases and, consequently, Drop-Sel reduces the probability of selecting a TCP packet as a victim.

For scenario S2, in which a dumbbell topology with different end-to-end delays is considered, Drop-Sel also improves the fairness between the traffic classes. To show this

Fig. 3 Complex topology for scenario S3

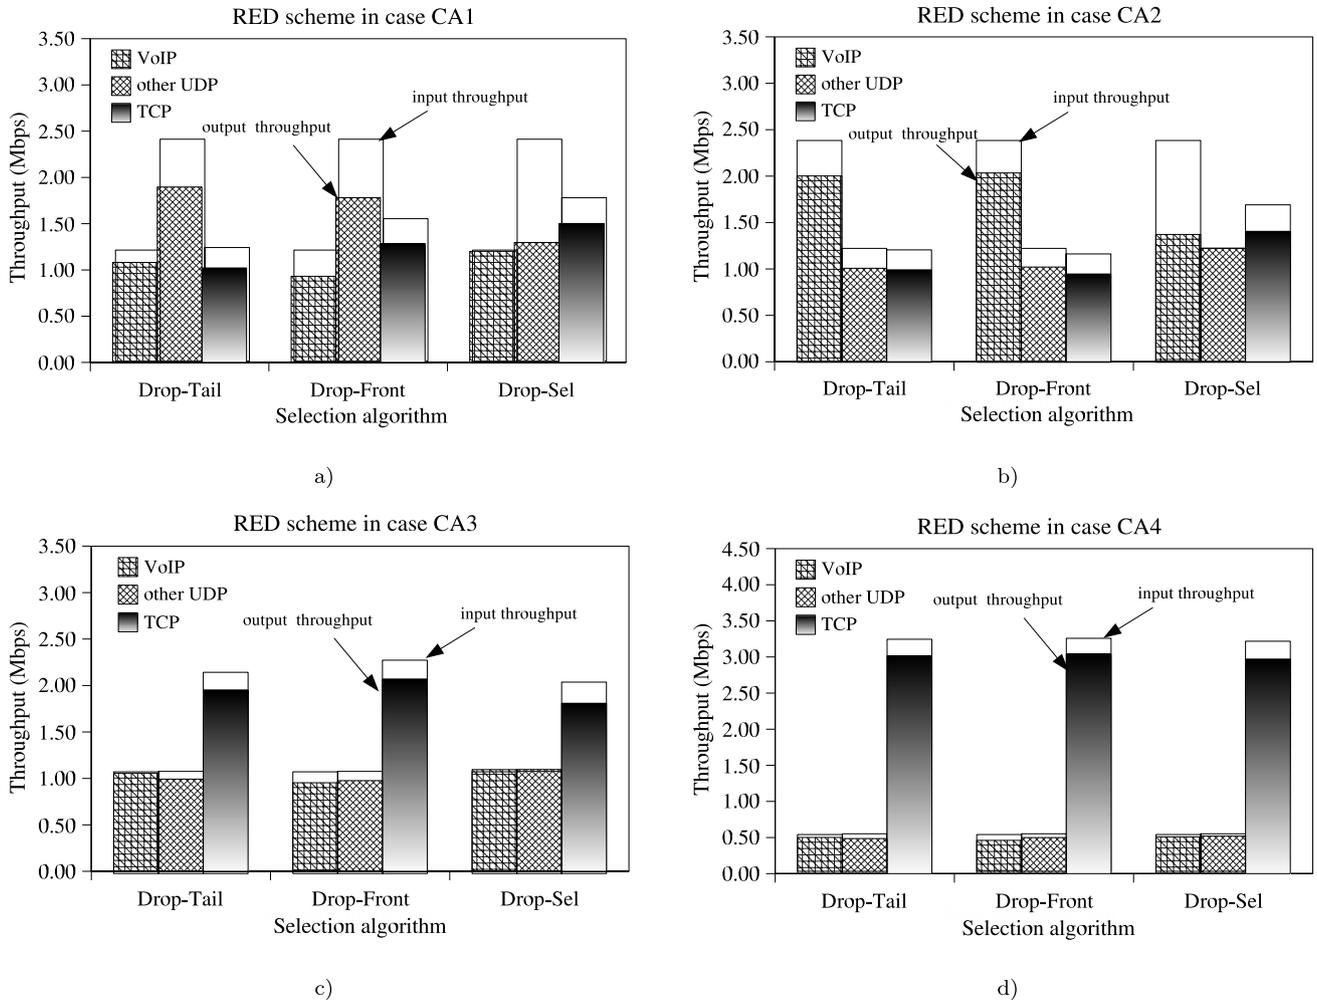
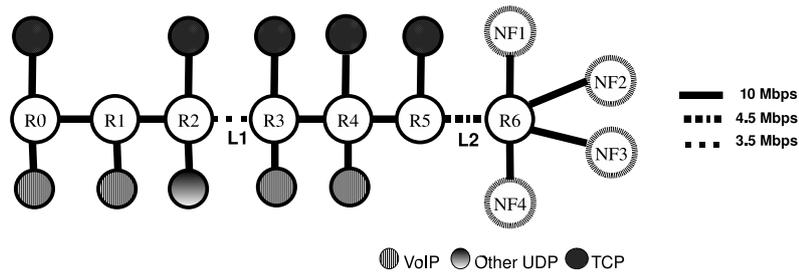


Fig. 4 Average throughput for victim selection algorithms in S1, for cases CA1, CA2, CA3 and CA4

we have considered the impact on the average performance when RED is adopted. These results are depicted in Figs. 5a, 5b, 5c, and 5d (cases CA5, CA6, CA7, and CA8 of Table 3).

As a first result, it can be observed that the TCP throughput obtained with Drop-Sel is higher or equal than the resulting throughput of the other schemes (Fig. 5). This occurs for cases CA5, CA6, and CA7, because TCP is not severely penalized since the UDP or VoIP packets occupy the majority of the queue. Consequently, TCP packets are discarded less often, the TCP sources do not notice the congestion, and

they do not drop their sending rate. When TCP dominates (case CA8), it can be checked that VoIP and O-UDP obtains the bandwidth that they require, without harming the TCP traffic, since their throughput is lower than the proportional share (Fig. 5d).

Note that, since the packets are generated with a random pattern, the output throughput does not achieve the exact distribution between the traffic sources. Nevertheless, note in Fig. 5 that Drop-Sel results in a more equitable use of the bandwidth than Drop-Tail and Drop-Front.

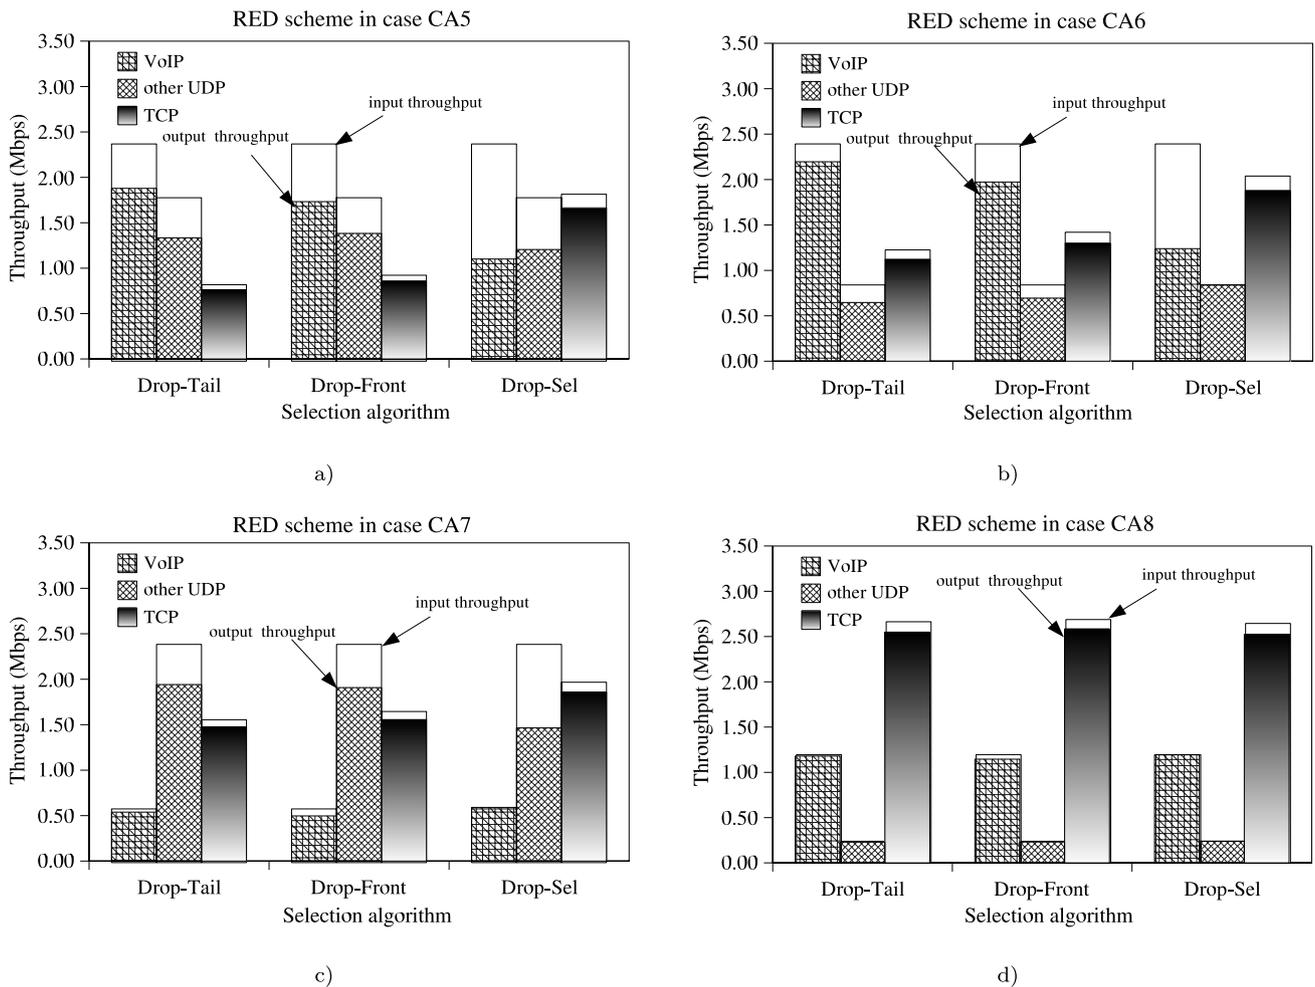


Fig. 5 Average throughput for victim selection algorithms in scenario S2, for cases CA5, CA6, CA7 and CA8

From these results we can conclude that, with the use of Drop-Sel, the TCP traffic is not significantly penalized and the VoIP and O-UDP traffic approaches an equitable share of the throughput.

Drop-Sel can also improve the fairness in other AQM schemes. To provide some insights, we consider for example the impact on the average throughput in AVQ and REM. For instance, Fig. 6a corresponds to case CA6 of scenario S2 (see Table 3) in which the generated VoIP traffic load dominates (54%), while O-UDP and TCP sources respectively generate the 19% and 27% of the overall load. In this case, when AVQ adopts Drop-Tail scheme, the VoIP traffic approximately gets a double bandwidth share compared to TCP. However, Drop-Sel respectively provides 40% and 39% of the bandwidth for TCP and VoIP flows, while O-UDP sources get what they need (21%). Similarly, Drop-Sel equitable behavior is also hold for REM scheme, as shown in Fig. 6b for case CA7 in which other UDP traffic load dominates.

As conclusion, the obtained results indicate that Drop-Sel is TCP-friendly since it improves the fairness for RED, AVQ and REM AQM schemes, when different traffic sources compete for the same network resources without noticeable impact on responsive sources even under traffic conditions for which TCP sources dominate.

To finish the Drop-Sel fairness evaluation, we now consider the PUNSI AQM scheme [20]. PUNSI was specifically designed for dealing with unresponsive flows. It prevents that unresponsive flows monopolize the available bandwidth and harm responsive sources. PUNSI penalizes unresponsive flows with a probability higher than those from responsive sources in accordance with its burstiness.

We simulate the PUNSI scheme for case CA2 in scenario S1 (Table 2). In this case, because PUNSI only considers two different classes, VoIP and O-UDP flows are aggregated into one traffic class. In Fig. 7a we plot the PUNSI input and output throughputs, and for comparison purposes we also depict results from RED with Drop-Sel and Drop-Tail. Note that, since PUNSI does not differentiate between types

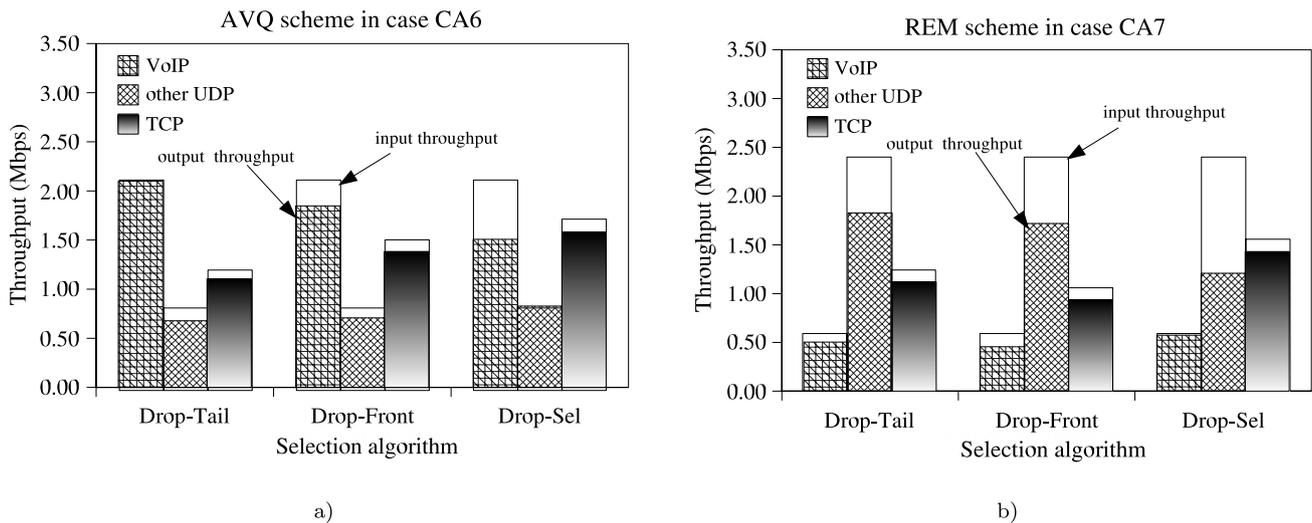


Fig. 6 Average throughput for victim selection algorithms in scenario S2, for AVQ and REM schemes

of UDP packets, the O-UDP traffic is penalized. On the other hand, Drop-Sel RED provides O-UDP sources with the throughput they need, achieving an equitable distribution between VoIP and TCP at the same time.

For better clarification we consider a final scenario in which the unresponsive UDP sources generate the overall 66% of workload. Fig. 7b shows the results obtained. As it can be seen, whenever unresponsive traffic dominates PUNSI drastically penalizes it, whereas Drop-Sel RED provides a fairness share.

4.3 Drop-Sel network-level QoS evaluation

The provided QoS in VoIP applications strongly depends on the end-to-end packet delay. For each packet, the network dynamics, and more specifically, the router queue waiting time intervals, have a significant influence on the end-to-end latency.

4.3.1 Packet delay

In [22] it is shown that the probability of a packet being delayed longer than a given value in a system with front dropping is lesser than or equal to that in a system with rear dropping. Therefore, Drop-Sel discards the packet (for the selected traffic class) nearest to the front of the queue. In this way, the dropped packet generates an empty slot, a queue shift and, consequently a reduction of the queue waiting time for all the packets behind it. Front dropping thus leads to an overall end-to-end delay reduction, what is particularly significant for VoIP traffic given that it decreases the number of useless packets at the receiver. Additionally, note that for responsive sources front dropping accelerates the congestion detection and, implicitly its reaction.

Figure 8 shows the cumulative distribution function (CDF) of the end-to-end packet delay for each of the considered victim selection algorithms. In this case, we assume that RED AQM is adopted. Four different traffic workloads are considered (cases CA1, CA2, CA3 and CA4, see Table 2). In Fig. 8a it can be seen that Drop-Sel reduces the end-to-end latency compared to the other victim selection procedures. Similar trend is observed for cases CA2, CA3, and CA4 in Figs. 8b, 8c, and 8d. This behavior can be explained because Drop-Sel effectively reduces the average queue delay (see Table 5) and selects the proper traffic class to be dropped, independently of the load conditions. In spite of Drop-Front always drops the nearest packet to the output of the queue, it provides greater average end-to-end delay compared to Drop-Sel. For instance, in cases CA3 and CA4 (Figs. 8c and 8d) although TCP traffic load dominates, Drop-Front potentially reduces the responsive traffic dropping rate to prevent congestion. This fact further increases the TCP sources rate, increasing the queue occupancy, and rising the average queuing delay.

Therefore, we conclude that Drop-Sel improves the QoS for VoIP traffic classes regardless the traffic load. Additionally, note from Fig. 4 that Drop-Sel does not noticeably dismiss the throughput of the other (responsive and unresponsive) contending sources.

4.3.2 Packet loss probability

In general terms and, for VoIP traffic in particular, it is also important to evaluate the loss rate because of the final experienced quality will also depend on this parameter. There are two events that cause packet losses: first, packets are dropped in the AQM router for notifying and preventing congestion; and secondly, useless packets—those which ac-

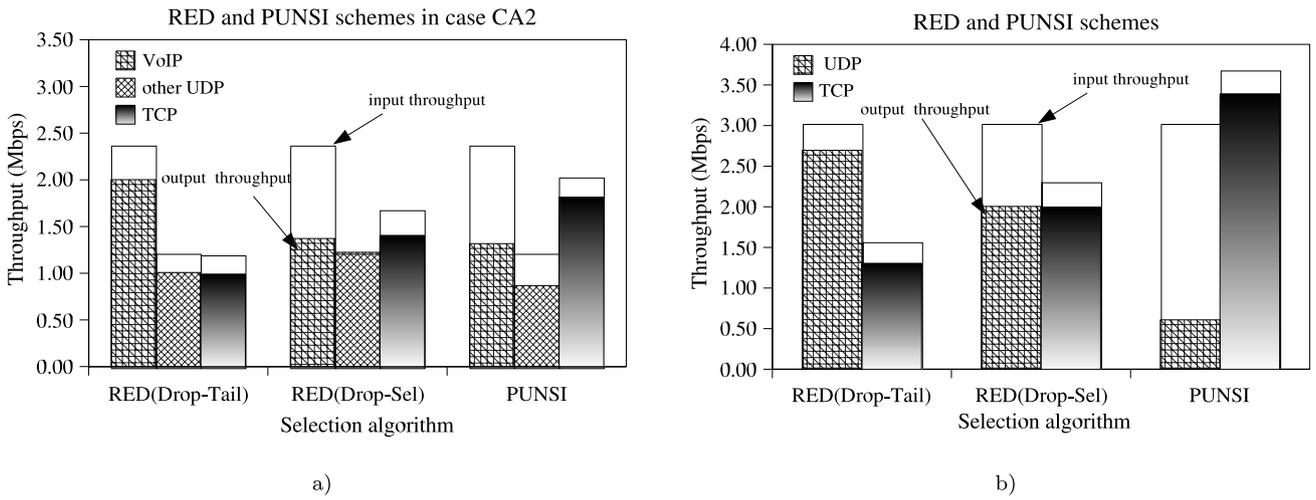


Fig. 7 Average throughput for victim selection algorithms in RED and PUNSI schemes

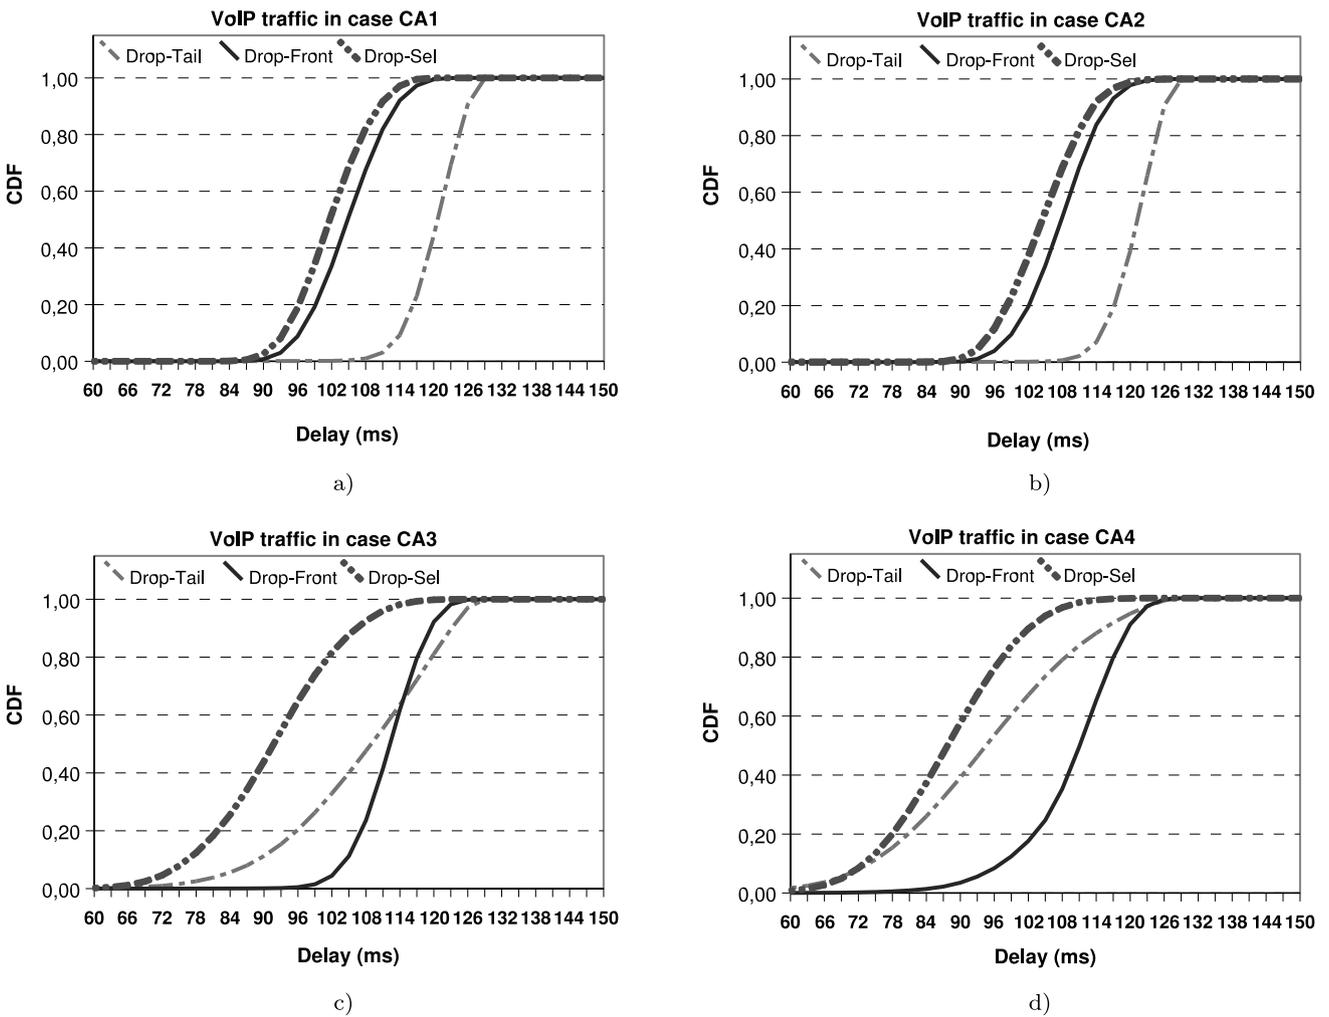


Fig. 8 Cumulative distribution fraction of VoIP packet delay in scenario S1, for the RED scheme

Table 5 Average delay in scenario S1 (dumbbell topology)

Case	Victim selection algorithms								
	Drop-Tail			Drop-Front			Drop-Sel		
	Queuing delay (ms)	End-to-end delay (ms)	End-to-end variance	Queuing delay (ms)	End-to-end delay (ms)	End-to-end variance	Queuing delay (ms)	End-to-end delay (ms)	End-to-end variance
CA1	83	123	22×10^{-3}	68	108	42×10^{-3}	65	105	40×10^{-3}
CA2	84	124	19×10^{-3}	71	111	42×10^{-3}	67	107	48×10^{-3}
CA3	71	110	17×10^{-2}	75	114	34×10^{-3}	55	94	13×10^{-2}
CA4	58	97	25×10^{-2}	73	112	83×10^{-2}	51	91	12×10^{-2}

Table 6 VoIP packets loss rates for scenario S2 (dumbbell topology)

Case		AQM Schemes								
		RED	RED	RED	AVQ	AVQ	AVQ	REM	REM	REM
		DropTail	DropFront	DropSel	DropTail	DropFront	DropSel	DropTail	DropFront	DropSel
CA5	R0 queue (%)	21.03	23.79	54.01	2.36	24.19	29.37	5.23	26.45	53.21
	final user (%)	11.94	7.13	4.33	0.07	0.40	0.57	7.91	5.88	3.90
	Total	32.97	30.92	58.34	2.43	24.59	29.94	13.14	32.33	57.11
CA6	R0 queue (%)	7.88	11.59	48.17	1.26	16.57	38.94	3.17	16.98	47.82
	final user (%)	14.19	13.29	5.60	0.12	0.19	0.69	9.05	11.99	5.26
	Total	22.07	24.88	53.77	1.38	16.76	39.63	12.22	28.97	53.08
CA7	R0 queue (%)	8.80	16.14	0.00	0.35	15.51	0.00	2.22	13.64	0.74
	final user (%)	11.84	11.78	4.24	0.43	0.39	0.00	8.71	6.23	1.41
	Total	20.64	27.92	4.24	0.78	15.90	0.00	10.93	19.87	2.15
CA8	R0 queue (%)	1.21	4.31	0.26	0.18	5.80	0.02	0.72	5.56	0.28
	final user (%)	5.63	11.29	3.23	0.12	0.07	0.10	4.04	8.19	2.89
	Total	6.84	15.60	3.49	0.30	5.87	0.12	4.76	13.75	3.17

accumulate end-to-end delay greater than 300 ms—are also dropped at the final users.

To show the impact of AQM packet victim selection procedures on VoIP packet loss probability, Table 6 provides the obtained results for different traffic workloads in the dumbbell topology. In this case, the results for the different cases of scenario S2 are reported. In terms of the overall loss probability, RED is the least suitable scheme for all the VoIP traffic rates considered. It achieves the slowest reaction to congestion. In consequence, the average enqueueing time is high and, therefore, the loss rate at the final user is significantly increased. For instance, in case CA6 the Drop-Tail RED rate of useless packets at the final user is 14.19%. However, with the AVQ and REM algorithms, this value falls to 0.12% and 9.05%, respectively.

Furthermore, results for cases CA7 and CA8 of Table 6 indicate that Drop-Sel reduces the total VoIP loss rate at the final user. It can be explained because of Drop-Sel produces the lowest queuing delay for this traffic, given that

TCP sources dominates and accordingly they are the most penalized, reducing thus the queue occupancy.

In addition, we can conclude that if VoIP flows are not taking more than their fair share, Drop-Sel minimizes the queue drop rate. For this situation, Drop-Sel offers potential performance guarantees for VoIP traffic class without considering explicit resource reservation approach. For example, in case CA7 Drop-Tail RED produces a queue drop rate of 8.80%, Drop-Sel RED reduces it to 0.00%. If cases CA7 and CA8 are considered for REM and AVQ, interestingly, note that again the Drop-Sel method leads to improvements both in the drop rate at the AQM router and in that of useless packets at the final user.

On the other hand, if VoIP sources dominates the traffic load, Drop-Sel will severely penalize it. This is because of the Drop-Sel TCP-friendly fairness property. We shall now see the consequences of this trend for VoIP traffic. For instance, for case CA6, Drop-Tail RED dropping rate is equal to 7.88%, whereas Drop-Sel raises it up to 48.17%.

Table 7 VoIP packets loss rates for scenario S3 (complex topology)

Case		AQM Schemes								
		RED	RED	RED	AVQ	AVQ	AVQ	REM	REM	REM
		DropTail	DropFront	DropSel	DropTail	DropFront	DropSel	DropTail	DropFront	DropSel
CA9	R2 queue (%)	1.44	5.16	0.00	0.05	5.65	0.01	0.47	4.81	0.10
	R5 queue (%)	12.02	14.12	0.00	1.87	11.99	0.01	2.00	13.82	0.12
	final user (%)	8.86	1.85	0.00	0.00	0.00	0.00	8.29	1.11	0.20
	Total	22.32	21.13	0.00	1.92	17.64	0.02	10.76	19.74	0.42
CA10	R2 queue (%)	0.01	0.01	0.00	0.01	0.00	0.00	0.00	0.00	0.00
	R5 queue (%)	9.61	12.15	0.00	2.14	11.07	0.03	1.88	10.32	0.05
	final user (%)	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
	Total	9.62	12.16	0.00	2.15	11.07	0.03	1.88	10.32	0.05

Table 7 shows the results obtained when the complex scenario S3 with multiple bottleneck links are considered. As it can be observed, the impact of AQM packet victim selection procedures on VoIP packet loss probability, is similar to the effects observed in the dumbbell topology.

For the reported experimental results, we can conclude that in AQM, better fairness can be provided if the dropping procedure observes the network traffic and selects the packet from the heaviest traffic. Additionally, for VoIP traffic the selection procedure can also improve the QoS in terms of end-to-end delay and packet loss rate probability whenever VoIP sources do not dominate the traffic load.

4.4 Perceptual evaluation of Drop-Sel

To provide a complete evaluation of the perceived quality obtained by the proposed mechanism, we provide two measures of perceptual quality: the Mean Opinion Score and the intelligibility level of the VoIP flows.

4.4.1 E-model quality evaluation

To assess the perceptual quality obtained with Drop-Sel, we use the E-model and ITU-T Recommendation G.107 [30]. 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 proposed in [31], and obtain the R factor using (1), defined as:

$$R = 94.2 - 0.11 \cdot (d - 177.3) \cdot H - 0.024 \cdot d - 30 \cdot \log(1 + 15 \cdot p) \tag{1}$$

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

$$H = \begin{cases} 0 & \text{if } (d - 177.3) < 0, \\ 1 & \text{if } (d - 177.3) \geq 0. \end{cases} \tag{2}$$

To provide more readable subjective evaluations, the R factor can be mapped to MOS punctuation [30]. Figure 9 shows the MOS values obtained for different traffic load conditions (cases CA5, CA6, CA7, CA8, CA9, and CA10 of Tables 3 and 4). In general, it can be observed that RED obtains the lowest MOS score regardless of the VoIP traffic load.

According to the objective evaluation reported in Tables 6 and 7, as it could be expected, AVQ provides the best subjective performance. This can be explained because of AVQ strategy reduces the queuing AQM delay and eventually the total packet loss rate.

Excepting cases CA5 and CA6, Drop-Sel improves the final user experienced quality for all the simulated AQM schemes. For instance in cases CA7, CA8, CA9, and CA10 it does not just improve the fairness, but it additionally has a noticeable positive impact on the provided subjective quality (Figs. 9c, 9d, 9e, and 9f). Note that in these cases responsive or O-UDP sources dominate the traffic load. However, when the VoIP sources cause the congestion because they dominate the traffic load, Drop-Sel penalizes the VoIP traffic (Figs. 9a and 9b). This occurs for cases CA5 and CA6, since the generated VoIP traffic percentage is greater than 33.33% of the available bandwidth, the fair quota for every competing type of traffic to use the 100% of the bandwidth.

The experimental evaluation shows a rational correlation between network-level and user-level quality estimations. We conclude, therefore, that Drop-Sel improves the fairness

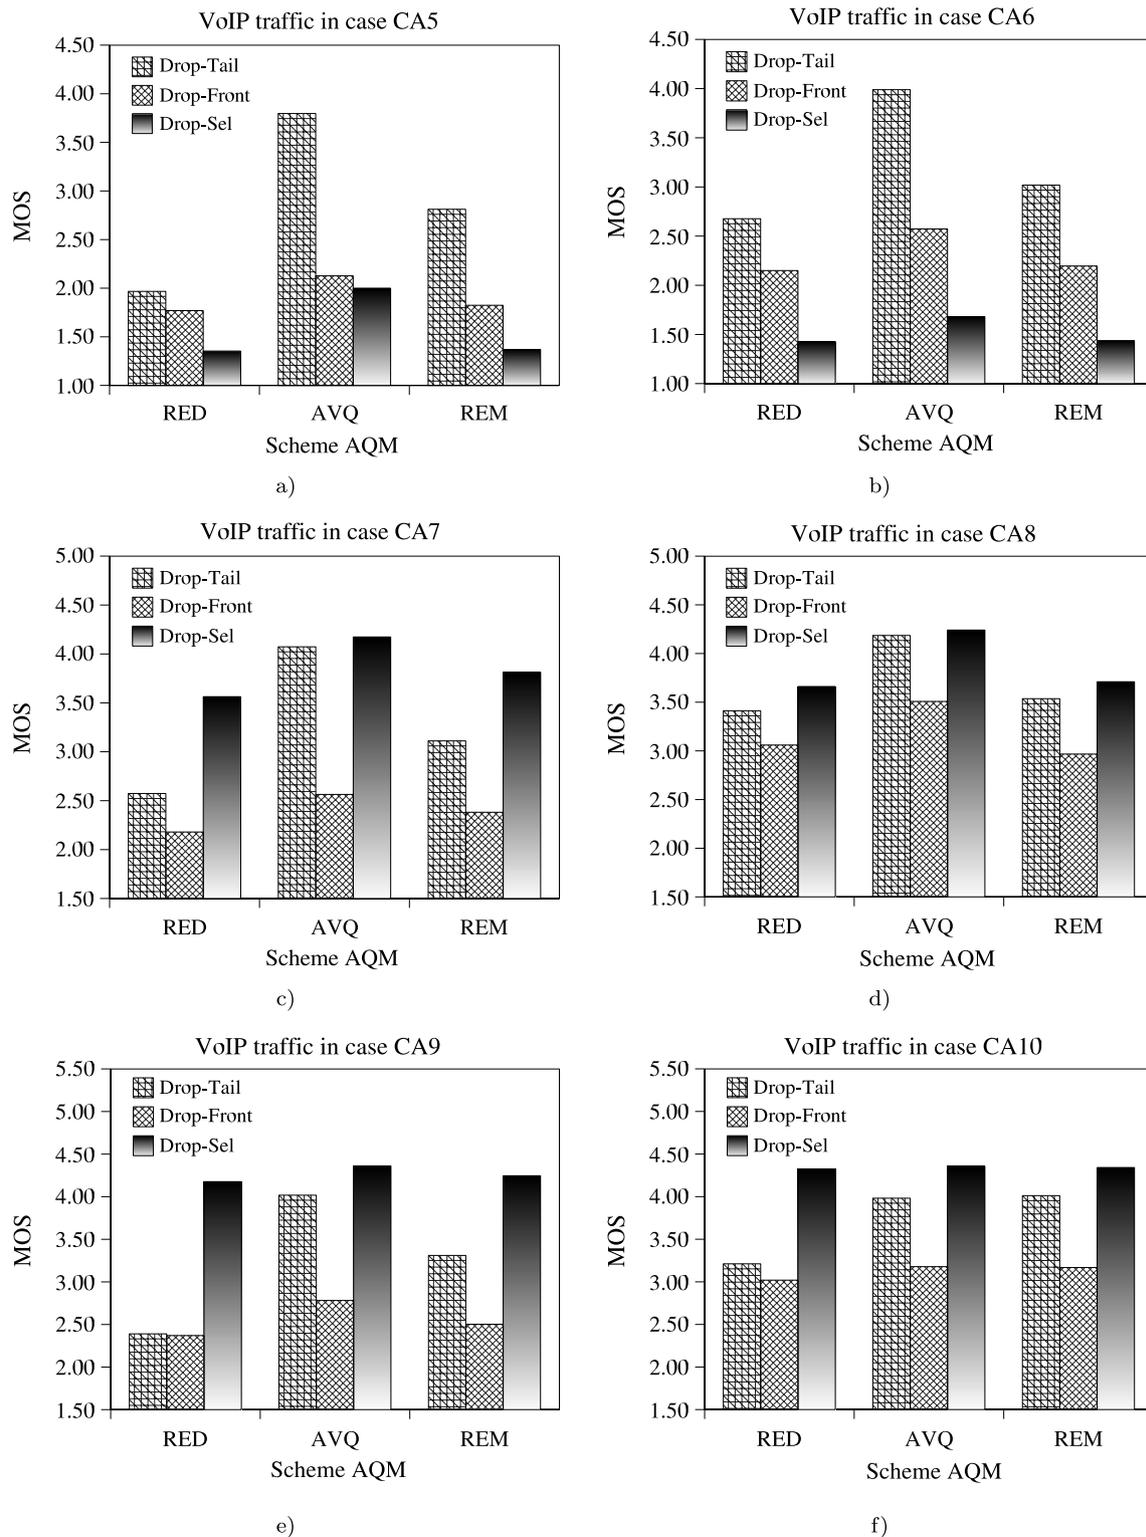


Fig. 9 E-MODEL based MOS evaluation of Drop-Front, Drop-Tail, and Drop-Sel under different traffic conditions

and the perceived quality, specially when the VoIP traffic percentage does not damage the fair quota of the remaining traffic. In addition, when the VoIP traffic predominates, for

scenarios in which fairness is not required and VoIP quality is the priority, we suggest to add automatic mechanisms to switch from Drop-Sel to other algorithm.

4.4.2 ASR-based perceptual evaluation

Additionally, to complete our evaluation, a third alternative score is considered to definitively establish the impact of Drop-Sel on the end-user perceived quality.

In this section, we evaluate the goodness of the proposed Drop-Sel mechanism by estimating the end-to-end voice intelligibility. To do that, we adopt an automatic speech recognition engine. The end-to-end automatic continuous recognition score (measured in terms of word or sentence accuracy rates) can be used to assess the quality of the reconstructed speech [32, 33].

Word and sentence accuracy rates are directly related to the final intelligibility perceived by the user. Thus, increasing intelligibility is a target for any proposed voice transmission algorithm. We believe this evaluation methodology is an essential tool for testing the improvements achieved with any proposed voice transmission scheme, on the basis of the following advantages that have been identified: firstly, the ASR score takes into account issues that would otherwise be difficult to measure, while no implicit model need be adopted, as speech recognition technology is now sufficiently mature. Additionally, it is low cost and provides highly reproducible measures.

The recognition task employed is based on the Aurora 2 speech database [34], a corpus that consists of connected digit sequences for American English speakers. After transmission, the speech signal is processed in order to reduce its inherent variability, and a feature extractor then segments the received speech signal into overlapped frames of 25 ms every 10 ms. Each speech frame is represented by a 14-dimension feature vector containing 13 Mel Frequency Cepstrum Coefficients (MFCCs) plus log-Energy. Finally, the feature vectors are extended with their first and second derivatives. Our speech recognizer machine is based on Hidden Markov Models (HMM). It uses eleven 16-state continuous HMM word models (plus silence and pauses, which have 3 and 1 states, respectively), with 3 Gaussians per state (except silence, with 6 Gaussians per state). The HMM models were off-line trained with a set of 8440 noise-free sentences, while the corpus test comprised 4004 noise-free out-of-training sentences. The above-mentioned experimental conditions were adopted for the experiments reported.

(1) Word accuracy

To evaluate the speech recognizer's performance, we define the Word Error-Rate (WER) as:

$$\text{WER} = \frac{n_i + n_s + n_d}{n_t} \times 100 \quad (3)$$

where n_s is the number of substituted words, n_i is the number of spurious words inserted, n_d is the number of deleted

words and, n_t is the overall number of words. Prior to counting the substitutions, deletions and insertion errors, dynamic programming is used to align the recognized sentence with its correct transcription.

Based on WER, the word accuracy (WA) is defined by

$$\text{WA} = 100 - \text{WER}. \quad (4)$$

The WA rates obtained from the simulations of case CA7 of scenario S2 are shown in Table 8. Specifically, several TCP and O-UDP sources compete with a total of 10 VoIP flows, generated by the 3 VoIP applications described in Table 1. In particular, application A generates 2 flows (labeled A1 and A2), application B generates 6 flows (B1, B2, B3, B4, B5, and B6), and application C generates 2 flows (C1 and C2).

In general, the scores resulted from Drop-Sel outperform those of Drop-Tail and Drop-Front. These improvements are noteworthy for flows B4, B5, and B6, which are the farthest sources, and which are the most prone to generate late packets. In particular, when RED is adopted, Drop-Sel improves an average of a 51.1% the WA score obtained with Drop-Front, and a 12.2% the score obtained with Drop-Tail for B4, B5, and B6. This improvement achieved by Drop-Sel is due to its lower number of VoIP packet discards at the router, and its lower generation of late packets at the end user side.

(2) Correct-sentence rate

The correct-sentence rate provides a complementary measure of intelligibility. A sentence is said to be correctly recognized whenever no word insertions or word substitutions are incurred.

Table 9 shows the correct-sentence rates obtained for case CA7. As it can be seen, the Drop-Sel rates are always better than or equal to the Drop-Tail and Drop-Front results. Thus, we experimentally show that considering the proposed victim selection scheme improves the end-to-end intelligibility. This fact is more significant for flows with worst quality, that is to say, for speech flows generated at distant sources, which are the most impaired, due to the impact of the delay on the E-model MOS score. For these remote sources the interactive nature of the speech is a major concern. For instance, flows B4 and B5 (Table 9) obtain rates of 53.92% and 53.52% respectively with RED Drop-Tail, whereas the Drop-Sel scheme raises the correct sentence rate to 76.47% and 76.72%, respectively. In Drop-Tail REM, flow B4 gets a rate of 76.27% and flow B5, 75.79%. However, Drop-Sel provides rates of 92.96% and 92.21% respectively, a very significant intelligibility improvement.

Summing up, compared to other victim selection procedures, after the conducted evaluation it can be concluded that for VoIP traffic the Drop-Sel algorithm improves the

Table 8 Word Accuracy rates of VoIP packet by flow for case CA7 in scenario S2

Flow	AQM Schemes								
	RED	RED	RED	AVQ	AVQ	AVQ	REM	REM	REM
	DropTail	DropFront	DropSel	DropTail	DropFront	DropSel	DropTail	DropFront	DropSel
A1	98.58	98.18	99.02	99.02	97.88	99.02	98.96	98.31	98.99
A2	98.73	98.21	99.02	99.02	98.42	99.02	99.00	98.56	99.02
B1	98.60	97.39	99.02	99.01	96.77	99.02	98.95	97.71	98.97
B2	98.35	97.57	99.02	99.02	96.49	99.02	98.94	97.82	98.99
B3	98.50	96.89	99.02	98.99	96.44	99.02	98.90	97.38	98.92
B4	78.54	38.85	90.50	98.04	91.48	98.59	85.18	71.77	97.19
B5	78.30	37.81	90.43	98.09	91.34	98.62	84.98	70.79	96.87
B6	78.55	39.41	91.11	98.10	94.77	98.62	85.40	72.76	97.10
C1	96.70	93.44	99.02	98.78	96.38	99.02	97.65	93.87	98.56
C2	96.98	91.42	99.02	98.59	95.53	99.02	97.65	92.88	98.49

Table 9 Correct Sentence rates of VoIP packet by flow for case CA7 in scenario S2

Flow	AQM Schemes								
	RED	RED	RED	AVQ	AVQ	AVQ	REM	REM	REM
	DropTail	DropFront	DropSel	DropTail	DropFront	DropSel	DropTail	DropFront	DropSel
A1	95.73	94.78	97.08	97.08	93.79	97.08	96.85	95.06	96.98
A2	96.20	94.73	97.08	97.08	95.38	97.08	97.00	95.68	97.08
B1	95.90	92.49	97.08	97.08	90.84	97.08	96.88	93.26	96.90
B2	95.08	93.11	97.08	97.08	90.41	97.08	96.85	93.66	97.03
B3	95.55	91.04	97.08	96.98	90.41	97.08	96.70	92.36	96.78
B4	53.92	18.91	76.47	94.71	78.95	95.88	76.27	57.07	92.96
B5	53.52	18.86	76.72	94.86	78.47	96.05	75.79	56.67	92.21
B6	54.09	19.86	78.02	94.81	86.31	96.00	76.20	57.84	93.01
C1	90.61	82.85	97.08	96.28	89.74	97.08	93.06	83.37	95.70
C2	91.46	78.80	97.08	95.75	87.81	97.05	92.86	81.62	95.50

end-to-end intelligibility besides of providing previously reported network level improvements.

5 Conclusions

To prevent network congestion, active queue management has been widely investigated. As a result, a number of interesting schemes have been proposed. In general terms, if different types of traffic (responsive and unresponsive sources) share the output AQM queue, they should be distinguished and differently processed. In general terms, the fairness and the final end-to-end subjective quality depend on the adopted victim selection procedure.

In this paper, we evaluate how the network performance (packet delay and loss rate) can be improved, in terms of both objective and subjective measurements, by simply observing and selecting the network traffic class to penalize for preventing congestion.

With this aim in mind, we have devised a straightforward scheme to be applied in AQM schemes. The proposed approach considers the congestion control demands of the application. In particular, we put forward a simple procedure for selecting the packet to be dropped at the AQM router. To detect the traffic that causes the congestion, we provide an algorithm which identifies the traffic class that will be accordingly penalized. Among the simulated approaches of reference, we experimentally prove that the proposed algorithm provides fairness and improves the end-to-end quality.

The proposed scheme, referred to as Drop-Sel, is evaluated in conjunction with a set of relevant AQM schemes: RED, REM, AVQ. To provide evaluations with other cutting-edge schemes, PUNSI is also simulated. We have considered a number of scenarios with different topologies and traffic loads, in which both UDP and TCP flows are generated to share the limited network resources. In our simulations, different traffic sources and destinations are included

at diverse locations in order to provide a wide range of possibilities. The experiments performed showed that Drop-Sel achieves a significant fairness improvement.

For multimedia traffic, and particularly for VoIP, we believe that achieving an excellent low-level QoS might be useless if this is not correlated to the (high-level) end user's satisfaction. Accordingly, to complete the evaluation we subjectively measured the benefits of adopting the proposed scheme. In this case, the E-model is used for evaluation purposes. Specifically, we estimated the MOS score under a wide range of experimental conditions. Additionally, we also show that for VoIP traffic the proposed scheme (Drop-Sel) also improves the end-user intelligibility. This is shown by adopting a methodology based on end-to-end automatic speech recognition measurements.

In summary, this work has demonstrated how our lightweight victim selection algorithm can provide both a fair service to VoIP and other types of UDP traffic, without harming the TCP throughput. In cases in which the VoIP traffic is not the predominant, Drop-Sel also raises the perceptual quality of the VoIP flows, both the MOS and the intelligibility scores, and benefits the TCP sources. If the VoIP traffic prevails, VoIP traffic is penalized accordingly in order to keep the fairness.

Acknowledgements We thank the reviewers, whose comments helped to improve the manuscript.

References

- Floyd, S., & Jacobson, V. (1993). Random early detection gateways for congestion avoidance. *IEEE/ACM Transactions on Networking*.
- Reguera, V. A., Álvarez Paliza, F. F., Godoy, W. Jr., & García Fernández, E. M. (2008). *Computer Communications*, 31, 73–87.
- Hashem, E. (1989). *Analysis of random drop for gateway congestion control*. MIT-LCS-TR-465.
- Jain, R., Ramakrishnan, K. K., & Chiu, D. (1988). Congestion avoidance in computer networks with a connectionless network layer. In *Proceedings of SIGCOMM '88*.
- Floyd, S., Gummadi, R., & Shenker, S. (2001). Adaptive RED: An algorithm for increasing the robustness of RED's active queue management. Available from: <http://www.icir.org/floyd/papers/adaptiveRed.pdf>, August 1, 2001.
- Jinsheng, S., Ko, K.-T., Guanrong, C., Chan, S., & Zukerman, M. (2003). PD-RED: to improve the performance of RED. *Communications Letters, IEEE*, 7, 406–408.
- Hadjadj Aoul, Y., Mehaoua, A., & Skianis, C. (2007). A fuzzy logic-based AQM for real-time traffic over Internet. *Computer Networks*, 51, 4617–4633.
- Kunniyur, S., & Srikant, R. (2004). An adaptive virtual queue (AVQ) algorithm for active queue management. *IEEE/ACM Transactions on Networking*, 12, 286–299.
- Wydrowski, B., & Zukerman, M. (2002). GREEN: an active queue management algorithm for a self managed Internet. In *Proceedings of IEEE international conference on communications, ICC'2002* (Vol. 4, pp. 2368–2372), April 2002.
- Athuraliya, S., Low, S. H., Li, V. H., & Qinghe, Y. (2001). REM: active queue management. *IEEE Network Magazine*, 15(3), 48–53.
- Sun, J., & Zukerman, M. (2007). RaQ: a robust active management scheme based on rate and queue length. *Computer Communications*, 30, 1731–1741.
- Lin, D., & Morris, R. (1997). Dynamics of random early detection. In *Proceedings of ACM SIGCOMM '97* (pp. 127–137).
- Pan, R., Breslau, L., Prabhakar, B., & Schenker, S. Approximate fairness through differential dropping. *Proceedings of ACM SIGCOMM Computer Communications Review*, 33(2), 23–39.
- Chung, J., & Claypool, M. Dynamic-CBT and ChIPS- Router support for improved multimedia performance on the Internet (Tech. Rep.). Computer Science, Worcester Polytechnic Institute. Available from: <http://web.cs.wpi.edu/claypool/papers/dcbt-chips/>.
- Stoica, I., Shenker, S., & Zhang, H. (1998). Core-stateless fair queuing: achieving approximately fair bandwidth allocations in high speed networks. In *Proceedings of ACM SIGCOMM*.
- Cao, Z., Wang, Z., & Zegura, E. (2000). Rainbow fair queuing: fair bandwidth sharing without per-flow state. In *Proceedings of INFOCOM* (pp. 922–931).
- Chan, S., Kok, C., & Wong, A. K. (2005). Multimedia streaming gateway with jitter detection. *Proceedings of Transactions on Multimedia*, 7(3), 585–592.
- Phirke, V., Claypool, M., & Kinicki, R. (2002). RED-Worcester traffic sensitive active queue management. In *Proceedings of 10th IEEE international conference on network protocols*. Available from: <http://csdl2.computer.org/comp/proceedings/icnp/2002/1856/00/18560194.pdf>.
- Pan, R., Prabhakar, B., & Psounis, K. (2000). CHOKe: a stateless active queue management scheme for approximating fair bandwidth allocation. In *Proceedings of IEEE INFOCOM '00* (pp. 942–951).
- Yamaguchi, T., & Takahashi, Y. (2007). A queue management algorithm for Fair bandwidth allocation. *Computer Communications*, 30, 2048–2059.
- Perkins, C., Hodson, O., & Hardman, V. (1998). A survey of packet-loss recovery techniques for streaming audio. In *Proceedings of IEEE Network* (pp. 40–48).
- Yin, N., & Hluchyj, M. G. (1993). Implication of dropping packets from the front of a queue. *Proceedings of IEEE Transactions on Communications*, 41(6), 846–851.
- Lakshman, T. V., Neidhardt, A., & Ott, T. J. (1996). *The drop from front strategy in TCP and in TCP over ATM*. New York: IEEE.
- Network Simulator ns2. Available from: <http://www.isi.edu/nsnam/ns/>.
- Schulzrinne, H., Casner, S., Frederick, R., & Jacobson, V. (2003). Rfc3550 RTP: A transport protocol for real-time applications, July 2003.
- ITU-T, Recommendation, G. 711, Pulse code modulation (PCM) of voice frequencies.
- Iacovoni, G., Morsa, S., Parisi, D., & Pierotti, D. (2002). *Broadband FRA scenario and traffic modeling*. Information Society Technologies.
- Liu, S. G., Wang, P. J., & Qu, L. J. (2005). Modeling and Simulation of self similar data traffic. In *Proceedings of the fourth international conference on machine learning and cybernetics* (pp. 3921–3925).
- Fomenkov, M., Keys, K., Moore, D., & Claffy, K. (2004). Longitudinal study of Internet traffic in 1998–2003. In *Proceedings ACM international conference* (Vol. 58).
- ITU-T, Recommendation, G.107, The E-model, a computational model for use in transmission planning, March 2005.
- Cole, R. G., & Rosenbluth, J. H. (2001). Voice over IP performance monitoring. *Proceedings of SIGCOMM Computer Communication Review*, 31(2), 9–24.

32. Ramos-Muñoz, J. J., Lopez-Soler, J. M., & Gomez, A. M. (2005). Intelligibility evaluation of a VoIP multi-flow block interleaver. In *Proceedings of IFIP IWAN*, Sophia-Antipolis France, November 2005.
33. Jiang, W., & Schulzrinne, H. (2002). Speech recognition performance as an effective perceived quality predictor. In *Tenth IEEE international workshop on quality of service* (pp. 269–275).
34. Hirsch, H. G., & Pearce, D. (2000). The AURORA experimental framework for the performance evaluation of speech recognition systems under noisy conditions. In *Proceedings of ISCA ITRW ASR: automatic speech recognition: challenges for the next millennium* (pp. 181–188).



Erika P. Alvarez-Flores received her B.Eng. degree in Electronic Engineering from the Hermosillo Technological Institute (ITH), Mexico (1992). In 2009 she got her Ph.D. degree by the University of Granada. Her primary fields of research include Active Queue Management and Multimedia Traffic.



Juan J. Ramos-Munoz received the B.Eng. degree in Computer Systems in 1999 and the Computer Sciences Engineering degree in 2001 from the University of Granada, Spain. Since 2004, he works as Assistant Professor in the Department of Signals Theory, Telematics and Communications of the University of Granada, where he has got his Ph.D. (2009). In 2006, he stayed at the Computer Network Department of the University of Basel, Switzerland as visiting researcher. His research interests are focused on VoIP,

Quality of Experience assessment, and adaptive mechanisms for providing QoS to Real-Time multimedia flows.



Pablo Ameigeiras received his M.Sc.E.E. degree in 1999 by the University of Malaga, Spain. He carried his Master Thesis at the Chair of Communication Networks, Aachen University (Germany) in the field of Random Access Mechanisms for Packet Services in UMTS. In 2000, he joined the Cellular System group at the Aalborg University (Denmark) where he carried out his Ph.D. thesis on Radio Resource Management and Quality of Service for packet services in HSDPA. After finishing his Ph.D., Pablo worked in Optimi Corporation investigating radio planning and optimization techniques for 3G networks. He recently joined the Department of Signal Theory, Telematics, and Communications at the University of Granada. His research interests include radio resource management, QoS, QoE and network performance.



Juan M. Lopez-Soler is Associate Professor at the Department of Signals, Telematics and Communications of the University of Granada (Spain). He teaches undergraduate and graduate courses in Computer Networks, Data Transmission, and Multimedia Networking. He received the B.Sc. degree in Physics (Electronics). In 1991-92, he stayed as Visiting Faculty Research Assistant in the Institute for Systems Research (formerly SRC) at the University of Maryland (USA). In 1995 he received the Ph.D. degree from the University of Granada. He has participated in more than 10 public and private funded research projects, being the coordinator in 6 of them. He has published more than 30 papers in international conferences and journals. His research interests include multimedia networking and real-time middleware.