Error Tolerant MAC Extension for Speech Communications over 802.11 WLANs

Antonio Servetti and Juan Carlos De Martin*
Dipartimento di Automatica e Informatica / *IEIIT-CNR
Politecnico di Torino, corso Duca degli Abruzzi 24, 10129 Torino, Italy
Phone: +39-011-564-7036 / -5421, fax: +39-011-564-7099
E-mail: [antonio.servetti | juancarlos.demartin]@polito.it

***Abstract*— The IEEE 802.11 MAC standard currently tolerates no errors in delivered packets. In this paper we study the performance of an error tolerant extension to the standard MAC that exploits the unequal perceptual importance of speech bitstreams. More in detail, we differentiate the CRC coverage on the speech payload in order to enable the delivery of partially corrupted packets. In our experiments network drivers of a wireless receiver have been modified to test three different CRC strategies. GSM AMR-WB speech transmissions in various channel conditions have been evaluated in terms of packet losses, goodput, and user perceived quality. Results show that accepting erroneous packets more than halves the number of lost packets with respect to the standard MAC implementation. However error detection on the speech payload can not be completely disabled. Only if the most perceptually important bits are protected speech quality is improved in every channel condition achieving quality gains up to 0.4 points of the MOS scale.**

I. Introduction

In recent years, the widespread adoption of IEEE 802.11 wireless LAN's is creating the basis of a new scenario for speech communications. However several challenges need to be addressed to provide successful speech services over a network originally intended for generic data traffic and characterized by potentially high error rates. While for data transfers, in fact, throughput is the main parameter for measuring the network performance, multimedia applications depend on strict quality of service (QoS) requirements in terms of packet losses and delay. In the current 802.11 standard, packet losses are also due to the need of data integrity during transfers, i.e., every hop has to discard all packets affected by channel errors, irrespective of the amount of corrupted data. This approach does not exploit modern multimedia compression algorithms that provide a certain degree of error resilience, so that the decoder can still benefit from corrupted packets. As a consequence a new error tolerant extension to the MAC layer is advisable for multimedia transmissions in wireless environments.

IEEE 802.11 medium access control (MAC) layer [1] provides a checksum to prevent forwarding of erroneous frames: if a bit or more are corrupted the packet is discarded and the sender will retransmit the data until a maximum retransmission limit is reached. Since speech data bits are known to have different perceptual importance [2], they can be packed in sensitivity order and a checksum applied only to the most important subset. Partial checksum will prevent useful frames from being dropped and will reduce the number of retransmissions, thus reducing the network load and delay.

Previous works proposed to implement selective error detection in the UDP protocol [3]. In the so-called UDP-Lite [4], the UDP checksum coverage is restricted to the most error sensitive part of the packet. Thus the receiver network stack forwards to the application also partially corrupted packets, provided that errors occur only in the non-checksummed data. According to the current IEEE 802.11 MAC standard, straightforward application of UDP-Lite to wireless networks is not possible because erroneous frames are dropped by the link layer before reaching the UDP layer. Also the recent 802.11e extension to the MAC standard [5], specific for multimedia applications, does not allow the adoption of a partial checksum technique.

However the idea of an error tolerant 802.11 network has already been discussed in the literature. While some paper suggest to totally disable the data integrity checks in the MAC layer [6] (thus disabling the retransmission mechanism too), [7] and [8] take a step forward introducing the idea of reflecting the UDP-Lite policy of sensitive and non-sensitive data in the MAC protocol. That permits to detect, discard, and retransmit heavily damaged packets not only at the receiver-end, but also at every wireless hop. In case of high delay scenarios, where end-to-end retransmissions would not be applicable, hop-by-hop retransmission may enable the delivery of an "acceptable" packet, and with lower overall delay. Results obtained by means of network simulations [9][10] show that allowing bit errors in WLAN speech transmissions decreases packet losses, delay and the number of retransmissions. This results in better call quality and, in some cases, in the ability to support a larger number of calls.

In this paper we present the results of actual speech transmissions over an 802.11b network where wireless device drivers have been modified to test a receiver with MAC partial checksum. Our main objective is to assess the performance of this technique in a real environment with particular attention to the speech quality perceived by the user. With respect to

This work was supported in part by MIUR project FIRB-PRIMO, http://primo.ismb.it

0-7803-8887-9/05/$20.00 (c)2005 IEEE

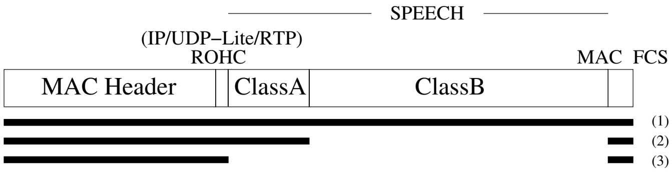

Fig. 1. Total (1), partial (2), and header-only (3) error detection coverage.

previous works our original contribution is in the simulation and evaluation methods. Instead of a model we use real 802.11b experiments. In addition, our performance analysis considers both packet loss rate (that is clearly reduced when partial checksum is used) and user satisfaction after speech decoding. Network experiments show that speech distortion introduced by decoding partially corrupted packets is clearly lower than the distortion that would have been caused by their discarding. Quality gains that exceed 0.3 points of the MOS scale are obtained. We also verify that error detection of the speech payload cannot be totally disabled: accepting more packets but with most sensitive bits damaged is, in fact, harmful to the speech quality.

The paper is organized as follows. In Section II, we describe the proposed selective bit error checking MAC protocol for wireless multimedia. Performance evaluation and quality results are presented in Section III. Conclusions are drawn in Section IV.

II. Error Tolerant Extension to 802.11 MAC

In a wireless environment packet communications may be affected by high error rates because of interference with other signal sources as well as other radio problems. To overcome this problem the IEEE 802.11 MAC [1] requires that each correctly received protocol data unit (MPDU) be acknowledged. To verify the MPDU integrity the frame format provides a 4-byte Frame Check Sequence (FCS) field: if one or more bits are corrupted the packet is discarded and the sender will retransmit the data until a maximum retransmission limit is reached. For each retransmission, to reduce the probability of repeated collisions, a longer random backoff time is selected thus incrementing the network delay for a successful transmission.

The adoption of an error tolerant extension in the MAC layer, via selective error detection, can produce several positive effects on the network. However to allow bit errors in the packet payload, the IEEE 802.11 standard must be modified. In the MAC layer header control bits can be introduced to indicate the portion of the packet protected by the CRC thus enabling an adaptive cross layer definition of the sensitive part of the payload. As already noted in [10] the main disadvantage of this modification is that the new control headers are not backward compatible with other standard devices.

To illustrate the behavior of the proposed technique let us consider n-bit long packets where only m bits of data ($m < n$) are checksummed, and so verified at the receiver. For simplicity's shake, and only in the current section, we assume an i.i.d. error source with mean p. The packet loss rate (PLR) is then $PLR_m = [1-(1-p)^m]$, independently of n. In a full checksum scheme, instead, m is equal to n. It is then evident that selective error detection decreases the PLR. But for evaluating its performance in case of multimedia communications we must account also for the presence of *corrupted* packets, i.e. packets where errors occur only outside the checksum coverage. The packet corruption rate (PCR), for the case of a single transmission per packet, is then $PCR_m = (1-p)^m \cdot [1-(1-p)^{(n-m)}]$. Thus, with no retransmissions (e.g., in multicast mode), we note that the sum of corrupted and lost packets with selective detection shall be the same as the number of lost packets without it.

If the retransmission limit is set to $N-1$ then the packet loss rate decreases with increasing values of N: $PLR_m(N) = (PLR_m)^N$. Also the expected number of transmissions T for each packet is lower:

$$T = \sum_{i=1}^{N-1} i(1-PLR_m)PLR_m{}^{i-1} + N \cdot PLR_m^{N-1} = \frac{1-PLR_m^N}{1-PLR_m}.$$

Consequently end-to-end delay and network load are reduced with positive effects on the quality of service that can be achieved by all the transmissions in the network.

The complex effect of retransmissions on the number of corrupted packets and on the perceived quality will be further investigated in the experiments in Section III. Retransmissions, in fact, play a fundamental role in favoring full or partial checksum techniques. The latter solution guarantees a lower PLR, often accepting corrupted packets, while the former presents more losses, but ensures the integrity of the received ones. In the following, we drop the rather unrealistic assumption of uniform bit error probability, and we derive bit error rate and error location from actual transmission experiments.

III. Performance evaluation

Performance of the proposed error tolerant extension to the 802.11 MAC is here evaluated for speech transmissions using the wideband GSM Adaptive Multi Rate (AMR) coder [11]. For a study on the same idea applied to video communications, see [12].

A. CRC Strategies

Figure 1 shows three different strategies for error detection coverage of an 802.11 data MPDU with speech payload. Strategy 1 (T-CRC) represents the typical IEEE 802.11 CRC as defined in the standard. In this case the MAC checksum covers the whole packet, and the packet is dropped wherever a bit error occurs. In strategy 2 (P-CRC), the payload is divided in two classes. Class A, that contains the speech bits most sensitive to errors, is covered by the CRC. The other bits are left unprotected, thus packets with damaged Class B can be saved for the speech decoding process. Strategy 3 (H-CRC), limits the checksum coverage to the MAC, IP, UDP-Lite, and RTP headers. Packets with errors in the header are dropped, triggering retransmissions.

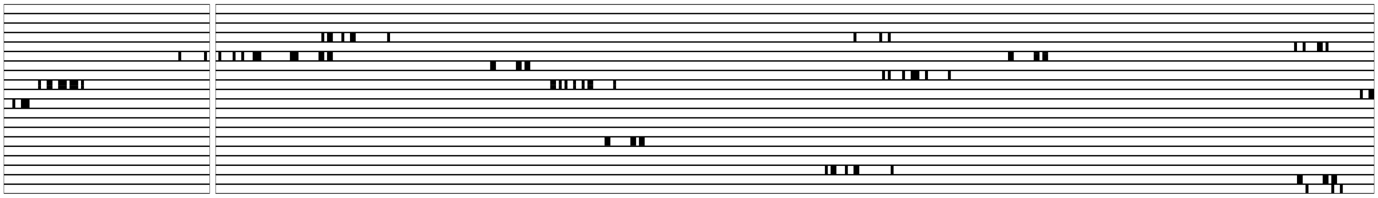

Fig. 2. Excerpt of the a measured 802.11 error pattern mapped on transmitted speech data. Each line corresponds to a 23.85 kb/s GSM AMR-WB frame (the block on the left represents Class A bits). Erroneous bits are marked as black points.

In the following experiments these strategies are applied on the 20 ms speech frames generated by the 23.85 kb/s GSM AMR-WB coder. Speech encoder output bits are ordered according to their subjective importance and divided in two classes: Class A (first 72 bits) and Class B (last 405 bits), as defined in the standard. Any error in Class A bits typically results in a corrupted speech frame which should not be decoded without applying appropriate error concealment. Errors in Class B gradually reduce the speech quality, but decoding of an erroneous speech frame is usually possible without annoying artifacts. Besides the RTP, UDP-Lite, and IP protocols headers, for which Robust Header Compression is assumed, an additional two-byte field is introduced in the 24-byte MAC header to specify the number of bits covered by the checksum. The checksum value is then expressed as usual as the last four bytes of the packet.

B. Wireless transmission experiments

Characterizing the error behavior of the 802.11 channel is a fundamental issue for assessing the performance of the proposed technique. While it is well known that wireless links typically have higher error rates than their wired counterparts, the detailed characteristics of wireless errors are not easy to reproduce in a computer simulation [13].

Our study is then based on actual 802.11 WLAN transmissions. We transmitted a well known packet stream over an 11 Mb/s 802.11b wireless network using specially formatted UDP packets. Their payload includes information for error detection such as a redundant sequence number and a repeated signature. The sequence number repetitions are used to estimate, via a majority criterion, the original sequence number at the sender. The signature is used to filter, among all received packets, the ones related to the experiment. This is necessary because the network driver is working in promiscuous mode.

Due to our interest in studying the distribution of bit errors inside corrupted packets, we modified a wireless device driver in order to collect data for every received packet, including erroneous transmissions. More specifically, the receiver was a Linux box using Prism 2 wireless 802.11b PCMCIA card and the wlan-ng (ver 0.2.1-pre20) device drivers. When in monitor mode the modified device drivers deliver all packets to the upper network layers, thus the traces collected at the client by the network sniffer (ethereal ver 0.10.0a, with libpcap 0.7.2) included both error-free and error prone transmissions. Traces can then provide bit-level error information by bit-wise comparing sent and received packets. The bursty nature of the wireless channel is visible in Figure 2, where part of a trace is shown and bit errors, inside speech frames, are represented by black boxes.

Figure 3 shows two sets of measurements from a sample experiment. Packets simulating a 23.85 kb/s voice communication are transmitted between two wireless hops to show the effect of the proposed error tolerant extension on the packet loss rate. Analysis of the number of erroneous bits in received packets is used to trace the time evolution of the channel bit error rate. We compare the percentage of packets discarded by the receiver after the error detection process for the three strategies under investigation: clearly partial error detection always guarantees a reduction in the PLR. An interesting result is achieved at low bit error rates (e.g., below 10^{-3}) when only few MPDU bits are corrupted. In this case the standard MAC strategy (T-CRC) prevents decoding of received speech frames, while other strategies present no losses. We will see in this

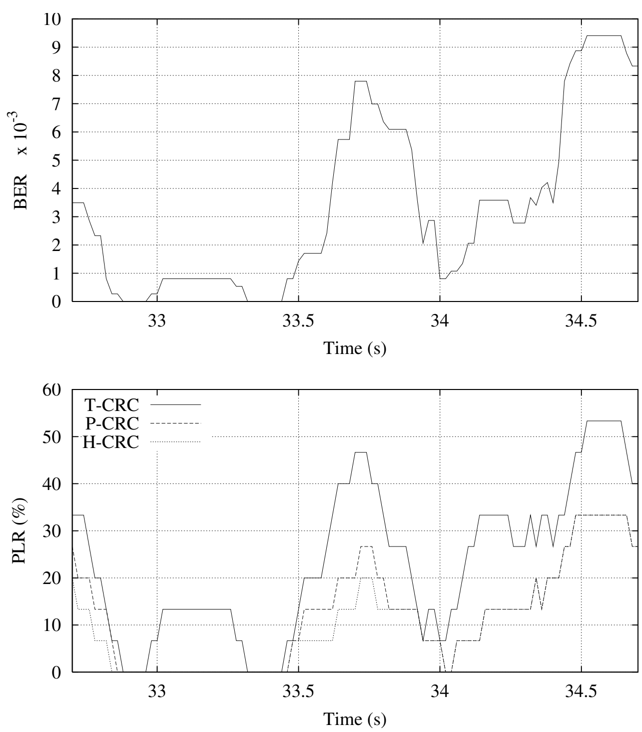

Fig. 3. Multicast transmission of GSM AMR-WB speech at 23.85 kb/s. Time evolution of the packet loss rate for the three strategies under investigation is related to the channel bit error rate. Measurement results are averaged over a sliding window of 300 ms.

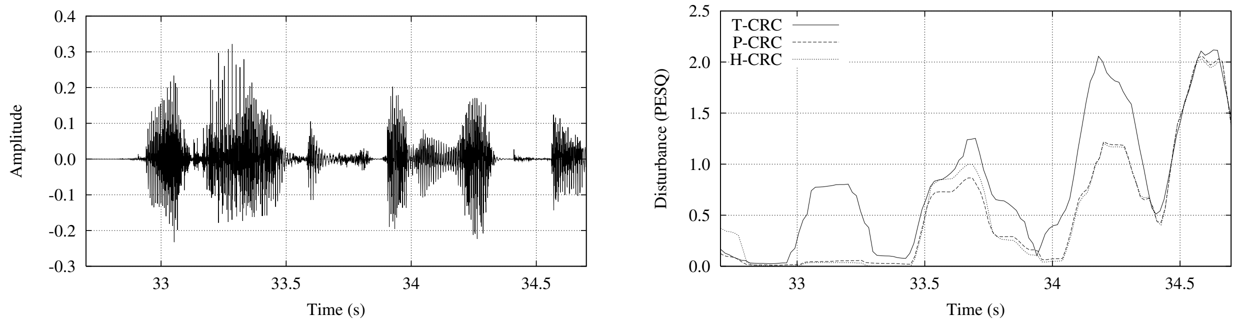

Fig. 4. Multicast transmission of GSM AMR-WB speech at 23.85 kb/s. The disturbance *(right)* perceived by the user by dropping or decoding corrupted frames is illustrated as a function of time for the speech sample in *(left)*.

condition the benefit of decoding slightly damaged packets.

C. Speech Quality Assessments

MPDU transmission experiments between two wireless 802.11b hops have been recorded in different channel conditions. Lost and corrupted packets, as well as retransmissions, are measured at the receiver depending on the error coverage of the different strategies. Corrupted speech frames are decoded without further processing by the AMR-WB decoder, lost frames are concealed as defined in the standard. Perceived quality results are then expressed by means of the objective speech quality measure given by the PESQ-MOS algorithm [14]. Speech samples have been taken from the NTT Multilingual Speech Database. We chose 24 sentence pairs spoken by two English speakers (male and female). 9600 packets are transmitted in each experiment for a total of 192 seconds of speech, silence included.

In the first scenario packets are sent in multicast mode, so no link-level retransmission is used. Figure 4 presents the effects of allowing bit errors in speech data. Perceived quality is evaluated for the the transmission in Fig. 3. In the vertical axis we represent, as a function of time, the disturbance value extracted from the PESQ algorithm for each analysis frame. These values will then be aggregated over the whole speech signal to generate the PESQ score. We note that decoding also partially corrupted frames generally reduces the overall level of disturbance, thus increasing the user satisfaction.

Additional experiments are illustrated in Table I that lists packet loss rate, and objective speech quality (on a MOS scale) for different average bit error rates[1]. Packet corruption rate, in this case, can be derived by subtracting the measured PLR from the PLR of the T-CRC strategy. With the error tolerant extensions (H-CRC, P-CRC) the percentage of lost frames is consistently lower than in the reference standard CRC. The PESQ score of the corresponding decoded speech confirms that, for the scenarios under consideration, it is better to receive and decode partially corrupted packets than to lose them altogether. Improvements that range from 0.2 to 0.5 on the PESQ-MOS scale are clearly noticeable by users, proving partial checksum particularly efficient at high error rates. This quality gain is motivated by the presence of a great number of bits only lightly sensitive to errors in the speech frame. No perceptible difference is, however, appreciated between H-CRC and P-CRC.

In the second scenario we test the effect of retransmissions. In this case, upon packet reception, the wireless hop has an additional chance with respect to the multicast case. Instead of simply dropping or forwarding a damaged packet, it can wait for a retransmission that may provide, at the expenses of an higher network delay, an error free packet. To study this scenario our experiments considered 802.11 transmissions with the retransmission limit set to four. From Fig. 5, we note that the three strategies present not only different loss behavior and quality, but also different channel utilization given a particular BER. Considering only the packet loss rate, both partial checksum strategies P-CRC and H-CRC appear promising: the number of lost packets halves with respect to the standard implementation and the channel goodput (the ratio of correctly-received packets to the number of transmitted packets) increases. The modified MAC, in fact, neither discards a packet if it is only "slightly" damaged, nor attempts a retransmission. However, we cannot assume that the best decoded speech quality is always achieved by the strategy with the lowest PLR. Our experimental results show, in fact, that the H-CRC strategy can even deliver lower quality than the standard MAC checksum implementation. The proof is

Trace	**T-CRC**		**P-CRC**		**H-CRC**	
BER	**PLR**	**MOS**	**PLR**	**MOS**	**PLR**	**MOS**
$3.19 \cdot 10^{-4}$	2.40	3.56	1.38	3.74	1.14	3.71
$4.89 \cdot 10^{-4}$	3.32	3.39	1.99	3.58	1.70	3.60
$1.33 \cdot 10^{-3}$	5.94	3.03	3.87	3.26	3.30	3.29
$1.41 \cdot 10^{-3}$	9.16	2.92	5.21	3.23	4.32	3.24
$3.94 \cdot 10^{-3}$	17.70	2.33	11.22	2.67	9.69	2.69
$3.16 \cdot 10^{-3}$	22.53	2.20	12.41	2.75	10.22	2.75

TABLE I

FULL AND PARTIAL CHECKSUM PERFORMANCE FOR GSM AMR-WB AT 23.85 KB/S. LINK-LEVEL RETRANSMISSIONS ARE DISABLED.

[1]The percentage of lost frames is not always proportional to the bit-error rate because of the non-uniform nature of wireless errors, i.e., for the same BER, a lower number of packets is lost if bit errors are more bursty.

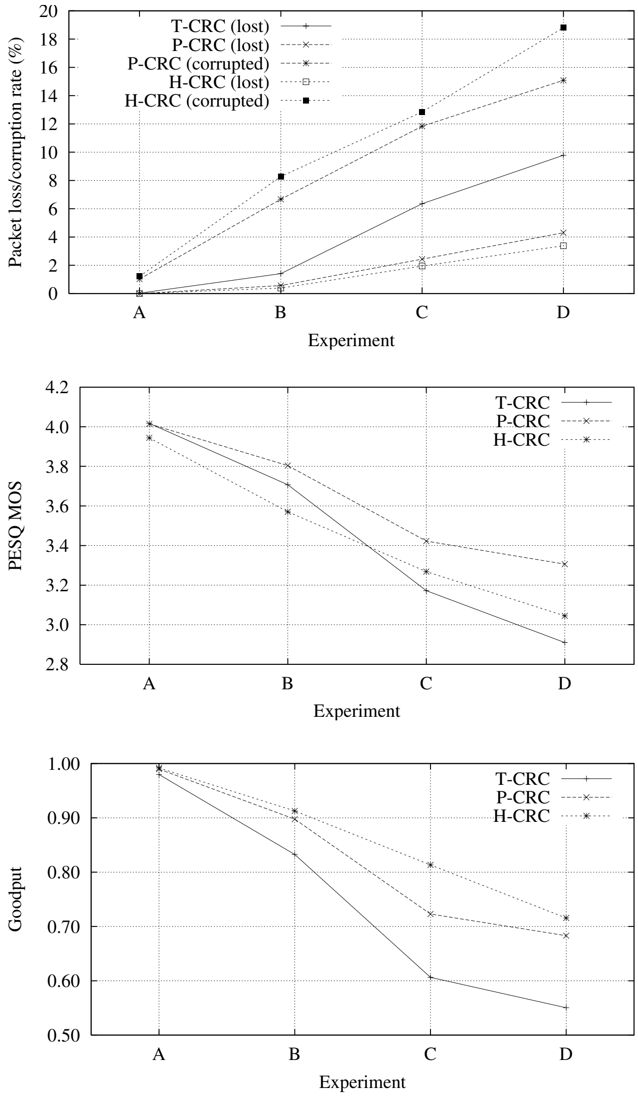

Fig. 5. Performance in terms of packet loss rate *(top)*, PESQ score *(middle)*, and goodput *(bottom)* for IEEE 802.11 GSM AMR-WB speech transmission at 23.85 kb/s. Four channel conditions are considered with average BER of $2.6 \cdot 10^{-4}$, $3.5 \cdot 10^{-3}$, $1.2 \cdot 10^{-2}$, and $1.9 \cdot 10^{-2}$ respectively. The maximum number of link-level retransmissions is set to four.

evident in the PESQ score of Fig. 5. When the bit error rate is low the corrupted frames delivered by H-CRC introduce a distortion that outruns the advantage of receiving more packets than the other strategies. P-CRC instead, assuring that the most perceptually important bits are always correct and useful, well combines positive effects on network load and improved speech quality: the MOS is significantly increased of more than 0.2 and channel goodput confirms a higher level of network utilization.

IV. Conclusions

Three selective error detection strategies have been investigated to enhance IEEE 802.11 link-layer effectiveness in supporting speech communications: standard coverage, protection of the most perceptually sensitive speech bits only, no CRC on the speech payload. Packets are then retransmitted only if errors occur inside the checksum coverage, otherwise erroneous speech frames are decoded as they are. Experiments with GSM AMR-WB in various channel conditions show that the assumption that bit corruption results in only minor distortion is valid only if the most sensitive part of the payload is checksummed. In this case speech quality is improved for every channel bit error rate achieving quality gains up to 0.4 points on the MOS scale. Furthermore, the proposed error detection technique also reduces the network load, since successful packet delivery requires, on average, less retransmissions.

References

[1] ISO/IEC, "Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications," *ANSI/IEEE Std 802.11*, 1999.

[2] K. Swaminathan, A.R. Hammons Jr., and M. Austin, "Selective error protection of ITU-T G.729 codec for digital cellular channels," in *Proc. IEEE Int. Conference on Acoustics, Speech, and Signal Processing*, Atlanta, Georgia, USA, May 1996, pp. 577–580.

[3] L.-A. Larzon, M. Degermark, and S. Pink, "UDP lite for real-time multimedia applications," in *Proc. QoS mini-conference of IEEE Int. Conference on Communications (ICC)*, Vancouver, Canada, June 1999.

[4] L.-A. Larzon, M. Degermark, S. Pink, L.-E. Jonsson, and G. Fairhurst, "The UDP-lite protocol," *draft-ietf-tsvwg-udp-lite-02.txt*, August 2003.

[5] IEEE 802.11 WG, "Draft supplement to standard for telecommunications and information exchange between systems - LAN/MAN specific requirements - part 11: Wireless medium access control (MAC) enhancements for quality of service (QoS)," *IEEE 802.11e/D5.0*, August 2003.

[6] S. Khayam, S. Karande, H. Radha, and D. Loguinov, "Performance analysis and modeling of errors and losses over 802.11b LANs for high-bit-rate real-time multimedia," *Signal Processing: Image Communication*, vol. 18, no. 7, pp. 575–595, August 2003.

[7] A. Servetti and J.C. De Martin, "Link-level unequal error detection for speech transmission over 802.11 networks," in *Proc. Special Workshop in Maui - Lectures by Masters in Speech Processing*, Maui, Hawaii, USA, January 2004.

[8] H. Dong, D. Chakares, A. Gersho, E. Belding-Royer, and J. Gibson, "Selective bit-error checking at the MAC layer for voice over mobile ad hoc networks with IEEE 802.11," in *Proc. IEEE Wireless Communications and Networking Conference (WCNC)*, Atlanta, GA, USA, March 2004, pp. 1240–1245.

[9] A. Servetti and J. D. Martin, "802.11 MAC protocol with selective error detection for speech transmission," in *Proc. 3rd International Workshop on QoS in Multiservice IP Networks*, Catania, Italy, February 2005.

[10] I. Chakeres, H. Dong, E. Belding-Royer, A. Gersho, and J. Gibson, "Allowing errors in speech over wireless LANs," in *Proc. 4th Workshop on Applications and Services in Wireless Networks (ASWN)*, Boston, MA, USA, August 2004.

[11] ETSI, "AMR speech codec, wideband; general description," *ETSI TS 126 171 version 5.0.0*, August 2002.

[12] E. Masala, M. Bottero, and J.C. De Martin, "Link-level partial checksum for real-time video transmission over 802.11 wireless networks," in *Proc. 14th Int. Packet Video Workshop*, Irvine, CA, USA, December 2004.

[13] S. Khayam and H. Radha, "Markov-based modeling of wireless local area networks," in *Proc. 6th ACM Int. Workshop on Modeling, Analysis, and Simulation of Wireless and Mobile Systems*, San Diego, CA, USA, September 2003, pp. 100–107.

[14] ITU-T, Recommendation P.862, "Perceptual evaluation of speech quality (PESQ), an objective method for end-to-end speech quality assessment of narrowband telephone networks and speech codecs," February 2001.