

Measuring VoIP Performance in IEEE 802.11p Vehicular Networks

Francesca Martelli, M. Elena Renda, Paolo Santi

IIT - CNR, Pisa, Italy

E-mail: {francesca.martelli, elena.renda, paolo.santi}@iit.cnr.it

Marco Volpetti

Dipartimento di Informatica

Università degli Studi di Pisa, Pisa, Italy

Abstract—This paper studies, for the first time to our best knowledge, VoIP performance over V2V IEEE 802.11p links. VoIP performance is evaluated in terms of both throughput and jitter through on-the-road measurements, investigating also the effects of inter-vehicle distance and speed on VoIP performance. Results of our study show that VoIP performance is quite poor, especially due to an excessively large number of lost packets under all investigated scenarios. In-lab experiments performed to investigate possible causes of this poor performance reveal that throughput degradation is likely to be due to inefficiencies in implementation of the vehicular networking stack, rather than to poor communication quality.

I. INTRODUCTION

Vehicular Ad Hoc Networks (VANet) are expected to play a major role in a near future, most importantly due to their potential of increasing safety conditions on the road and improving traffic conditions. However, VANets can be used also to realize advanced applications and services aimed at improving the driving experience in general. For this reason, the recently released IEEE 802.11p standard uses one control channel and up to six service channels, that can be used for realizing non safety related applications.

One of the most popular infotainment applications is Voice Over IP (VoIP), due to its advantages with respect to traditional circuit-switched telephony: high compression of voice data, bandwidth efficiency, and low cost for the user. In vehicular environments, VoIP can be used for instance to setup voice calls between a group of friends traveling in nearby vehicles towards the same destination.

The main purpose of this paper is testing VoIP performance in IEEE 802.11p-based VANets through on-the-field measurements. Our target, motivated by a group voice call scenario as the one described above, is vehicle-to-vehicle (V2V) communications in a typical car following situation. So far, all measurement studies for VANets we are aware of were concerned with either evaluation of the vehicular channel PHY layer properties [1], [2], [3], or with safety application performance evaluation [4], [5]. VoIP application requirements, though, are very different from those of safety applications, since they need a larger amount of bandwidth, and they are subject to strict timing constraints to obtain an acceptable voice quality. In particular, the most important metric for evaluating a VoIP communication is *jitter* [6], defined as the variation of end-to-end delay from one packet to the next one within the same flow. Furthermore, from a

networking perspective VoIP applications use IP-based unicast communications, while safety applications use broadcast communications and do not necessarily use IP formatting. Given the above, existing measurement studies cannot be used to gain an understanding on how VoIP applications are expected to perform in a real-world scenario.

There are some papers in which the capacity (namely, the number of concurrent voice calls supported) of VoIP over IEEE 802.11 is considered by means of analysis and/or simulation [7], [8], [9], [10]. In [8], two problems are addressed: increasing the VoIP capacity in a WLAN environment, and improving the VoIP performance in presence of TCP traffic, for different codecs, considering both CBR and VBR voice streams, and different 802.11 standards (a/b/g/e). Similarly, in [7], a comparison among IEEE 802.11a/b/g technologies is presented, with both CBR and VBR. In [9] and [10], the voice capacity limit is studied considering both IEEE 802.11e and 802.11b WLAN environments. Papers [11], [12] are based on experimental works, by measuring effective throughput and achieved capacity in IEEE 802.11 WLANs with one or more access points.

Very few papers exist about VoIP transmissions over MANETs. In [13], a simulation study for establishing which audio codec is better under a particular topology creation protocol (WEAC) is presented; the authors conclude that G.729 codec is the most suitable (with respect to G.711 and G.723.1) in an ad hoc network using WEAC routing protocol. In [14], the deployment of VoIP applications over an IEEE 802.11b MANET using AODV routing protocol is analyzed by means of simulations, focusing on performance metrics such as jitter, one-way delay and frequency of service interruptions. A similar study has been conducted in [15], where a new strategy for route selection and playout control is proposed to improve the user perceived quality.

The only experimental work about VoIP transmission over a MANET we are aware of is presented in [16], where the authors have considered AODV and OLSR routing protocols, and different audio codecs to establish the best configuration w.r.t. jitter, end-to-end delay and throughput.

Thus, to the best of our knowledge, ours is the first experimental study of VoIP performance over V2V IEEE 802.11p links. The on-the-road measurements we report here reveal several interesting insights on VoIP performance in V2V 802.11p networks. Most importantly, our study shows

that particular attention should be paid in the implementation of the IP protocol stack in vehicular environments, since approaches based on encapsulation/decapsulation of IP packets into WAVE-like packets, as done by the IEEE 802.11p compliant NEC LinkBird-MX units used in this study, appears to be inefficient.

II. HARDWARE, SOFTWARE AND MEASURES

A. Physical Setup

In order to implement a small VANet, in our measurements we used two IEEE 802.11p compliant NEC LinkBird-MX v3 devices (called LinkBird in what follows). The LinkBirds units are embedded Linux machines (kernel 2.6.19) based on a 64 bits MIPS processor working at 266Mhz. The PHY layer parameters of the IEEE 802.11p LinkBird network interface are summarized in Table I.

A small size (108 mm), omni-directional WiMo antenna operating in the 5.5–5.8 GHz frequency range, with 5dBi gain, was connected to the LinkBird device and mounted on the roof of the vehicles used during our V2V measurements.

Due to the limited computational power of the LinkBirds, we developed the VoIP application on a laptop, connected with the LinkBird through an Ethernet cable for transmitting/receiving packets. The laptop is also connected to a GPS receiver in order to collect/log relevant data (time, vehicle speed, latitude and longitude), used in the data post-processing phase (see Section II-B). The onboard equipment on each vehicle is composed of a LinkBird, a GPS receiver, a laptop, and the rooftop antenna. For further details and pictures on the physical setup, please refer to [4].

Parameter	Details
Frequency/Channel	5725 – 5925 MHz
Bandwidth	10 – 20 MHz
Version	IEEE802.11p Draft 3.0, July 2007
Transmit Power	Max 21 dBm
Bitrates (10MHz)	3, 4.5, 6, 9, 12, 18, 24, 27 Mbps
Bitrates (20MHz)	6, 9, 12, 18, 24, 36, 48, 54 Mbps

TABLE I
LINKBIRD-MX: 802.11P DETAILS.

B. VoIP Application

We developed a simple simulator of VoIP traffic. The application generates IP packets with a given – user defined – application data rate and packet size. The data rate should be selected in accordance with a VoIP codec (e.g., 8 kbit/s for G.729). The packet size is a parameter used to reflect the fact that, in order to improve efficiency, the very small packets generated by a typical VoIP codec (e.g., 10 bytes for G.729) are assumed to be packed in larger packets of size P , where P is a tunable parameter and is assumed to be a multiple of the codec packet size.

The application was entirely developed in Java, version 1.6 under the OS Linux (kernel 2.6.35), and uses the Car-2-X SDK software toolkit, provided together with the LinkBirds, to realize V2V communications. In particular, the Car-2-X SDK software includes an abstracted protocol stack, providing several communication primitives designed for VANets. The

IP packet is not among the primitives included in the C2X stack, but it is encapsulated within a specific C2X message (single hop broadcast).

In accordance with VoIP literature, we assume that a VoIP call alternates between talk and silence periods, whose lengths are exponentially distributed random variables with mean τ_t and τ_s , respectively. Following [17], we set $\tau_t = 1500$ msec and $\tau_s = 2000$ msec, and imposed a lower bound of 250 msec to the talk duration to avoid unnaturally short talk periods.

In our study, we aim also at evaluating the effect on VoIP performance of coexistence with safety-related vehicular applications. To this purpose, we have also used a simple beaconing application, which periodically sends packets (called *beacons*) reporting vehicle status data for safety purposes. In accordance with recommendations for active safety applications [18], we set the beaconing period to 100 msec and the beacon packet size to 80 bytes. The VoIP application can be set to run in isolation, or concurrently with the beaconing application.

During the simulation, the application logs, for each transmitted packet, its progressive number, the GPS time, and the vehicle status (speed, latitude, and longitude). Similar information is logged at the receiver side upon each packet reception. Recording of GPS time at both sides of the communication channel allows time alignment of the sender and receiver traces, which, in turn, enables computation of the VoIP performance parameters (*jitter* and *throughput* as defined below), and of the relative distance between vehicles.

C. Performance parameters

VoIP application performance parameters will be measured in terms of throughput and jitter. The throughput measures the number of bits per second transmitted/received by the VoIP application, and it is computed by dividing the duration of an experiment into time intervals of 1 second, and by recording the number of packets transmitted/received in each interval. Since the throughput is typically expressed in bits per seconds, we multiply the number of packets transmitted/received in a one second time interval by the packet size. Furthermore, we compute the ratio between the throughput at the transmitter and at the receiver end of the communication, in order to better understand the V2V channel behavior. Note also that in a VoIP context, the duration of the silence periods strongly influences the throughput. In particular, a low/zero throughput recorded at the receiver side might be due to a silence in transmission, or to a bad communication quality. To better distinguish between these two situations, we have also computed the *real* receiver-side throughput, by considering only those time intervals where the transmitter was not silent for the entire duration of the interval.

According to the real-time transport protocol (RTP) - RFC 3550 [6], the interarrival jitter is an estimate of the statistical variance of the RTP data packet interarrival time, measured in timestamp units and expressed as an unsigned integer. The interarrival jitter J is defined to be the mean deviation (smoothed absolute value) of the difference D in packet spacing at the receiver compared to the sender for a pair of packets. If S_i is

Test	PhDR	Pkt Size	ADR	Sender Position	Beaconing
1	3	80	8	Standard	Off
2	3	160	8	Standard	Off
3	3	320	8	Standard	Off
4	3	640	8	Standard	Off
5	3	80	8	Inverted	Off
6	3	160	8	Inverted	Off
7	3	160	8	Standard	On
8	3	80	8	Standard	On
9	6	160	8	Standard	Off
10	3	160	64	Standard	Off

TABLE II

FOR EACH TEST WE REPORT: PHYSICAL DATA RATE (PHDR - MBS), PACKET SIZE (B), APPLICATION DATA RATE (ADR - MBS), VEHICLE CONFIGURATION, AND BEACONING APPLICATION STATUS.

the RTP sending timestamp for packet i , and R_i is the time of arrival of packet i at the receiver expressed in RTP timestamp units, then for two packets i and j , D can be expressed as: $D(i, j) = (R_j - R_i) - (S_j - S_i) = (R_j - S_j) - (R_i - S_i)$, which is equivalent to the difference in the “relative transit time” for the packets i and j . The interarrival jitter is computed continuously as each data packet i is received, using this difference D for that packet and the previous packet $i - 1$ in order of arrival (not necessarily in sequence), according to the formula: $J(i) = J(i - 1) + (|D(i - 1, i)| - J(i - 1))/16$. This algorithm is the optimal first-order estimator and the gain parameter $1/16$ gives a good noise reduction ratio while maintaining a reasonable rate of convergence [19]. In the following, the jitter is meant to be the average jitter, expressed in milliseconds, computed within each one second interval.

III. VOIP MEASUREMENTS

A. Data Collection

We performed both static (in the lab) and dynamic (on the road) measurements. The outdoor measurements were performed on a 4 lanes road (two lanes for each direction) surrounding the CNR Research Area in Pisa. The road segment used for the experiments is delimited by two roundabouts, and has a length of approximately 1.5 Km . In order to investigate the degree of correlation between vehicle speed and VoIP performance, each experiment was composed of 5 sets of measurements, each performed by traveling along the selected road segment in both directions at a fixed speed v . The considered values of v in each experiment were $v = 50, 60, 70, 80$ and 90 Km/h .

In order to investigate the effects of PHY and application layer parameters on VoIP performance, we have performed a total of 10 experiments, whose main parameters are reported in Table II. Most of the experiments were performed taking G.729 as the reference codec, with a 8 Kbit/s data rate and varying packet sizes. In a final experiment, we also considered a less compressed audio codec such as G.711, which has a 64 Kbit/s data rate. In most cases, the transmitter vehicle is in front, and the receiver in back position. In order to investigate reciprocity of the V2V channel, we have also performed two experiments in which the positions of transmitter and receiver were inverted. During the experiments, vehicles mostly remained in the typical car following situation, empirically trying to vary the inter-vehicle distance while remaining

Test	Jitter (msec)			Throughput Receiver (bit/s)		
	Min	Max	Avg	Avg	Real Avg	Ratio
1	3.44	458.35	3.78	1596.14	4354.12	0.32
2	4.42	850.04	10.98	3069.12	4634.70	0.58
3	5.45	902.18	18.15	3651.86	4718.53	0.67
4	4.18	917.67	65.58	3331.59	4380.57	0.62
5	2.22	695.35	7.90	1609.92	4079.0	0.33
6	2.60	845.12	13.29	2601.41	4641.45	0.50
7	2.01	746.88	19.99	2146.38	4384.71	0.41
8	4.07	477.05	7.74	1590.78	4512.46	0.31
9	3.04	861.60	15.44	1526.15	4719.99	0.27
10	1.039	475.69	0.90	9795.67	33758.24	0.23

TABLE III
VOIP PERFORMANCE RESULTS.

within transmission range. However, since experiments were performed with normal traffic conditions, in many cases other vehicles were in between the transmitter and receiver vehicle, possibly leading to a degraded communication quality due to NLOS conditions.

B. Results

Table III reports, for each test, the minimum, maximum and average jitter. The table also reports the average receiver throughput, the average real receiver throughput, and the ratio between the sender and the receiver throughput, which can be considered as a good approximation of the average packet delivery rate recorded in the experiment.

It is interesting to observe that, while the average jitter always remained well below the 100 msec value considered to be acceptable for (buffered) VoIP communications, the maximum jitter value was always well above this threshold. However, relatively few jitter peaks above 100 msec were observed in the experiments – see Figure 2. On the other hand, VoIP performance in terms of throughput was unsatisfactory: in all the experiments, the packet delivery rate was well below 100%, with values as low as 23%. This is clearly unacceptable for VoIP applications, which are designed to tolerate only minimal packet losses (in the order of 1% at most). It is worth observing, though, that our experiments were performed without using any transport-layer packet retransmission mechanism such as TCP, since these mechanisms are known to perform very bad in challenging wireless environments such as VANets.

Varying packet size. The effect of varying packet size on VoIP performance can be evaluated by comparing results of tests 1–4. As expected, increasing packet size tends to increase jitter, due to the relatively larger delay variation induced by a relatively larger packet size. On the other hand, throughput performance seems to be relatively oblivious to packet size, except for the smaller packet value. In this case (test 1), the experienced throughput was considerably worse than with larger packet size – see Figure 2. To investigate the reasons of this performance degradation with smaller packet size, we have performed a number of experiments in the lab, in idealized conditions for what concerns communication quality. We have verified that performance degradation is due to the inability of the LinkBird units to keep the pace with the encapsulation/decapsulation operations required to send an IP




Fig. 1. Packet flow in our VoIP experiments.

packet over the VANet¹. This can be seen from the results in Table IV, reporting the losses experienced in the various legs of the communication path between the VoIP sender and receiver application (see Figure 1): packets are lost in the LinkBird units at both ends of the communication link. Since relatively more packets are lost for relatively smaller packet size, our belief is that losses are due to encapsulation/decapsulation operations performed by the C2X stack running on the LinkBirds: with smaller packet sizes, relatively more packets per second are managed by the LinkBird units (notice, in fact, that the application layer data rate is fixed), and relatively more losses are experienced.

Varying application data rate. The effect of varying the application data rate (Figure 3) is obtained by comparing test 2 with test 10. Although the values of jitter are lower in case of 64kb/s than with lower data rate, the ratio between transmitter and receiver throughput is sensitively higher in case of 8kb/s. As observed above, this is likely due to the inefficiency of encapsulation/decapsulation of IP packets performed in the LinkBird units. Our results clearly show that more compressed audio codec are preferable in vehicular environments.

Varying physical data rate. In test 9, the PHY data rate has been changed to 6Mbps, leaving all the other parameters unchanged w.r.t. test 2. In this case, both jitter and throughput were negatively affected (see Table III), revealing a worse performance of 6Mbps PHY data rate w.r.t. 3Mbps data rate.

Beaconing and VoIP coexistence. The beaconing application has been activated during tests 7 and 8, and can be compared with, resp., tests 2 and 1. The results in Table III show that the coexistence of the two applications degrades VoIP performance in terms of both jitter and throughput, most noticeably with packet size of 160B.

Inverting vehicle positions. Comparing the results of tests 1 and 2 with those of tests 5 and 6, we can conclude that inverting vehicle position has only a marginal effect on VoIP performance, thanks confirming the reciprocity of the V2V channel already observed in the literature – see, e.g., [4].

Correlation with vehicles’ distance and speed. Table V reports, for each test, the correlation of jitter and throughput w.r.t. the vehicles’ distance and speed. As expected, there is a clear correlation between inter-vehicle distance and VoIP application for both jitter and throughput. Less expected was the correlation between VoIP performance metrics and speed, which is comparable to the distance correlation.

IV. DISCUSSION AND FINAL REMARKS

The study reported in this paper discloses several interesting insights into VoIP performance in VANets. The most important

¹Note that in previous on-the-road and in-lab measurements performed sending single hop broadcast packets [4], the Linkbirds were able to support a much larger number of pkt/sec.

Pkt Size (B)	Packets sent/received through links:						
	A	B	C	D	E	F	G
80	320/320	320/320	219/320	219/219	209/219	209/209	209/209
160	146/146	146/146	121/146	121/121	111/121	111/111	111/111
320	112/112	112/112	88/112	88/88	77/88	77/77	77/77
640	24/24	24/24	24/24	24/24	16/24	16/16	16/16

TABLE IV
PACKETS PASSING THROUGH THE FLOW DEPICTED IN FIGURE 1.

Test	Jitter vs		Throughput vs	
	Distance	Speed	Distance	Speed
1	0.14	0.19	0.67	0.69
2	0.43	0.17	0.43	0.63
3	0.39	0.16	0.38	0.54
4	0.52	0.33	0.29	0.47
5	0.33	0.18	0.55	0.72
6	0.50	0.25	0.44	0.64
7	0.38	0.29	0.52	0.62
8	0.19	0.29	0.70	0.73
9	0.39	0.28	0.57	0.69
10	0.16	0.24	0.73	0.70
Avg	0.34	0.24	0.53	0.64

TABLE V
CORRELATION OF JITTER AND THROUGHPUT WITH DISTANCE AND SPEED.

observation is that performance is unsatisfactory, especially in terms of throughput. As we have shown in the paper, this is likely be due to the fact that IP packets are not “native” packets for vehicular networks, but they are encapsulated within “native” vehicular network packets. In the specific case of our experiments, IP packets were encapsulated within single hop broadcast packets provided by the C2X protocol stack. We believe that, although our experiments were performed with a specific hw/sw solution for vehicular networking (LinkBird units with C2X protocol stack), similar results are likely to be obtained with other solutions as well. In fact, IP-based protocol stack will not be the primary stack for communication in VANets, due to its inability of successfully deal with the requirements of active safety applications which will be predominant in vehicular environments.

Given the above, we believe significant progresses are still to be made before VoIP can be considered a viable solution in VANets. Progresses might consist in improving codec technology, especially for what concerns tolerance to packet losses and even lower data rates, as well as in improving efficiency of IP implementation in vehicular networking protocol stacks.

REFERENCES

- [1] L. Cheng, B. Henty, R. Cooper, D. Stancil, and F. Bai, “Multi-path propagation measurements for vehicular networks at 5.9 ghz,” in *IEEE Wireless Communications and Networking Conf.*, 2008, pp. 1239–1244.
- [2] L. Cheng, B. Henty, F. Bai, and D. Stancil, “Highway and rural propagation channel modeling for vehicle-to-vehicle communications at 5.9 ghz,” in *IEEE Symp. on Antennas and Propagation*, 2008, pp. 1–4.
- [3] I. Tan, W. Tang, K. Laberteaux, and A. Bahai, “Measurement and analysis of wireless channel impairments in dsrc vehicular communications,” in *IEEE Int. Conf. on Communications*, 2008, pp. 4882–4888.
- [4] F. Martelli, M. Renda, and P. Santi, “Measuring IEEE 802.11p performance for active safety applications in cooperative vehicular systems,” in *73rd IEEE Vehicular Technology Conf.*, 2011, pp. 1–5.
- [5] A. Vinel, D. Staehle, and A. Turlikov, “Study of beaconing for car-to-car communication in vehicular ad-hoc networks,” in *VehiMob Workshop*, Dresden, Germany, 2009.
- [6] H. Schulzrinne, S. Casner, R. Frederick, and V. Jacobson, “RTP: A transport protocol for real-time applications. RFC 3550,” July 2003. [Online]. Available: <http://tools.ietf.org/html/rfc3550>

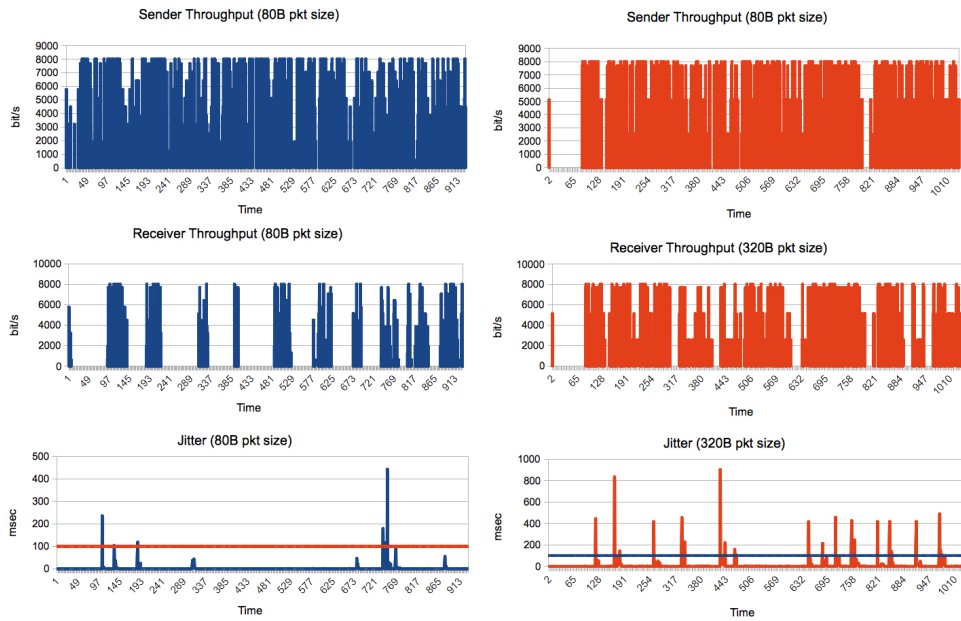


Fig. 2. Effect of varying the packet size: 80 and 360 bytes.

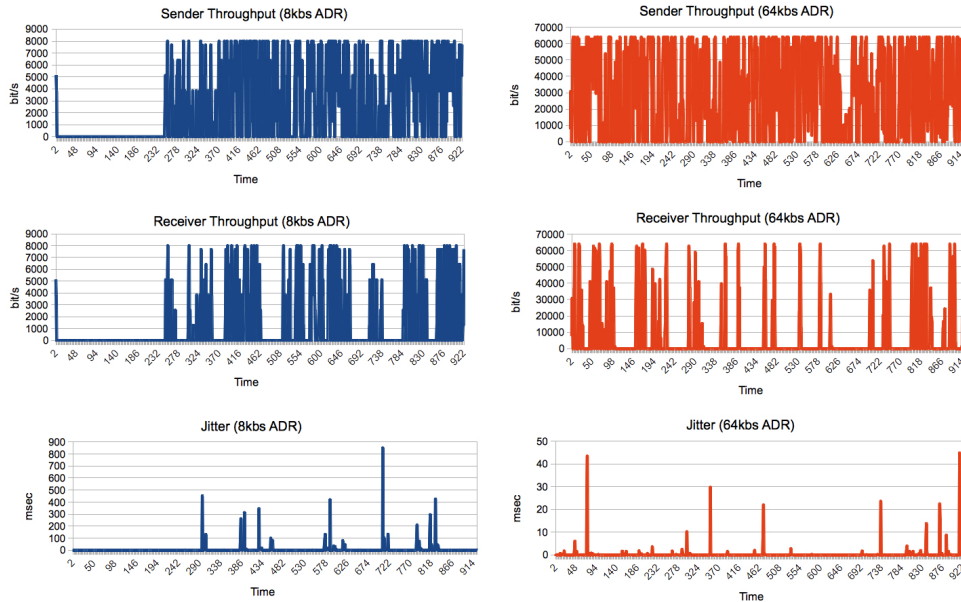


Fig. 3. Effect of varying the application data rate: 8 and 64 kb/s.

- [7] M. Anis, M. Khan, J. Inam, and R. Azmatullah, "Comparison between the IEEE 802.11x standards of voip using the CBR and VBR voice schemes," in *IEEE Radio and Wireless Symp.*, 2008, pp. 207–210.
- [8] W. Wang, S. C. Liew, and V. Li, "Solutions to performance problems in VoIP over a 802.11 wireless LAN," *IEEE Trans. on Vehicular Technology*, vol. 54, no. 1, pp. 366–384, Jan. 2005.
- [9] K. Stoeckigt and H. Vu, "VoIP capacity – analysis, improvements, and limits in IEEE 802.11 wireless LAN," *IEEE Trans. on Vehicular Technology*, vol. 59, no. 9, pp. 4553–4563, Nov. 2010.
- [10] Y. Jeong, S. Kakumanu, C.-L. Tsao, and R. Sivakumar, "VoIP over Wi-Fi networks: Performance analysis and acceleration algorithms," *Mobile Networks and Applications*, vol. 14, pp. 523–538, 2009.
- [11] S. Garg and M. Kappes, "An experimental study of throughput for UDP and VoIP traffic in IEEE 802.11b networks," in *IEEE Wireless Communications and Networking*, vol. 3, March 2003, pp. 1748–1753.
- [12] S. Shin and H. Schulzrinne, "Measurement and analysis of the VoIP capacity in IEEE 802.11 WLAN," *IEEE Trans. on Mobile Computing*, vol. 8, pp. 1265–1279, Sept. 2009.
- [13] G. Kaosar, T. Sheltami, and A. Mahmoud, "Delay and power efficient voice transmission over MANET," in *IEEE Int. Conf. on Communications*, 2007, pp. 3313–3318.
- [14] E. Thibodeau, M. Youssef, and A. C. Houle, "Investigating MaNet performance in a VoIP context," in *Conf. on Electrical and Computer Engineering*, 2006, pp. 920–923.
- [15] L. Atzori, F. Boi, and G. Nonnis, "IP telephony over mobile ad hoc networks: Joint routing and playout buffering," in *IEEE Int. Conf. on Communications*, 2008, pp. 1861–1865.
- [16] S. Armenia, L. Galluccio, A. Leonardi, and S. Palazzo, "Transmission of VoIP traffic in multihop ad hoc IEEE 802.11b networks: experimental results," in *1st Int. Conf. on Wireless Internet*, 2005, pp. 148–155.
- [17] K. Sriram and W. Whitt, "Characterizing superposition arrival processes in packet multiplexers for voice and data," *IEEE J. Selected Areas in Communications*, vol. 4, no. 6, pp. 833–846, 1986.
- [18] J. Kenney, in *VANET: Vehicular Applications and Inter-Networking Technologies* (Editors: H. Hartenstein and K. Laberteaux). Chichester, UK: John Wiley and Sons, 2010.
- [19] J. A. Cadzow, *Foundations of digital signal processing and data analysis*. New York, USA: Collier Macmillan, 1987.