Simulation Study of Cross-Layer Wi-Fi Rate-Adaptation Mechanisms for Routing Algorithms with Realistic Vehicular Traces

Olivier Mehani
University of New South Wales, Australia, and
Mines ParisTech, CAOR — Centre de Robotique, Mathématiques et Systèmes, France,
olivier.mehani@student.unsw.edu.au

***Abstract*—In the past decade, two research topics in communication have received increased attention. Both cross-layer designs and vehicular communications are currently hot topics in telecommunication research and applications. However, it is not certain whether the former can provide improvements in real environments. In this paper, we propose to simulate a vehicular ad-hoc network (VANET) based on realistic traces, and study the impact of cross-layer rate adaptation mechanisms for IEEE 802.11 on common ad-hoc routing protocols. We find that enhanced MAC protocols improve quite a few metrics of interest, but sometimes fail to provide a better end-to-end throughput, even reducing it by an non-negligible amount.**

***Index Terms*—cross-layer MAC rate adaptation, ad-hoc routing, VANET, simulation**

I. Introduction

With the increasing availability of off-the-shelf IEEE 802.11-based wireless communication devices, it has become increasingly desirable create networks of vehicles. The applications of such networks are manifold: infotainment, driver information and assistance or even autonomous vehicles collaborating with each other to improve their knowledge of the environment.

Due to the inherent mobility of road vehicles, standard centralised wireless access modes are usually replaced by self-organising ad-hoc networks. The Vehicular Ad-hoc NETworks (VANET) are thus a subset of the wider Mobile Ad-hoc NETworks (MANET). In order to maintain IP connectivity, specific MANET routing protocols have been proposed that propagate reachability information between physical neighbours. Each node updates their local routing tables based on this information.

VANETs do not present harder conditions at the network layer only. The high dynamicity of the vehicles in a wide range of environments a challenge at the link and physical layers as well. Indeed, the connection time between two nearby nodes may only last for few tens of seconds, mush less, even, if the cars are driving along opposite lanes. Also, as they move continuously, channel conditions are likely not to remain the same, and be influenced by external factors such as tunnels, trees or other vehicles.

It becomes necessary to make use of the network conditions to the best of what they can support, as soon as and for as long as they can. The idea of crossing layers to pass more information and allow a better adaptation to the current conditions can be applied to that effect. Several solutions have been proposed, linking the two lower layers of the OSI stack, to improve the use of the currently available resources or conditions.

We intend to evaluate by simulation the combination of these techniques and their impact on the performance of a VANET. More specifically, we are interested in how relevant cross-layers designs are in more complex, real situations than the tightly controlled environments and scenarios their authors designed them for. Indeed, in [2], Kawadia *et al.* showed that there could exist situations in which bad interactions lead such designs into highly inefficient. They noticed even poorer performances than the non-optimised protocols in many cases.

In the following, we investigate such scenarios to confirm whether or not such bad interactions are likely to happen or could be safely ignored. Before presenting our simulation setup in section III, we give an overview of the techniques involved in section II. The results are presented in section IV. We finally summarise and conclude this work in section V.

II. Related Work

The proposed evaluation is based on various researches spanning the three lower layers of the OSI stack. This section introduces those in use throughout this paper and gives a short overview of their characteristics.

A. Wireless Channels

Networks using wired links are easy to simulate as the dedicated medium guarantees proper reception of any sent packet. The situation is very different in the case of wireless communication. Here, we review the most common wireless propagation models used for simulation.

This paper was written in the context of an assignment in the UNSW TELE9756 Course on Advanced Networking [1]. It mostly aims at presenting the choice and use of a network simulation tool. The emphasis here is on the simulation approach and setup. The results given here should only be re-used with caution, if at all.

The air medium has a higher attenuation than the usual copper cable and obstacles along the line-of-sight (LoS) between sender and receiver further reduce the strongest component of the signal. Additionally, as the communication is made in a broadcast manner — the signal is sent all around the transmitting implement — there is the problem of multipath fading: several component of the signal, being reflected on various features of the environment, further reduce the receive signal strength, thus reliable long range communication. In this context, it is important to model the wireless channel as accurately as possible [3].

1) Free Space: The first model to be used was assuming a simplistic circular coverage area. Modelling a free space environment, the attenuation increases with the distance d from the sender. The received power is

$$P_r(d) = \frac{P_t G_t G_r \lambda^2}{(4\pi)^2 d^2 L}, \tag{1}$$

where P_t is the transmitted signal power, G_x the antenna gains (usually taken equal to 1), λ the wave length and L the system loss (≥ 1).

2) Two-ray Ground Reflection: Obviously, the previous model doesn't take reflections into account. Thus, the two-ray ground reflection model was introduced. As its name hints, it considers two paths: a direct LoS ray, and an additional path reflecting on the ground. It is expressed in terms of the heights of the antennas h_x as

$$P_r(d) = \frac{P_t G_t G_r h_t^2 h_r^2}{d^4 L}. \tag{2}$$

The attenuation with this model increases faster with distance than (1) and is more accurate for long distances. However, due to the varying effect of the rays combinations for small distances, the Free Space model is to be preferred for short ranges. The reader can refer to [3] for more details about this cross-over distance.

3) Probabilistic Shadowing: The above two models are deterministic, they do not account for random fading of the signal as it is observed in real measurements. The shadowing model, however, represents fading as a random variable. This model decomposes the attenuation in two terms: the mean signal strength, called the *path loss*, and its variation as a Gaussian variable with zero mean, $X_{\text{dB}} = \mathcal{N}(0, \sigma^2_{\text{dB}})$. Its general expression with respect to a known attenuation at distance d_0 and the path loss exponent β is

$$\left[\frac{P_r(d)}{P_r(d_0)}\right]_{\text{dB}} = -10\beta \log\left(\frac{d}{d_0}\right) + X_{\text{dB}}. \tag{3}$$

4) Ricean Fading: This model takes into account the effect of multipath fading on the main signal component. It accounts, in more general a way than the Free Space model, for the various reflected copies of the signal which may randomly cancel each other. The Ricean model is parametrised by the Ricean factor K which is the ratio of the power in the main signal component to the reflected ones. The PDF of the received signal power can be written as

$$P[P_r] = \frac{(1+K)\mathrm{e}^{-K}}{\bar{p}} \exp\left(\frac{1+K}{\bar{p}} p\right) I_0\left(\sqrt{\frac{4K(1+K)p}{\bar{p}}}\right), \tag{4}$$

where $\bar{p}$ is the local mean of the power and $I_0(\cdot)$ is a Bessel function.

In 2000, [4] presented a way to implement Ricean fading into packet-based simulators. Due to its relative accuracy to model channel fading conditions at small time and space scales, this model would be the best suited to present a realistic physical channel to our simulated adaptive MACs.

B. Wi-Fi MACs

This section reviews Media Access Control (MAC) of Wi-Fi networks. More specifically, it focuses on the IEEE 802.11 standard and some proposed improvements to better adapt the rate to the physical conditions of the channel.

1) Standard IEEE 802.11: The 802.11 standard is based on works on the Wavelan II wireless card [5]. It standardises the physical and link layers over a wireless medium. However, the Wavelan II was operating in the 900 MHz band variants of the 802.11 use frequencies in the 2.4 GHz (802.11b/g), 5 GHz (802.11a) unlicensed band [6]. The more recent 802.11p amendment makes use of the 5.9 GHz licensed band for vehicular communications. At the link layer, message formats are given to implement a CSMA/CA[1] MACs.

Several modes of operation can be used In the *managed mode*, all communication goes through a station called the *access point*, which forwards MAC frames to the Layer-2 destination. The *ad-hoc mode* removes the need for this intermediary, with each station directly communicating with their L2 neighbours. They are more suited for dynamic networks, hence MANETs and VANETs.

Despite initially operating at a single rate of 1 or 2 Mbps at first, the standard supports data rates up to 54 Mbps. MAC messages are always sent at the base rate to ensure a better coverage. It also allows proper synchronisation of the stations before negotiating a higher rate for data packets.

Before sending any packet, the MAC listens on the channel for transmissions from other nodes. If the medium is busy, the station backs off for a random amount of time before repeating the process. Though this mechanism works properly for small networks where all nodes are in range of each other, problems appear in when there are "hidden" stations.

The hidden node problem occurs when a station B has a packet for A, which is in its range, but can't hear that a node C, which isn't, is already transmitting data to A. In this case, a packet collision occurs at A, which can't decode any of the two packets it receives at once. The RTS/CTS mechanism of 802.11 is designed to overcome this problem.

[1] Carrier Sense Multiple Access with Collision Avoidance

When a packet is ready to be sent, a node can send a Request To Send (RTS) MAC frame informing the receiver of the time the medium will be occupied during this transmission. If the medium is clear, the receiver sends a Clear To Send (CTS) packet. This CTS packets propagates the information to all nodes in range of the receiving node, even if they can't hear traffic from the sending node. As the CTS also carries information about the media occupation, these neighbouring nodes learn haw long they have to back off for, thus avoiding any collision during this transmission.

When a multirate mode is selected, each node determines their sending rate according to how their previous transmissions went. The most commonly used rate selection algorithm is the Auto Rate Fallback (ARF): if the last 10 packets where transmitted properly, regardless of the destination, the next higher data rate is selected. If more than 2 packets transmission have failed (as noticed due to the lack of acknowledgements because *e.g.*, the signal wasn't strong enough on the receiver's side), the data rate is reduced to the next lower setting [5].

As mentioned, this data rate adaptation is done locally in each node, and regardless of the destination. Thus, the network performance may be pulled down as soon as a single node experiences bad channel conditions. The two following cross-layer adaptations of 802.11 have been proposed to overcome this issue.

2) Receiver-Based Auto Rate Control: The Receiver-Base Auto Rate control (RBAR) [7] builds on top of the CTS/RTS mechanism of 802.11 to adjust the rate to the channel conditions. It however requires an incompatible modification of the standard. It needs replacing the *Duration* field of these MAC frames with information on the chosen data rate and frame size.

When a packet is ready to be sent, the sender sends its RTS frame with the data rate it has chosen. The receiver monitors the signal strength at which it received the request. Depending on the channel conditions, it can decide whether the transmission can be made at a higher rate or should be slowed down. The receiver sends its CTS accordingly.

A new Reservation Sub-Header (RSH) is introduced, to be sent by the sender before its data packet to confirm the new setting. Based on the data rate and frame size contained in the CTS and RSH, all neighbouring nodes can compute the time period during which the medium will be occupied and they have to back off, thus preserving the initial protection against the hidden node problem.

RBAR increases the network performance by allowing condition-adaptive per-station rate control. It allows using the medium at the best possible rate between each node, thus reducing the transmission time and the risk of decoding errors upon reception. This proposal however doesn't maintain equal air-time access between nodes. Thus, the performance improvement will be noticed at a global scale, with nodes with less connectivity having more time to communicate, but not locally, with node with a good signal to each other only sending their packets more quickly before relinquishing the channel.

3) Opportunistic Auto Rate Control: Based on the previous proposal, the Opportunistic Auto Rate control (OAR) [8] restores the air-time equality. It is based on the assumption that the channel coherence time is longer than that needed to send a packet. Thus, receiver with a good channel should use it to send more packets if they can.

In essence, OAR adapts the data rate as RBAR does using RTS/CTS/RSH. In contrast, if the data rate has been increased, the sender preserves the initial duration of the transmission by increasing the data size in its CTS frame.

If the sender has more packets to send, it can send them back-to-back. The *more fragments* flag is used on any non-final packets that more are to be expected. Each fragment is acknowledged immediately, allowing the receiver to both update the rate if the channel conditions have changed and inform its neighbour about the continuing medium occupation.

OAR allows a pair of nodes with good channel conditions to keep hold of the medium in order to transmits more data, thus increasing their network performance without impacting that of other nodes by holding the medium more than would be fair. As for RBAR, the adaptation also allows to avoid transmitting frames at a rate higher than the current conditions support. Thus, less retransmissions can be expected.

C. MANET Routing Algorithms

We now overview some MANET routing protocols that have been used for VANETs. Table-based ad-hoc routing algorithms can be separated into two classes: *proactive* and *reactive*.

Proactive protocols are quite similar to those used to exchange routes between or within Autonomous Systems on the Internet. They maintain an up-to-date map of the topology at all time by periodically exchanging control messages. On the contrary, reactive protocols only search for a route to a destination for which there are packets to be sent.

This section presents one protocol of each class, that will be evaluated in the rest of this paper.

1) Ad-hoc On-Demand Vector: The Ad-hoc On-Demand Vector protocol (AODV) [9] is one of the reactive class of protocols. When the network layer receives a packet for a destination it doesn't know a path to, it broadcast a *Route Request* (RREQ).

Each node receiving an RREQ can behave in two ways. Either they do not have a routing table entry for the destination, in which case they rebroadcast the request after a random time during which they didn't hear it retransmitted. If they know a route to the destination, they send a *Route Reply* (RREP) back to the sender informing them of the path. All nodes along the path update their routing tables according to the RREP. The next packets are then sent along the discovered path.

AODV also uses periodical unsolicited *Hello* messages to keep the local neighbourhood information updated.

2) Optimized Link State Routing: The Optimized Link State Routing protocol (OLSR) [10] is from the proactive class. Each node participates to the algorithm by periodically sending *Hello* messages, in a way similar to AODV. These messages contain information about themselves and their direct one-hop neighbours.

The nodes thus learn their two-hop neighbourhood. Aggregated *Topology Control* (TC) messages are then used to relay this information further to the rest of the network. Instead of having each node take care of it (flooding), which would be quite network intensive and inefficient, this task is delegated to only these nodes that were designated as *MultiPoint Relays* (MPRs). Nodes are selected to act as MPRs, and propagate TC messages, by their one-hop neighbours to cover all their two-hop neighbours.

III. Simulation Setup

In this section, we present a simulation setup[2] which investigates how combinations of the solutions just presented, each of them separately providing improvements, perform and confirm the occurrence of the bad interactions mentioned in [2].

A. Tools and Protocols

The following simulations are made using the *ns*-2.34 simulator[3] from the VINT project. It has been chosen as it is a well know and studied simulator as well as because implementations of all the previously mentioned are available.

An implementation of standard IEEE 802.11 is already shipped with the simulator. Implementations of RBAR for *ns*-2.27[4] and OAR for *ns*-2.1b7[5] have been ported to the latest version of the simulator. The OAR patch also provides a per-channel implementation of Ricean fading.

As for 802.11, an implementation of AODV is included with the simulator. The OLSR implementation from Inria[6] has been ported from *ns*-2.27.

Next, we describe the simulation scenarios.

B. Scenarios

Vehicles have mobility patterns which can't be accurately simulated with models such as random way-points or Brownian movement. Indeed, they follow rather strict constraints, starting with having to follow roads. It is then necessary to model vehicular moves as accurately as possible. Several approaches can be taken, from experimental collection to synthetic creation of vehicular traces.

In [11], the authors presented a generator of realistic vehicular traces based on microscopic agent interactions. They generated traces based on read maps of Switzerland which they made publicly available in several formats[7]. We propose to reuse their *ns*-2-formatted traces to simulate real road traffic for the VANET. Though quite similar to [11] in the simulation approach, the novelty of our work lies in the orientation towards cross-layer MACs and the evaluation of their performance when coupled to specific routing algorithms.

Each vehicle in these traces has been configured as a wireless station. Communication pairs have been formed between vehicles in order of first appearance in the trace file. To evaluate the network performance, 20 nodes in each simulation were attached to a CBR[8] agent generating traffic to be sent to another peer.

Each trace studied has been used to simulate VANETs equipped of a coherent combination of one of the MAC and routing protocols. The Ricean fading model has been used to simulate the physical channel. The next section compares the performance of each combination. For plain 802.11 simulations, the data rate is set to 2 MBps.

TABLE I
Configuration of the Shadowing model in *ns*-2.

	Parameter	Urban	Highway
Desired	Per-packet success probability	75 %	
	Frequency	2.4 GHz	
	Range	400 m	500 m
Mod.	β	3.5	2.5
	σ_{dB}	6	4
ns-2	`RXThresh_`	3.92926×10^{-15}	4.68225×10^{-12}

C. Channel Model

Whereas it was first tried to use the Ricean fading model, aberrant results were observed. It was thus decided to resort to using the shadowing model instead.

The model is parametrised similarly to what is described in [11], using the `RXThresh_` parameters of the simulator's physical channel. There are two sets of parameters depending on whether the traces are urban or along highways. They are summarised in Table I.

IV. Results and Discussion

This section presents consolidated results from the simulations. In [2], the authors give the example of how, by extending the range of wireless nodes, RBAR disrupts the operation of transports when a minimal-hop routing protocol is in use. Using a lower number of hops, the routing protocol establishes longer range links which can't be operated at high data rates.

The metrics given here cover both the network and transport layers in order to determine the network performance of the various combinations of OAR and RBAR with OLSR and AODV. Simulation results with plain 802.11b are used as the baseline for comparison.

[2]The simulation scripts are available at https://scm.narf.ssji.net/unsw/browser/tele9756/assignment_a.
[3]http://nsnam.isi.edu/nsnam/
[4]http://perform.wpi.edu/downloads/rbar/
[5]http://www-ece.rice.edu/networks/software/OAR/OAR.html
[6]http://hipercom.inria.fr/OOLSR/downloads.html
[7]http://www.lst.inf.ethz.ch/research/ad-hoc/car-traces/index.html#traces
[8]Constant Bit Rate

TABLE II
ROUTING OVERHEAD [%]

		AODV avg.	AODV σ	AODV-cmp avg.	OLSR avg.
Urban	**802.11b**	54.4	19.7	13.6	39.8
	RBAR	53.7	18.4	15.1	39.8
	OAR	48.5	26.0	15.1	39.7
Hwy.	**802.11b**	72.0	18.8	not simulated	
	RBAR	75.5	17.4		
	OAR	79.0	15.6		

TABLE III
ROUTE ESTABLISHMENT DELAY [s]

		AODV avg.	AODV σ	AODV-cmp avg.	AODV-cmp σ	OLSR avg.	OLSR σ
Urban	**802.11b**	11.9	1.9	3.1	2.3	0.014	0.005
	RBAR	11.7	2.9	5.0	3.8	0.014	0.004
	OAR	10.6	3.5	5.0	3.8	0.014	0.003
Hwy.	**802.11b**	11.6	2.1	not simulated			
	RBAR	12.3	1.6				
	OAR	12.4	1.7				

TABLE IV
HOP COUNT

		AODV avg.	AODV σ	AODV-cmp avg.	AODV-cmp σ	OLSR avg.	OLSR σ
Urban	**802.11b**	1.3	0.29	1.0	0.0	1.0	0.016
	RBAR	1.3	0.24	1.0	0.0	1.0	0.015
	OAR	1.2	0.22	1.0	0.0	1.0	0.015
Hwy.	**802.11b**	1.3	0.22	not simulated			
	RBAR	1.3	0.22				
	OAR	1.3	0.21				

Due to its proactive nature, OLSR generates a lot of unsolicited *Hello* packets. With the large number of simulated vehicles, this leads to a combinatorial explosion in the number of simulated packets in the air, the memory used and the overall time of the simulation. To maintain these conditions within simulable ranges, it has been necessary to focus on the trace with the smallest number of cars only — as few as 60 vehicles already took more than 4 months for only 26 s of simulated time. In addition, the propagation model had to be reversed back to Two-Ray Ground, as the more elaborate models proved too time consuming to compute for each packets in those runs. Finally, To retain some comparison possibilities on the routing algorithms, AODV simulations where also run with these parameters and are presented separately as "AODV-cmp" in the following.

A. Routing Metrics

1) Routing Overhead: The routing overhead is the ratio of the number of packets exchanged by the routing protocols to the total number of frames that were put on the networks. It gives the portion of the network traffic that was not used for data traffic. Table II compares this metric for the various protocol combinations.

It appears from these results that the use of RBAR, and that of OAR even more, reduces the routing overhead. As these protocols allow to reach further nodes, this results may be due to the fact that less signalling packets need to be forwarded, thus leaving more time to carry application data packets. As expected, the proactive nature of OLSR has a clear impact on the overhead of this protocol.

2) Route Establishment Delay: The route establishment delay is the time it takes for the first packet to reach its destination to do it. It is mostly relevant in the case of reactive routing protocols. However, in these simulations, the OLSR nodes start with no knowledge about the rest of the topology. Furthermore, the CBR traffic is enabled at the very beginning of the simulation. Table III thus gives the average time it takes AODV to find a route through the VANET, and that it takes for OLSR to learn about them.

Being able to reach hops further away from the source, both RBAR and OAR reduce the route establishment delay by allowing to find the destination of the data packets more quickly.

As for OLSR, the delay is much lower than for AODV, all other parameters kept identical. This can be explained by the fact that OLSR's proactive *Hello* messages are sent in parallel within the entire network at the same time, whereas AODV's *RREP/RREQ* have to be forwarded sequentially along the network path and back before the route can be established.

3) Hop Count: Once a route has been found, a simple metric to assess its quality is the number of HOPs. In traditional wired networks, the smaller this metric the better. This rationale changes when considering wireless networks. As outlined in [2], in radio networks with link quality decreasing with distance from the source, it may be better to establish more, shorter links. This would trade off the shortness of the path for better per-link packet error rate and modulation. Such an approach mould benefit the overall quality of the path in terms of bandwidth and jitter. Regardless of the interpretation of this metric depending on the context, it is a important to have it. It is shown in Table IV.

In these simulations, it appears that the hop count is fairly stable and close to one hop. Out of all the simulations, the longest path found was 48 hops. Both routing algorithms however did a good job of reducing the number of hops regardless of the MAC protocol.

B. End-to-End Metrics

1) Packet Delivery Ratio: Data packets can be lost for several reasons. Be it due to weak signals, lack of buffer space or expiration of their Time To Live (TTL), this is always an event to be minimised. Table V gives the ratio of the number of dropped packets to that of initially sent packets. It does not take forwarded packets into account as this would bias the results by increasing the PDR.

Changing the MAC layer to RBAR, then OAR has a beneficial impact on the overall PDR of the simulated scenario. This is likely explained by a more accurate choice of the data rate, resulting in less undecodable frames. In

TABLE V
PACKET DELIVERY RATIO [%]

		AODV		AODV-cmp	OLSR
		avg.	σ	avg.	avg.
Urban	**802.11b**	14.3	6.0	23.9	24.7
	RBAR	13.2	5.3	23.9	24.7
	OAR	14.4	6.8	23.9	24.8
Hwy.	**802.11b**	11.5	5.4	not simulated	
	RBAR	11.6	5.3		
	OAR	12.2	5.7		

TABLE VI
PER-FLOW OVERALL THROUGHPUT [BPS]

		AODV		AODV-cmp		OLSR	
		avg.	σ	avg.	σ	avg.	σ
Urban	**802.11b**	883.2	402.1	689.6	82.9	583.8	5.2
	RBAR	843.1	370.6	811.9	259.7	584.2	4.73
	OAR	840.4	343.1	811.9	259.7	586.2	5.0
Hwy.	**802.11b**	882.5	359.5	not simulated			
	RBAR	1004.6	421.2				
	OAR	1076.6	479.3				

any case, it is still well below 50 %. This could be justified by the mobility pattern of the vehicles which forces routes to be constantly updated.

2) Per-Flow Throughput: The throughput is one of the most high-level metrics, which makes it of particular importance to evaluate the performance. This is one of which that all components of a network stack try to increase, both locally and globally. Table VI lists the average of the per-flow throughputs over the full time of the connection.

The global per-flow average throughput is increased by the rate adaptive MAC protocols except in urban environments. This can be explained by the fact that there were more direct links between the vehicles that the routing protocols could use for shorter paths. This would confirm the unintended consequences of cross-layer designs exposed in [2] in this case. In the highway simulations, the average throughput was increased by almost 25 %. The shorter OLSR and AODV-cmp simulations also show an increase.

V. CONCLUSION

We have ported various cross-layer MAC rate control enhancements and ad-hoc routing protocols to the latest version of *ns*-2. All combinations of these have been evaluated in the context of VANETs using precomputed realistic traces to describe the mobility patterns.

The main objective of this study was to provide some insight into the statement from [2] that unintended interactions of cross-layer designs, breaking common assumptions on which other layer algorithms are built, could ultimately decrease the performance of the system they were introduced to improve. The trend of our results confirms such an impact on the throughput of simulated CBR data flows, but only in urban environments. All other considered metrics were still found to be improved by the cross-layer MACs.

These simulations however proved to take much longer than what was usually expected, which reduced the latitude to run more experiments to study a specific behaviour or obtain more statistically significant results.

Future work should thus start by adapting the simulations scenarios so that they can be run on more manageable time scales. Additionally, only overall means were derived in this paper. More attention to studying the correlation of metrics is also expected to raise more insight into the impact of cross-layer MAC protocols on data flows. This would allow to focus only on those which exhibit unexpected behaviours in order to confirm their causes.

REFERENCES

[1] R. Malaney, "TELE9756 — Advanced networking course," Sydney, NSW, Australia, Feb. 2008.

[2] V. Kawadia and P. R. Kumar, "A cautionary perspective on cross-layer design," *IEEE Wireless Communications*, vol. 12, no. 1, pp. 3–11, February 2005. [Online]. Available: http://dx.doi.org/10.1109/MWC.2005.1404568

[3] "Radio propagation models," in *The ns Manual (formerly ns Notes and Documentation)*, K. Fall and K. Varadhan, Eds., January 2009, ch. 18, pp. 178–182. [Online]. Available: http://www.isi.edu/nsnam/ns/doc/node216.html

[4] R. J. Punnoose, P. V. Nikitin, and D. D. Stancil, "Efficient simulation of Ricean fading within a packet simulator," in *VTC 2000-Fall, 52th IEEE Vehicular Technology Conference*, vol. 2, IEEE Vehicular Technology Society. Piscataway, NJ, USA: IEEE Computer Society, September 2000, pp. 764–767. [Online]. Available: http://dx.doi.org/10.1109/VETECF.2000.887108

[5] A. Kamerman and L. Monteban, "WaveLAN-II: A high-performance wireless LAN for the unlicensed band," *Bell Labs Technical Journal*, vol. 2, no. 3, pp. 118–133, August 1997. [Online]. Available: http://dx.doi.org/10.1002/bltj.2069

[6] IEEE Std 802.11-2007, "IEEE standard for information technology — Telecommunications and information exchange between systems — Local and metropolitan area networks-specific requirements — Part 11: Wireless LAN medium access control (MAC) and physical layer (PHY) specifications," LAN/MAN Standards Committee, New York, NY, USA, June 2007. [Online]. Available: http://dx.doi.org/10.1109/IEEESTD.2007.373646

[7] G. Holland, N. Vaidya, and P. Bahl, "A rate-adaptive MAC protocol for multi-hop wireless networks," in *MobiCom 2001, 7th Annual International Conference on Mobile Computing and Networking*, ACM SIGMOBILE. New York, NY, USA: ACM, July 2001, pp. 236–251. [Online]. Available: http://dx.doi.org/10.1145/381677.381700

[8] B. Sadeghi, V. Kanodia, A. Sabharwal, and E. Knightly, "OAR: An opportunistic auto-rate media access protocol for ad hoc networks," *Wireless Networks*, vol. 11, no. 1-2, pp. 39–53, January 2005. [Online]. Available: http://dx.doi.org/10.1007/s11276-004-4745-x

[9] C. E. Perkins and E. M. Royer, "Ad-hoc on-demand distance vector routing," in *WMCSA 1999, 2nd IEEE Workshop on Mobile Computing Systems and Applications*. Los Alamitos, CA, USA: IEEE Computer Society, February 1999, pp. 90–100. [Online]. Available: http://dx.doi.org/10.1109/MCSA.1999.749281

[10] P. Jacquet, P. Mühlethaler, T. Clausen, A. Laouiti, A. Qayyum, and L. Viennot, "Optimized link state routing protocol for ad hoc networks," in *INMIC 2001, IEEE International Multi Topic Conference*. Los Alamitos, CA, USA: IEEE Computer Society, December 2001, pp. 62–68. [Online]. Available: http://dx.doi.org/10.1109/INMIC.2001.995315

[11] V. Naumov, R. Baumann, and T. Gross, "An evaluation of inter-vehicle ad hoc networks based on realistic vehicular traces," in *MobiHoc 2006, 7th ACM international symposium on Mobile Ad Hoc Networking and Computing*, ACM SIGMOBILE. New York, NY, USA: ACM, May 2006, pp. 108–119. [Online]. Available: http://dx.doi.org/10.1145/1132905.1132918