Experimental Analysis of Broadcast Reliability in Dense Vehicular Networks

Kishore Ramachandran and Marco Gruteser
WINLAB, Rutgers University
Tech. Center of NJ, North Brunswick, NJ 08902-3390, USA
Email: {kishore, gruteser}@winlab.rutgers.edu

Ryokichi Onishi and Toshiro Hikita
Toyota InfoTechnology Center
One Telcordia Dr., Piscataway, NJ 08854, USA
Email: {onishi, hikita}@us.toyota-itc.com

***Abstract*—Dedicated Short Range Communications (DSRC)-based communications enable novel automotive safety applications such as an Extended Electronic Brake Light or Intersection Collision Avoidance. These applications require reliable wireless communications even in scenarios with very high vehicle density, where these networks are primarily *interference-limited*. Given the uncertainties associated with current simulation models, particularly their interference models, it is critical to experimentally validate network performance for such scenarios.**

Towards this goal, we present a systematic, large-scale experimental study of packet delivery rates in a dense environment of 802.11 transmitters. We show that even with 100 transmitters in communication range with a frame size of 128 bytes and a bit-rate of 6Mbps, (a) most receivers can decode over 1500 pps in a saturated network, which corresponds to a packet delivery rate of 45% and (b) the mean packet delivery rate, for 10 pps per node workload that emulates vehicular safety applications, is about 95%. These results demonstrate that a COTS 802.11 implementation can correctly decode many packets under collision due to physical layer capture and can serve as a reference scenario for validation of network simulators.

I. Introduction

Advances in vehicle-to-vehicle (V2V) communication enable novel safety, driver information, and entertainment applications by providing a low latency, high-capacity channel between vehicles. Safety applications such as an extended electronic brake light (EEBL) or intersection collision avoidance (ICA) promise to reduce vehicle accidents by transmitting warning messages between vehicles to notify following cars and their drivers of dangerous situations. To be effective, these applications require low-latency and highly reliability V2V communication protocols.

Messaging reliability may be affected by multipath fading, shadowing from roadside structures and other vehicles, and co-channel interference, among others. This paper concentrates on interference, whose effect is most significant in dense scenarios, for example on major highways with hundreds of Dedicated Short Range Communication (DSRC, currently under standardization) equipped cars in communication range. Evaluating messaging reliability and designing mechanisms to reduce interference and congestion in this environment, such as transmission power control, directional antennas, or admission control mechanisms, requires an in-depth understanding of MAC performance under congestion.

While several simulation-based studies have addressed this problem [1]–[5], to our knowledge no experimental validations of these results in high density scenarios exist. In particular, it remains unclear whether the effect of co-channel interference and physical layer capture (PLC) [6] is appropriately modelled in these simulations. These effects determine which stations can receive a packet when multiple senders are transmitting simultaneously. While many analytical results based on the well-known Bianchi's IEEE 802.11 saturation[1] throughput model [7] include the simplifying assumption that all colliding packets are lost, it is generally believed that tracking the signal-to-interference-plus-noise-ratio (SINR) yields the most accurate results (implemented by simulators such as Qualnet [8]). This model accepts a packet if its signal power exceeds the cumulative power of all interfering signals by a certain capture threshold. Recent experimental results with low-power sensor mote radios [9], however, suggest that this cumulative interference is not a good predictor of packet error rate. If similar observations hold for the 802.11 implementations, a model considering only the strongest interfering signal, similar in spirit to the basic ns-2 interference model, may yield more accurate results. In either case, there also exists further uncertainty about the exact capture threshold in current hardware, which will also affect performance.

New Contributions: To reduce these uncertainties, this paper presents preliminary results of an experimental validation of 802.11 MAC performance in high density vehicular networks. Although saturation throughput performance for single-hop IEEE 802.11 networks has been studied extensively [7], [10], [11], to our knowledge, this is the first large-scale experimental study of its kind that uses 100 802.11 nodes to measure performance for many-to-many broadcast applications. Using a laboratory setting with controlled interference, it allows in-depth analysis through repeatable experiments, with precisely known configurations. Key contributions of this paper include:

- Experimental analysis of saturation throughput in dense, single channel, single hop 802.11 networks with up to 100 transmitters, a packet size of 128 bytes, and a bitrate of 6Mbps. Even with significant packet collisions, most

This material is based upon work supported by Toyota Information Technology Center (ITC).

[1]Each transmitter always has a packet to send.

Fig. 1. ORBIT Testbed setup consisting of 400 small form-factor PCs with two 802.11 wireless NICs each.

receivers can correctly decode, on average, over 1500 packets per second (pps).

- Illustrating the significant effect of PLC on both throughput fairness as well as cumulative saturation throughput. Throughput predicted by models that do not account for packet capture, such as Bianchi's model does not correlate well with empirical observations since with 100 senders in saturation almost all packets collide. These results underline the importance of precisely modeling interference and PLC.
- Experimental analysis of packet delivery rates (PDR) in dense, single channel, single hop 802.11 networks with up to 100 transmitters. With a packet size of 128 bytes, a bitrate of 6Mbps, and a packet injection rate of 10 pps per sender (emulating planned safety applications), results show that the average PDR remains at about 95% even with 100 senders.

Roadmap: The remainder of this paper is organized as follows. In the next section, we explain our evaluation methodology. In section III, we characterize saturation throughput and PDR performance to emphasize the significant effect of PLC. This is followed by section IV, which presents a PDR comparison with earlier simulation results. Finally, section V concludes the paper.

II. Methodology

This work considers a many-to-many broadcast scenario in a single-hop 802.11a network with up to 100 stations. The experiments leverage the publicly accessible ORBIT testbed [12] (shown in Figure 1) to carry out systematic and controlled experiments. This testbed consists of 400 nodes (standard Linux PCs), each of which is equipped with two wireless 802.11a/b/g interfaces. The nodes are placed in a two-dimensional rectangular grid with 1m spacing (see Figure 1) and the antennas mounted on the sides.

Conducting experiments in a controlled laboratory setting (such as ORBIT) rather than a real road environment provides the following advantages. First, it allows experimentation with hundreds of stations with manageable effort. Second, it allows experiments with controlled propagation and interference characteristics. Third, experiments are repeatable, allowing easier isolation of root causes.

The limitations of this approach are that nodes remain stationary and the setup does not capture the time-varying channel characteristics of V2V communications. As such, the results can only characterize performance of a snapshot scenario, in which relative vehicle positions do not change. Further, all nodes being fairly close together, remain in communication range and no hidden nodes are present. While emulation of larger inter-node distances and hidden nodes is possible on ORBIT, in principle, we omitted them because validation of this approach requires substantial further work. Even with these limitations, we believe that the current setup is of considerable value as a reference scenario for understanding the accuracy of high density V2V simulations and adequate modeling of interference and packet capture.

The experiments use an 802.11a MAC, on which the IEEE 802.11p standard [14] under consideration is based upon. The nodes use IEEE 802.11 ad-hoc mode while broadcasting packets at a fixed 6Mbps bit-rate. We carry out multiple runs of each experiment to ensure repeatability of results. Table I provides further details on configuration parameters.

We measure packet reception rates through a set of sniffers, configured to only receive packets[2]. One sniffer is chosen next to each sender, on the node with the highest SNR link to the sender, according to calibration measurements conducted before the experiments.

Note that the only significant source of packet loss in this experiment scenario is co-channel interference from other senders within the same experiment (i.e. packet collissions). The ORBIT testbed is largely shielded from outside interference and we have confirmed during the calibration phase that packet error rates on all links in our experiments are less than 1%. Further, note that only synchronous collisions, where at least two nodes select the same slot for transmission, can be expected to occur during our experiments. Asynchronous

[2]We noticed significant packet loss in the MadWiFi [13] driver, apparently due to software bugs, when the same NIC was used for both sending and receiving packets under heavy load. Thus, we use a seperate node for packet reception. Each sniffer is in ad-hoc mode (as opposed to monitor mode) and reports per-frame information (RSSI, PHY bit-rate, recv. timestamp, frame size) using a modified version of the driver.

TABLE I
Attributes Summary for Experiments

Attribute	**Value**
Radio Nodes	1GHz VIA C3 Processor, 512MB RAM, 20GB HDD
Wireless Interfaces	2 X Atheros AR5212 based mini-PCI 802.11a/b/g
PHY/LLC/MAC Used	IEEE 802.11a @ Channel 40
PHY Link Speed	6 Mbps (RTS and MAC retries disabled)
Wireless output power	18 dBm
OS Used	Linux 2.6.18
Wireless Card Driver	MadWifi [13] and Atheros HAL v.0.9.
Antenna	Omni-directional (6dBi gain)

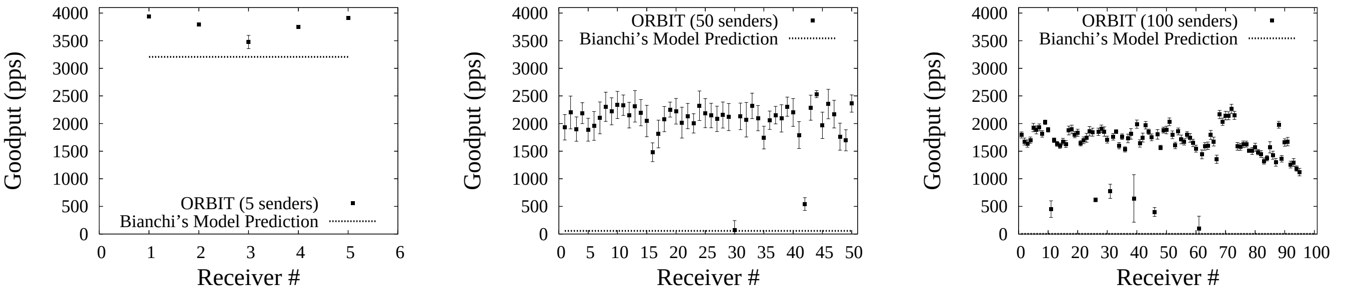

(a) Cumulative throughput (5 senders)   (b) Cumulative throughput (50 senders)   (c) Cumulative throughput (100 senders)

Fig. 2. Empirical mean and std. dev. in cumulative throughput (packets per sec.) at each receiver. Note that empirical throughput is much higher than what Bianchi's model predicts. In the case with 100 senders, the model predicts a throughput of 2 pps (not visible).

collisions, where one transmission begins while another is active, are suppressed by carrier sense.

Workload and Metrics: The experiments consider two different workloads: (i) a saturation workload, where each transmitter generates packets at the maximum possible rate and (ii) a periodic broadcast workload, where each transmitter generates 10 packets per second. In both cases, we use a modified version of the Unix *ping* utility to reliably generate small 128-byte packets (including MAC layer headers) at millisecond granularity. We choose an experiment duration of 120 seconds for each experiment configuration to ensure that each transmitter sends at least 1000 frames during the experiment. Packet size and broadcast frequency are in agreement with common assumptions about V2V safety applications.

As metrics, we choose cumulative (and per-user) goodput for the saturation workload. This allows quantification of available throughput and also enables comparison with results from Bianchi's well-known analytical saturation throughput model for the 802.11 MAC. For the periodic broadcast workload we measure the mean packet delivery rate per sender across all receivers. This metric characterizes messaging reliability in this scenario. It also illustrates how many packets were lost due to collision.

III. Saturation Throughput and Packet delivery rate in High Density Scenarios

Figure 2 presents the mean and standard deviation of cumulative saturation throughput measured at every receiver for 5, 50, and 100 senders. We also show the corresponding analytical prediction from Bianchi's model [7]. To enable this comparison, we modify the model for broadcast transmissions at 6Mbps[3] and simulate it in MATLAB [15].

Note the throughput gains compared to Bianchi's model, which are likely due to PLC [6]. Bianchi's model assumes that all frames involved in a collision are lost. With PLC, however, a receiver can decode one of the frames involved in a collision if its signal power is stronger than the other interfering transmissions. Moreover, the stronger frame is decoded even if it arrives after the other colliding frames provided it is within 128μs from the start of reception of the first received frame [16, pp. 202-203]). The existence of PLC in 802.11 radio modems has been verified experimentally [17] and it has been incorporated in analytical [6], [18] and simulation models [8], [19]. The use of Bianchi's model here serves as guidance to estimate the significance of this effect and to understand how frequently collisions occur.

Our main observations from these graphs are:

- Even with 100 transmitters, most receivers can correctly decode, on average, over 1500pps. With 50 transmitters, mean throughput at most receivers is over 2000pps. Hence, at a packet transmission rate of 10pps, the network should be able to accomodate at least 100 transmitting nodes. We investigate this further later on in this section.
- Relative to the 5 sender case (Figure 2(a)), mean throughput drops by 46% for the 50 sender case and 56% for the 100 sender case. From the analytical throughput curve, which does not take into account PLC, throughput is close to zero in both cases, meaning that almost all frames are involved in collisions. Thus, PLC recovers a frame in about 50% of the collisions in this scenario.
- Models that assume the loss of all colliding packets, such as [7], significantly underestimate maximum achievable throughput. With 100 transmitters, the mean cumulative throughput from our empirical observations is approx. 1600pps whereas Bianchi's theoretical model predicts a throughput of only 2pps.
- Throughput fairness decreases with increasing numbers of senders. This is evident in the increased variance in number of frames received across different receivers (even when not considering outliers, which may indicate software problems). Table II also reports Jain's fairness index (JFI) [20] at 2 receivers for different sets of senders.

To analyze PLC further, we compare throughput and RSSI

TABLE II
Empirical Fairness comparison for 2 receivers

Number of senders	5	10	30	50
JFI at Rcvr. 1	0.956679	0.918676	0.548513	0.383876
JFI at Rcvr. 2	0.955543	0.859743	0.619533	0.451277

[3] For each frame transmission, delays associated with SIFS, ACKs, retransmissions and exponential backoff are not used.

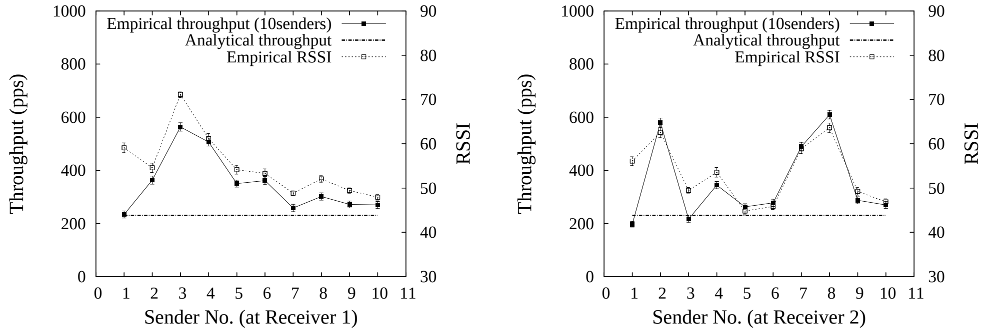

(a) Per-sender throughput and RSSI at rcvr. 1 (10 senders) (b) Per-sender throughput and RSSI at rcvr. 2 (10 senders)

Fig. 3. Empirical mean and std. dev. in per-sender throughput and RSSI at 2 receivers in the 10 sender, 10 receiver experiments. Note that there is good correlation between per-sender throughput and RSSI.

for different senders in the same experiment. Figure 3 shows the mean and standard deviation in per-sender throughput and RSSI for two receivers in the 10 sender, 10 receiver experiment.

From these results, we observe that

- a strong correlation between the per-sender throughput and RSSI exists—the sender with the highest RSSI at a particular receiver also has the highest throughput. We observe similar correlation in the 5, 30, 50 and 100 sender experiments (not shown here due to space constraints). This supports, that throughput differences can be attributed to physical layer capture.
- the throughput for the senders with an RSSI advantage shows a more significant difference between the empirical observation and the theoretical prediction.

These results highlight the importance of correctly modelling PLC in simulation scenarios with saturated channels. They also show that significant unfairness can exist even in vehicular environments that are predominantly line-of-sight.

We now consider the periodic broadcast scenario, where each sender generates a lower traffic load emulating potential safety applications. We measure packet delivery rate (PDR) across different node densities. Figure 4 presents the boxplot (max, median, min, inter-quartile range, and outliers) of PDRs for five randomly selected sample senders[4] across all receivers. Note that the majority of the receivers reach a PDR above 90% with the median PDR varying between 96% and 98.4%, depending on vehicle density (mean PDR varies between 94.4% and 97.6% and std. deviation in PDR varies between 2.4% and 5.1%). Overall, PDR even with 100 senders within transmission range remain much higher because the channel does not reach saturation.

Since Bianchi's model is only valid for saturated networks, it is more difficult to calculate collision rates and capture gains in this non-saturated scenario. For an Aloha-like[5] protocol in a similar scenario, Table III shows the expected packet delivery rates with collision and possible gains under idealized capture assumptions. The packet delivery rate with PLC is higher, similar to those obtained from our experimental results. These results imply that the negative effects of synchronous MAC collisions on the reliability of DSRC safety messages could be largely mitigated by PLC, even at high vehicular densities.

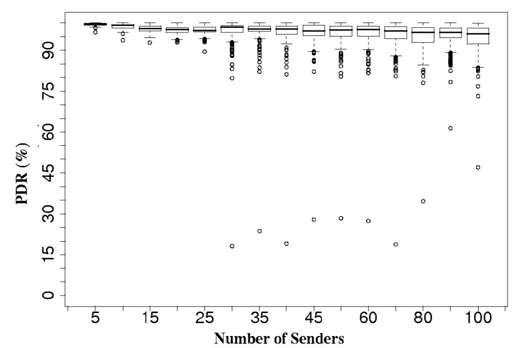

Fig. 4. Boxplot (five-number summary) of PDR for varying vehicular densities. Note that the boxplot is calculated over different time intervals and across different receivers.

IV. Discussion

To gain a preliminary understanding of how these results compare to those produced by state-of-the-art simulation models, let us compare with ElBatt and colleagues' high-density vehicular Qualnet [8] simulations [4], which best match our experimental scenario. ElBatt et al. describe a high density and a low density broadcast scenario using only a slightly larger packet size (160 bytes including all headers) and the same 10pps injection rate. The main difference between these simulations and our experimental setup lies in the spatial arrangement of the nodes, which follow a highway layout in the simulations and are separated by larger distances.

Table IV presents the vehicle density (number of vehicles in communication range) and measured packet delivery rate that we derived from the results reported in [4] for low and high density scenarios. The low density scenario, corresponds

[4]Since all senders are essentially homogenous, we expect this result to be quite similar even if we include the data from the rest of the senders.

[5]Nodes transmit for a 100 usec period at a random point in time every 100msec. All nodes are in range of each other and in the case of collision, the frame with the highest received power is captured.

to the periodic broadcast scenario reported in this paper and both results agree. The high-density includes a sufficient number of vehicles to (at least nearly) saturate the channel. Thus, the result may be compared with the saturated channel experiments reported in Fig. 3. Note that in this scenario the results differ. The experimental results with 50 and 100 senders already indicate PDRs of 54% and 44%, lower than in the simulation with 348 vehicles (about 60%).

This difference raises questions for further inquiry. We speculate that the difference in spatial arrangement of nodes accounts for some of this difference. It is also possible that the simulation model, especially its interference model, does not accurately match 802.11 behavior. A more detailed comparison between simulation and experimental results could shed light onto this.

V. Conclusions and Future Work

We have experimentally analyzed cumulative (and per-sender) throughput for a network saturation workload, and packet delivery rate (PDR) for a periodic broadcast workload emulating vehicular safety applications, in dense IEEE 802.11a networks with up to 100 senders. We highlighted the substantial effect of physical layer capture (PLC) on performance. Specifically, we conclude that:

- In a saturated network with up to 100 transmitters, a packet size of 128 bytes, and at a bitrate of 6Mbps, most receivers can correctly decode, on average, over 1500 packets per second (pps). This corresponds to a PDR of 45% and, in a preliminary comparison, appears lower than predicted by a state-of-the-art simulation in a similar (but not identical) scenario.
- Analytical models that assume the loss of all colliding packets, such as [7], significantly underestimate maximum achievable throughput. About 50% of collisions were recovered through PLC.
- For a workload that emulates vehicular safety applications, even with 100 senders within transmission range, PDR is substantially higher, about 95%, because the channel is not yet saturated.

We hope that researchers will find these results useful as a reference scenario to validate accuracy of simulation models. In future work, this study can be extended by more rigorous comparison with simulation results and an investigation of the accuracy of SINR-based interference models. It would also be interesting to experimentally investigate the effect of asynchronous collisions, due to hidden nodes, on the reliability of vehicular safety applications.

TABLE III
MEAN PDR (ACROSS 30 RUNS) FROM JAVA SIMULATIONS FOR AN ALOHA-LIKE MAC.

Number of Senders	10	30	50	70	100
PDR (%) (without PLC)	99.8	98.7	94.8	91.6	82.5
PDR (%) (with PLC)	99.9	99.1	97.9	95.6	91.3

TABLE IV
PDR FROM DSRC QUALNET SIMULATIONS [4] FOR LOW AND HIGH DENSITY SCENARIOS.

Node Density	Low	High
PDR (%)	96%	61.1%

References

[1] Q. Xu, T. Mak, J. Ko, and R. Sengupta, "Vehicle-to-vehicle safety messaging in DSRC," in *Proc. of VANET*. New York, NY, USA: ACM Press, 2004, pp. 19–28.

[2] J. Yin, T. ElBatt, G. Yeung, B. Ryu, S. Habermas, H. Krishnan, and T. Talty, "Performance evaluation of safety applications over DSRC vehicular ad hoc networks," in *Proc. of ACM VANET*. New York, NY, USA: ACM Press, 2004, pp. 1–9.

[3] X. Yang, J. Liu, F. Zhao, and N. Vaidya, "A vehicle-to-vehicle communication protocol for cooperative collision warning," in *Proc. of MOBIQUITOUS*, 2004.

[4] T. ElBatt, S. K. Goel, G. Holland, H. Krishnan, and J. Parikh, "Cooperative collision warning using dedicated short range wireless communications," in *Proc. of VANET*. New York, NY, USA: ACM Press, 2006, pp. 1–9.

[5] M. Torrent-Moreno, S. Corroy, F. Schmidt-Eisenlohr, and H. Hartenstein, "IEEE 802.11-based one-hop broadcast communications: understanding transmission success and failure under different radio propagation environments," in *Proc. of MSWiM*. New York, NY, USA: ACM Press, 2006, pp. 68–77.

[6] M. Soroushnejad and E. Geraniotis, "Probability of capture and rejection of primary multiple access interference in spread spectrum networks," in *IEEE Transactions on Communications*, June 1991, pp. 986–994.

[7] G. Bianchi, "Performance analysis of the IEEE 802.11 distributed coordination function," *IEEE Journal on Selected Areas in Communications*, vol. 18, no. 3, pp. 535–547, May 2000.

[8] S. N. Technologies, *Introduction to QualNet*, 2004.

[9] D. Son, B. Krishnamachari, and J. Heidemann, "Experimental analysis of concurrent packet transmissions in wireless sensor networks," in *Proc. of ACM SenSys*, November 2006, pp. 237–249.

[10] S. Choi, K. Park, and C. Kim, "On the performance characteristics of WLANs: Revisited," in *SIGMETRICS*, 2005.

[11] K. Ramachandran, H. Kremo, M. Gruteser, P. Spasojevic, and I. Seskar, "Scalability analysis of rate adaptation techniques in congested ieee 802.11 networks: An orbit testbed comparative study," in *Proc. of IEEE WoWMoM*, 2007.

[12] D. Raychaudhuri, I. Seskar, M. Ott, S. Ganu, K. Ramachandran, H. Kremo, R. Siracusa, H. Liu, and M. Singh, "Overview of the orbit radio grid testbed for evaluation of next-generation wireless network protocols," in *Proc. IEEE WCNC*, vol. 3, March 2005, pp. 1664–1669.

[13] MADWiFi, *Multiband Atheros Driver for WiFi*. http://madwifi.org, 2007.

[14] IEEE, *Wireless Access in Vehicular Environment (WAVE) in Standard 802.11 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 (802.11p)*, February 2006.

[15] "Matlab - the language of technical computing," http://www.mathworks.com/products/matlab/, 2007.

[16] "IEEE 802.11: Wireless LAN Medium Access Control (MAC) and Physical layer (PHY) Specifications," June 1997.

[17] A. Kochut, A. Vasan, A. U. Shankar, and A. Agrawala, "Sniffing out the correct physical layer capture model in 802.11b," in *Proc. of ICNP*. Washington, DC, USA: IEEE Computer Society, 2004, pp. 252–261.

[18] C. Ware, J. Chicaro, and T. Wysocki, "Modeling of capture behavior in IEEE 802.11 radio modems," in *Proceedings of the IEEE International Conference on Telecommunications (ICT'01)*, 2001.

[19] "Network simulator 2," http://www.isi.edu/nsnam/ns, 2005.

[20] R. Jain, D. Chiu, and W. Hawe, "A quantitative measure of fairness and discrimination for resource allocation in shared computer systems," DEC Tech Report TR-301, Tech. Rep., September 1984.