

Performance Anomalies of nonoptimally configured Wireless LANs

Pablo Serrano
Universidad Carlos III
Madrid, Spain
Email: pablo@it.uc3m.es

Albert Banchs
Universidad Carlos III
Madrid, Spain
Email: banchs@it.uc3m.es

Telemaco Melia
NEC Network Laboratories
Heidelberg, Germany
Email: melia@netlab.nec.de

Luca Vullero
Consorzio Nazionale CINI
Naples, Italy
Email: vollero@unina.it

Abstract—To this date, many works have been conducted to study the throughput and delay performance of WLANs under saturated conditions and to obtain the configuration that provides optimal performance under these conditions. From these previous works, however, it remains unclear whether this configuration is also appropriate for a WLAN operating under nonsaturation conditions. In this paper we present solid arguments which demonstrate that the optimal configuration resulting from saturation is also appropriate for a WLAN operating under nonsaturation conditions. Specifically, we show (via analysis and simulation) that a WLAN configured differently suffers from a number of performance anomalies when operating with finite sending rates. This is an important result for the configuration of WLAN parameters.

I. INTRODUCTION

The performance of Wireless LANs has been widely studied in the literature. Most analyses to this date have concentrated on the case when all stations always have packets ready for transmission [1]–[3]. This case is commonly referred to as *saturation conditions*. Assuming saturation conditions represents an important simplification and allows obtaining accurate results.

Starting from the above analyses, a number of papers have aimed at finding the configuration that optimizes performance under saturation conditions [3]–[5]. Hereafter, the expressions “optimal configuration” and “optimal configuration under saturation conditions” are used indistinctly to refer to this configuration. Indeed, with the recently approved 802.11e standard [6], a number of WLAN parameters are left open for configuration, and finding their optimal values is an important and yet unresolved research issue.

However, the convenience of using the above optimal configuration for realistic WLAN environments in which stations are typically not saturated is not clear and, in fact, this has been questioned by a number of researchers [7], [8]. With this paper we advocate the use of the optimal configuration

The work described in this paper is based on results of the IST FP6 Integrated Project DAIDALOS. It was performed during a visit of L. Vullero at the University Carlos III of Madrid funded by the IST FP6 Network of Excellence E-NEXT. DAIDALOS and E-NEXT receive research funding from the European Community’s Sixth Framework Programme. The views and conclusions contained here are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the DAIDALOS project, the Network of Excellence E-NEXT or the European Commission.

resulting from saturation conditions for any WLAN scenario (not necessarily saturated).

Our proposal to use this optimal configuration is sustained by our finding that a WLAN not configured for optimal configuration under saturation conditions suffers from a number of undesirable performance anomalies. We first illustrate these anomalies by means of simulation results and then present an analytical model that explains and corroborates the behavior observed via simulation.

By using our analytical model, we present solid arguments that demonstrate that the above anomalies do not occur in an optimally configured WLAN. This is further demonstrated by repeating the above simulations and showing that the anomalies are no longer observed when the WLAN is optimally configured.

Note that, while some of the anomalies shown here had already been observed in other papers [9], [10], and other analytical models had been proposed for WLANs under nonsaturation conditions [8], [11], the model proposed here to carry out the analysis is novel and gives an insightful understanding of the reasons and motivations for the observed anomalies that none of the previous papers provides. Our model also serves to show how these anomalies can be avoided by appropriately configuring the WLAN parameters. These results provide new insights into the behavior of WLANs that can be very useful and drive the design of configuration guidelines for 802.11e.

The rest of the paper is structured as follows. In Section II we briefly review the 802.11 DCF mechanism. In Sections III, IV and V we illustrate by means of simulation some performance anomalies (in terms of delay, stability and throughput) that this mechanism suffers from when configured according to the 802.11 standard. In these sections we also present an analytical model that explains the anomalies observed. In Section VI we propose an optimal configuration for WLANs based on our model, and show that the anomalies do not occur with our configuration. Finally, we conclude with some remarks in Section VII.

II. 802.11 DCF

This section briefly summarizes the DCF mechanism as defined in the 802.11 standard [12]. With this mechanism, a station with a new frame to transmit monitors the channel

activity. If the channel is idle for a period of time equal to *DIFS*, the station transmits. Otherwise, if the channel is sensed busy (either immediately or during the *DIFS* period), the station starts a backoff process.

Upon starting the backoff process, the station computes a random value uniformly distributed in the range $(0, CW - 1)$, and initializes its backoff time counter with this value. The *CW* value is called the contention window, and depends on the number of failed transmissions for the frame. At the first transmission attempt, *CW* is set equal to the minimum contention window parameter (CW_{min}).

As long as the channel is sensed idle the backoff time counter is decremented once every time interval σ . When a transmission is detected on the channel, the backoff time counter is “frozen”, and reactivated again after the channel is sensed idle for a certain period. This period is equal to *DIFS* if the transmission is received with a correct CRC, and *EIFS* otherwise.

As soon as the backoff time counter reaches zero, the station transmits its frame. A collision occurs when two or more stations start transmission simultaneously. An acknowledgement (Ack) frame is used to notify the transmitting station that the frame has been successfully received. The Ack is immediately transmitted at the end of the frame, after a period of time equal to *SIFS*.

If the Ack is not received within a timeout given by the *Ack_Timeout*, the station assumes that the frame was not received successfully and reschedules the transmission. The station then doubles *CW* (up to a maximum value given by the CW_{max} parameter), computes a new backoff time and reenters the backoff process. If the number of failed attempts reaches a predetermined retry limit *R*, the frame is discarded.

As it can be seen from the description given in this section, the behavior of a station depends on a number of parameters, including CW_{min} and CW_{max} . These parameters are fixed with the 802.11 standard [12], but can be configured to different values with the 802.11e standard [6]. The rest of the paper is devoted to the study of the optimal configuration of these parameters and the consequences of misconfiguring them.

III. DELAY ANOMALY

To illustrate the delay behavior of a nonoptimally configured WLAN, we consider the following scenario. We have a WLAN with 40 stations, each of them configured according to the 802.11b standard [13] ($CW_{min} = 32$ and $CW_{max} = 1024$). The transmission queue of all the stations is infinite, the packet length is equal to 1500 bytes and the traffic sources are CBR. Simulations are performed with an event-driven simulator that closely follows the details of the MAC protocol of 802.11 DCF for each independently transmitting station.

Let the saturation throughput be the throughput that each station obtains under saturation conditions [3]. We perform two tests. In the first test, all the stations send at a rate equal to 101% the saturation throughput ($r_{send} = 1.01 r_{sat}$). In the second test, all the stations send at a rate equal to 99% the




Fig. 1. Backoff delay.

saturation throughput ($r_{send} = 0.99 r_{sat}$). For both tests we measure the backoff delay, defined as the time elapsed since a packet starts its backoff process (i.e., it reaches the first position in the transmission queue) until it is successfully transmitted. Fig. 1 gives a 100 second plot of the backoff delays, averaged over 1 second intervals, obtained for the two tests.

The results of Fig. 1 are striking. Indeed, they show that a very small increase in the sending rate yields a huge increase (of about two orders of magnitude) of the backoff delay. This implies that subtle variations on the sending rate can severely degrade network performance in terms of delay. Hereafter we refer to this behavior as a *delay anomaly*.

The above behavior is rather surprising. While this behavior may be expectable for the queuing delay when the sending rate exceeds the service rate¹, a priori one would not expect such an impact on the backoff delay. In the following we present a model that explains the reasons for the observed behavior.

Following [3], we define a slot time as the time elapsed between two backoff time counter decrements and denote by τ the probability that a station transmits in a given slot time. Then, the throughput experienced by a station can be expressed as the average amount of information successfully transmitted in a slot time divided by the average duration of a slot time:

$$r = \frac{p_s l}{T_{slot}} \quad (1)$$

where p_s is the probability that a station transmits successfully in a given slot time, l is the average packet length and T_{slot} is the average duration of a slot time.

The probability that a station transmits successfully in a slot time corresponds to the case when this station transmits and no other does. The former occurs with probability τ and the latter with probability $(1 - \tau)^{N-1}$, N being the number of stations in the WLAN. Thus,

$$p_s = \tau(1 - \tau)^{N-1} \quad (2)$$

¹Hereafter, with sending rate we refer to the rate at which a station generates packets, and with service rate we refer to the rate at which the generated packets are served by the WLAN.




Fig. 2. $0.99r_{sat}$ vs. $1.01r_{sat}$.

Finally, the average slot time duration is computed as the probability that a slot time is empty, contains a successful transmission of some station, and contains a collision, multiplied by the respective slot time durations:

$$T_{slot} = (1 - \tau)^N \sigma + N\tau(1 - \tau)^{N-1} T_s + (1 - (1 - \tau)^N - N\tau(1 - \tau)^{N-1}) T_c \quad (3)$$

where T_s and T_c are the duration of a slot time that contains a successful transmission and a collision, respectively (see [3] for the formulae to compute T_s and T_c).

Fig. 2 illustrates r as a function of τ as given by Eq. (1) in the interval $[0, \tau_{sat}]$, τ_{sat} being the τ resulting from saturation conditions [3]. Note that τ can never take a value larger than τ_{sat} , since τ_{sat} corresponds to the case when all stations always have packets available for transmission.

Note that, in stable conditions, either the service rate is equal to the sending rate, or the service rate is smaller than the sending rate and the transmission queue grows indefinitely, which results in the station being saturated. Following this, it can be seen from Fig. 2 that the values of τ at which we operate in the two tests performed above are the following ones, respectively:

- $r_{send} = 0.99 r_{sat}$. In this test, the only value of τ for which the sending rate is equal to the service rate is the value marked as τ_a in the figure. Therefore, this is the τ of operation.
- $r_{send} = 1.01 r_{sat}$. In this test, the service rate under saturation (r_{sat}) is smaller than the sending rate. Therefore, the station operates under saturation: under these conditions, the queue grows indefinitely and as a result the station always has a packet available for transmission. The τ of operation is therefore $\tau_b = \tau_{sat}$.

The extreme difference observed in the backoff delay between the two tests is then explained by the fact that, although the sending rates are very close in both tests, the τ 's of operation are very different (see Fig. 2). The τ for the case $r_{send} = 1.01 r_{sat}$ is much larger, which yields a higher collision rate resulting in much longer backoff delays.

In order to validate our model, we calculated the average backoff delays resulting from the model. For the $r_{send} =$




Fig. 3. Unstable behavior.

$0.99 r_{sat}$ case, we computed τ_a by resolving numerically the nonlinear equation $r(Eq.(1)) = 0.99 r_{sat}$. Then, we computed the average backoff delay by using the model of [2] with $\tau = \tau_a$. Similarly, for the $r_{send} = 1.01 r_{sat}$ case, we also used the model of [2] with $\tau = \tau_{sat}$. The resulting average backoff delays are 2.30 and 71.35 ms, respectively, which matches the results obtained via simulation in Fig. 1.

IV. STABILITY ANOMALY

The experiments performed in the previous section with infinite queues show that small variations in the sending rate can have dramatic effects on the experienced delay performance. In this section we show that, with finite queues, dramatic changes of performance may occur continuously *while maintaining the sending rate constant*. We refer to this undesirable behavior as *stability anomaly*.

In order to illustrate the above effect, we consider the following scenario. We have a WLAN with 40 stations, each with a queue of 50 packets in our first test and of 10 packets in our second test. Traffic is generated at a rate of 110% the saturation throughput. The backoff delays resulting from these two tests, averaged over 10 second intervals, are given in Fig. 3.

The results of Fig. 3 show severe instability. At the beginning of the simulation, the backoff delay remains stable around a given value. At some point, it suddenly changes sharply to a new value. For the 50 packet queue length test, it then remains stable at the new value for the rest of the simulation run. For the 10 packet queue length test, the backoff delay keeps switching back and forth from one value to the other, although it stays most of the time around the high delay value. We note that the unstable behavior of these two tests is highly undesirable in networks and should be avoided, as it makes the resulting performance unpredictable.

The behavior observed is explained by our model as follows. The possible points of operation of the WLAN are the following ones (see Fig. 4):

- $\tau = \tau_a$. For this τ value the service rate is equal to the sending rate. This corresponds to a point of equilibrium in which all incoming packets are served.

- $\tau = \tau_b$. This also represents a point of equilibrium for the same reasons as above.
- $\tau = \tau_{sat}$. At this τ , the service rate is smaller than the sending rate which yields all the queues being filled up and the WLAN operating under saturation conditions.

However, when operating at the above τ 's, the WLAN does not remain fixed at the given τ but it suffers some perturbations around the operation point. These perturbations are caused by the probabilistic nature of the DCF mechanism. Indeed, a number of consecutive random collisions leads to an increase in the number of transmissions which yields a higher τ temporarily, and similarly a random absence of collisions during a certain period yields a smaller τ . In the following we analyze the impact of these perturbations on the equilibrium points identified above:

- $\tau = \tau_a$. At this τ value, a temporary increase of τ yields a service rate larger than the sending rate. If the service rate is larger than the sending rate, the number of packets available for transmission tends to decrease, and as a result the transmission probability τ also tends to decrease. Similarly, it can be seen that, in case of a temporary decrease of τ , the system tends to increase τ . This shows that τ_a represents a stable point of equilibrium, since upon suffering perturbations the τ tends to go back to the original value.
- $\tau = \tau_b$. At this τ , a temporary increase of τ yields a service rate smaller than the sending rate. With this service rate, not all the incoming packets are served and as a result the number of packets available for transmission increases, which yields an increase in the transmission probability τ . Under these conditions, we have that τ tends to increase continuously until reaching τ_{sat} . Similarly, it can be seen that a temporary decrease of τ leads to further decreasing τ until reaching τ_a . We conclude that τ_b represents an unstable equilibrium point and the system never operates at this τ except for very short time periods.
- $\tau = \tau_{sat}$. When operating at τ_{sat} , the service rate is smaller than the sending rate and therefore the queues become full. Perturbations around this τ do not move the system away from this point unless they empty the queues of all the stations.

From the above, the behavior observed in Fig. 3 is then explained as follows. Upon starting the simulation, τ increases until reaching τ_a and it remains stable around this value for which the service rate is high enough to serve all incoming packets.

However, as argued above, while operating at τ_a there are some perturbations around this value, and, at some point (when many collisions occur) these perturbations can bring τ to a value larger than τ_b . At this τ , the service rate is smaller than the sending rate and therefore the queues start to fill up. This further increases τ until reaching saturation, when all queues are nonempty.

Since the sending rate at saturation is larger than the service




Fig. 4. Unstability model.

rate, τ remains then stable at the saturation value, although with some perturbations as well. In the case of the 10 packet queue length, these perturbations can, with some probability, empty the queues of all the stations and bring the system back to τ_a . This is however unlikely and, in fact, the WLAN spends most of the time saturated. For the 50 packet queue length case, there are too many packets in the queues, which makes it much harder for the perturbations to empty them. In fact, this does not occur for the entire simulation run².

In order to validate our model, we calculated the analytical average backoff delays for the above test and compared them against the simulation results. From the model, we obtained an average delay of 2.68 ms for τ_a and the same value as in the previous section (71.35 ms) for τ_{sat} . These results match approximately the values we have obtained in the simulation of Fig. 3. We note, however, that the delay obtained via simulation for the 10 packet queue length test in the saturated periods is a bit smaller than the one predicted by the model. The reason is the packet drops caused by this small queue length, which is not accounted for in the model.

V. THROUGHPUT ANOMALY

We now study the following variation to the scenario of Section III. We have 40 stations with infinite queues, 39 of them transmitting at a rate of 99% the saturation throughput and the remaining one operating under saturation (i.e., always with a packet ready for transmission). Hereafter we refer to them as the nonsaturated stations and the saturated station, respectively. In Fig. 5 we plot the throughput obtained by the saturated station and one nonsaturated station, averaged over 1 second intervals.

We conducted a second test in which the 39 nonsaturated stations send at a rate of 101% the saturation throughput. In this test, all the 40 stations always have a packet available for transmission and obtain the same throughput. The throughput

²In some additional experiments, we observed that with a smaller number of stations (e.g., 10) we have the behavior of the 10 packet queue length case also for larger queues. Indeed, with fewer stations the total number of packets waiting for transmission when all queues are full is smaller, and therefore it is easier that the randomness of the mechanism leads to emptying all the queues.




Fig. 5. Throughput anomaly.

experienced by one station in this second test is also plotted in Fig. 5.

The results show that throughput performance suffers from a drastic degradation when the sending rate of the nonsaturated stations is increased by a very small amount. Indeed, while the throughput of the nonsaturated stations is kept to the same value in both tests, the throughput of the remaining station decreases sharply in the second test. We refer to this as a *throughput anomaly*.

We next present a model that explains the observed anomalous behavior. Let τ_{sat} and τ_{nonsat} be the probabilities that the saturated station and one nonsaturated station transmit in a given slot time, respectively. According to [3], the probability with which the saturated station transmits at a given slot time can be computed as

$$\tau_{sat} = \frac{2}{1 + CW_{min} + p(c)_{sat} CW_{min} \sum_{j=0}^{m-1} (2p(c)_{sat})^j} \quad (4)$$

where $CW_{max} = 2^m CW_{min}$ and $p(c)_{sat}$ is the probability that a transmission of the saturated station collides. A transmission of the saturated station collides if one or more of the $N - 1$ nonsaturated stations transmits in the same slot time. Thus,

$$p(c)_{sat} = 1 - (1 - \tau_{nonsat})^{N-1} \quad (5)$$

From the fact that the throughput experienced by the nonsaturated stations is equal to their sending rate (as otherwise they would be saturated) we have the following equation:

$$0.99 r_{sat,N} = \frac{\tau_{nonsat}(1 - \tau_{nonsat})^{N-2}(1 - \tau_{sat})l}{T_{slot}} \quad (6)$$

where $r_{sat,N}$ is the saturation throughput with N stations (computed according to [3]) and T_{slot} is computed as a function of τ_{sat} and τ_{nonsat} as follows:

$$T_{slot} = p_e \sigma + p_s T_s + (1 - p_e - p_s) T_c \quad (7)$$

with

$$p_e = (1 - \tau_{sat})(1 - \tau_{nonsat})^{N-1} \quad (8)$$

$$p_s = (N - 1)\tau_{nonsat}(1 - \tau_{sat})(1 - \tau_{nonsat})^{N-2} + \tau_{sat}(1 - \tau_{nonsat})^{N-1} \quad (9)$$

From the above, we have a system of equations on τ_{sat} and τ_{nonsat} that has only one solution in the region $\tau_{nonsat} \in (0, \tau_{sat,N})$ and $\tau_{sat} \in (\tau_{sat,N}, 1)$, $\tau_{sat,N}$ being the τ value with N saturated stations (computed from [3]). By resolving this system numerically we obtain τ_{sat} and τ_{nonsat} , from which the throughputs of the saturated and nonsaturated stations are then computed as follows:

$$r_{sat} = \frac{\tau_{sat}(1 - \tau_{nonsat})^{N-1}l}{T_{slot}} \quad (10)$$

$$r_{nonsat} = \frac{\tau_{nonsat}(1 - \tau_{sat})(1 - \tau_{nonsat})^{N-2}l}{T_{slot}} \quad (11)$$

If we perform the above computations for the first test considered in this section, we obtain a throughput of 1.79 Mbps for the saturated station and of 0.16 Mbps for each of the nonsaturated ones. The analysis for the second test is the one for saturated conditions [3] and results in a throughput of 0.16 Mbps for each station. These values match the simulation results obtained in Fig. 5, which validates our model.

VI. OPTIMAL CONFIGURATION

We now address the issue of finding the optimal configuration that avoids all the performance anomalies observed in the previous sections.

It can be seen from the model used throughout the paper that the cause for the anomalies is the ‘‘bell’’ shape of r as a function of τ (see Figs. 2 and 4). This shape yields that for a given sending rate (or two very close sending rates) there are several possible τ 's of operation very distant from each other. Indeed, all the anomalies observed are caused by the performance degradation resulting from moving from one of these operation points to the other.

The main idea behind the optimal configuration that we propose here lies in avoiding this ‘‘bell’’ shape by setting the CW_{min} and CW_{max} parameters such that τ_{sat} corresponds to the point of the curve $r(\tau)$ where r takes its maximum value (see [3] for the formulae to compute τ_{sat} as a function of CW_{min} and CW_{max}). This is illustrated in Fig. 6. We hereafter denote this value by τ_{max} . Note that with this limitation on the maximum allowed τ there is only one stable operation point given a certain sending rate, eliminating thus the cause for the anomalies observed. Specifically, note that with this configuration the transition from nonsaturation to saturation is smooth, as the τ of operation converges towards τ_{sat} upon approaching saturation (in the figure it can be observed that for $r = 0.99 r_{sat}$, τ_{nonsat} is very close to τ_{sat} , in contrast to Fig. 2).

Note that the above τ_{max} value is the one that maximizes the saturation throughput. Therefore, the optimal configuration




Fig. 6. Optimal Configuration.

that we are proposing here coincides with the one proposed in [3]–[5]. However, the fundamental difference between those papers and this one is that while their reason to propose this configuration is to *maximize performance when the WLAN operates under saturation*, our reason here is to *avoid the performance anomalies when operating under nonsaturation conditions*. Note that our scenarios, in contrast to those other papers, apply to the realistic case in which the WLAN stations send at a finite rate, possibly below the saturation throughput.

Based on the above, in this paper we reuse the work of [5] to find the optimal configuration of the CW_{min} and CW_{max} parameters as a function of the number of stations in the WLAN. The optimal configuration proposed in that paper for the case of $N = 40$ which we have taken as reference here is the following^{3,4}: $CW_{min} = 386$ and $CW_{max} = 2^5 CW_{min}$.

In order to confirm that a WLAN configured with the above parameters does not suffer from the performance anomalies that we have observed with the default 802.11b configuration, we have repeated the simulations performed in the previous sections but with the optimal configuration proposed above. The simulation results obtained (not reported because of space reasons) confirm the behavior predicted by our analytical model: they show that none of the throughput, stability and delay anomalies occur with the optimal configuration. We conclude that a WLAN configured in a way that its saturation throughput is maximized does not suffer from performance anomalies. Based on these results, we advocate this configuration for WLANs.

VII. CONCLUSIONS

In this paper we have studied a number of performance anomalies (some of which had already been observed in

³The optimal configuration given in [5] is left open as a function of m (where $CW_{max} = 2^m CW_{min}$). In this paper we have taken $m = 5$ in order to keep the good features of the exponential backoff increase algorithm.

⁴Note that the CW configuration given here is the result of having many stations (40) in the WLAN; with less stations, the optimal CW values would be smaller. One of the drawbacks of using large CW values is that delay performance is slightly worse under light load conditions. However, we argue that this is a small price to pay in order to guarantee good performance (including avoidance of the performance anomalies) under heavier loads.

previous papers) that misconfigured WLANs (and in particular configured according to the 802.11b standard) suffer from under certain conditions. The anomalies have been observed via simulation and analyzed by means of a model that we have developed.

Our model provides an insightful understanding of the behavior of the WLAN including the reasons and causes for the anomalies observed. The throughput and delay results obtained from the model match reasonably well those obtained via simulation.

By using our model, we have shown that a WLAN configured in a way that the throughput obtained under saturation conditions is maximized does not suffer from the studied performance anomalies. This is the optimal configuration that we propose in this paper.

While the configuration that we propose here coincides with the one proposed by other papers, the rationale behind our proposal is fundamentally different. The other papers propose the configuration because it is the one that optimizes performance under saturation conditions. We propose it because with this configuration we avoid the various performance anomalies observed in the paper when the stations are sending at a finite rate. Our scenario corresponds to a more realistic WLAN environment.

REFERENCES

- [1] H. Wu, Y. Peng, K. Long, S. Cheng, and J. Ma, "Performance of Reliable Transport Protocol over IEEE 802.11 Wireless LAN: Analysis and Enhancement," in *Proceedings of IEEE INFOCOM 2002*, June 2002.
- [2] A. Banchs and L. Vullero, "A Delay Model for IEEE 802.11e EDCA," *IEEE Communications Letters*, vol. 9, no. 6, June 2005.
- [3] 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, March 2000.
- [4] D. Qiao and K.G. Shin, "Achieving Efficient Channel Utilization and Weighted Fairness for Data Communications in IEEE 802.11 WLAN under the DCF," in *Proceedings of the Tenth International Workshop on Quality of Service*, May 2002.
- [5] A. Banchs and L. Vullero, "Throughput Analysis and Optimal Configuration of 802.11e EDCA," to appear in *Computer Networks*.
- [6] IEEE 802.11e, *Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) specifications: Medium Access Control (MAC) Enhancements for Quality of Service (QoS)*, Supplement to IEEE 802.11 Standard, November 2005.
- [7] D. Gao, J. Cai, and K. N. Ngan, "Admission Control in IEEE 802.11e Wireless LANs," *IEEE Network*, vol. 19, no. 4, July 2005.
- [8] B. Li and R. Battiti, "Analysis of the IEEE 802.11 DCF with Service Differentiation Support in Non-Saturation Conditions," in *Proceedings of QoIS'04*, September 2004.
- [9] R. Vijayakumar, T. Javidi, and M. Liu, "From saturation to nonsaturation: A study on 802.11 networks," Dept. of Electrical Eng., University of Michigan, Tech. Rep. CSPL-363, 2005.
- [10] B.-J. Kwak, N.-O. Song, and L. E. Miller, "Analysis of the Stability and Performance of Exponential Backoff," in *Proceedings of WCNC'03*, March 2003.
- [11] O. Tickoo and B. Sikdar, "Queueing Analysis and Delay Mitigation in IEEE 802.11 Random Access MAC based Wireless Networks," in *Proceedings of IEEE INFOCOM'04*, Hong Kong, China, March 2004.
- [12] IEEE 802.11, *Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) specifications*, Standard, IEEE, August 1999.
- [13] IEEE 802.11b, *Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) specifications: High-speed Physical Layer Extension in the 2.4 GHz Band*, Supplement to IEEE 802.11 Standard, September 1999.