

A Measurement Study on Link Capacity of a High Stress IEEE 802.11b/g Network

Michael Juang¹, Kuang-Ching Wang^{1*}, James Martin²

¹ Department of Electrical & Computer Engineering, Clemson University, Clemson, SC 29634 USA

² School of Computing, Clemson University, Clemson, SC 29634 USA

*Corresponding author

Abstract—The paper presents results from a link capacity measurement study conducted over an IEEE 802.11b/g network with highly stressed radio propagation conditions in a football stadium. Existing capacity studies for IEEE 802.11 networks have considered either stationary or statistical multipath fading channel conditions. In an environment with a high density of people and movement such as in a football stadium or a subway station, the radio characteristics vary much more vibrantly, and their impact on network capacity is not well understood. As a first step to better understand such network characteristics, the study examines the throughput of a mobile terminal over the stadium network in the presence of varying radio conditions due to movement of people and concurrent transmissions by other IEEE 802.11 devices. The study uses a novel approach to quantify network utilization by the experiment terminal as well as other contending terminals, using a recorded packet trace. From the measured utilization, throughput, and signal strengths, key observations made include the fast variations in signal strength due to crowd movement, throughput's dependency on signal strength and chosen link rate, and contention with other terminals. Based on the findings, the paper discusses useful insights for performance assessment, protocol enhancements, and improvements in measurement methodology for such networks.

Index Terms—WLAN Testbed, Performance Measurement, Radio Propagation.

I. INTRODUCTION

In recent years, numerous large scale IEEE 802.11-based wireless local area networks (WLANs) have been deployed to cover public spaces and entire cities to enable a wide range of mobile applications in domains of public safety, public service, and personal communications. As society becomes increasingly dependent on such infrastructures, it is imperative to assure that the networks deliver stable performance in all conditions. This study is interested in the performance of IEEE 802.11 networks that operate in public areas with a substantial number of people and potentially a large number of networking devices. Despite the already extensive literature on IEEE 802.11 network capacity analysis, a limited body of work has considered such a *high stress* condition.

In this paper, a high stress network is considered to have two defining characteristics: (i) rapidly changing radio channel conditions and (ii) a large number of contending terminals accessing the same network medium. In a public gathering place like a stadium or a subway station, a large crowd of people with frequent large movements (standing up and waving in a stadium or moving around a station) can potentially cause the radio propagation condition between a communicating terminal and the network infrastructure to change frequently. Also, a potentially large number of

network users may cause substantial contention within the network. To our knowledge, radio channel properties in such environments have not been studied systematically before.

As a first step to develop an accurate model for such environments, a measurement study was conducted over an operational IEEE 802.11 b/g network within Clemson University's football stadium (Clemson Memorial Stadium) during home games. With existing (uncontrolled) users accessing the network, an experimental laptop computer, referred to as a mobile terminal (MT), was used to generate saturating traffic for link capacity assessment. Configured in promiscuous mode, the MT recorded a trace of all transmitted, received, and overheard packet transmissions to measure its achieved throughput and link capacity. Since the measurement was completely based on the MT's acquired information, the procedure can generally be used by any terminal to assess the link capacity from its own perspective. The received packets' signal strengths were also studied. The network is part of the Clemson *iTiger* stadium network that is being designed to support simultaneous multimedia download for tens to hundreds of fans, for which a baseline performance study is required.

The rest of the paper is organized as follows. In Section II, previous WLAN performance analysis studies are reviewed. Section III gives an overview of the IEEE 802.11 medium access control (MAC) protocol as a basis for Section IV, where analytical expressions for packet transmission durations under different conditions are derived. Section V describes the measurement methodology, and Section VI presents the experimental results and discussions, while Section VII summarizes possible improvements for the measurement methodology. Finally, the paper concludes in Section VIII.

II. RELATED WORK

Many studies have been conducted to assess the performance and capacity of IEEE 802.11 networks. In the works of [3] and [5], the capacity of an IEEE 802.11 network under ideal channel conditions was derived analytically according to the distributed MAC protocol and terminals with asymptotic (constantly backlogged) transmissions; the capacity thus derived, as a result, was referred to as the saturation throughput. In [2] and [9], the effects of different link rates and protocol overheads in IEEE 802.11b/g were incorporated in the analytical models with single-client and multi-client scenarios respectively; idealized, lossless channel conditions were assumed. In [7] and [8], the effects of varying physical channel conditions were studied by applying different statistical multipath fading models to their testbed and

simulation experiments, respectively, considering no other interference or medium contention.

Already, a large number of measurement studies have been conducted on operational IEEE 802.11 networks serving real users (e.g., [1, 4]). These studies have commonly observed the achieved throughput of an experimental terminal to vary over time due to reasons such as terminal mobility, multipath fading, and random contention with other users. These studies, however, have also commonly ignored to collect information about the channel conditions and transmission activities of other terminals for analyzing the causes of the experimental terminal's throughput variations. These factors are crucial in understanding how actual networks behave differently than idealized network models in other works in the literature.

The Purdue eStadium project developed a stadium-wide IEEE 802.11b network as a live testbed for high stress wireless network research [10]. In [10], the network implementation and potential research directions were described. This measurement study was conducted in a similar stadium network, with extensive transmission history and link layer events recorded for systematic analysis of the link capacity, terminal achieved throughput, and the varying impacts from fluctuating radio conditions.

III. OVERVIEW OF THE IEEE 802.11 MAC PROTOCOL

The studied stadium network employs IEEE 802.11 b/g in infrastructure mode with the distributed coordination function (DCF) MAC protocol. In this section we briefly summarize the relevant channel access procedures as the basis for network utilization calculations in the paper. To avoid collisions among transmissions within the same IEEE 802.11 channel, the DCF protocol defines inter-frame spacing rules and an exponential random backoff procedure for each station and access point (AP) to observe.

Before transmitting a packet, a station must first sense the channel to determine if another station is already transmitting. If the channel is idle, the station waits for a distributed inter-frame spacing (DIFS) duration; once the channel remains idle for DIFS, the station proceeds with transmission of its data packet. If the channel is busy, the station waits for the channel to be idle for DIFS plus a random backoff duration. The backoff duration is randomly chosen between $[0, CW]$ with uniform probability where CW is the current contention window size in number of slots. A backoff timer counts down from the chosen duration when the medium is sensed to be idle. Once the timer expires and the medium is idle, the station can proceed to transmit. If a station's transmission fails, CW is increased according to the binary exponential algorithm as defined in [6] until CW reaches CW_{\max} .

The experimental MT has been configured to never send Request-to-Send (RTS) packets before its transmissions; nevertheless, other devices in the network did send out RTS and Clear-to-Send (CTS) packets prior to their data transmissions. Thus, in a data transmission sequence, first the data frame is sent after the DIFS wait. Then, the receiving station waits for a short inter-frame spacing (SIFS) duration before sending an acknowledgement (ACK) to the sender, completing the exchange. The station requires the correct

reception of an ACK to consider the transmission sequence a success. Broadcast and certain control frames do not require ACKs.

Failure can result from either (1) the receiving station not correctly receiving the data packet from the sending station and hence not sending out an ACK, or (2) the receiving station getting the data but the sending station not receiving the ACK correctly. In the first case, the station waits for an ACK timeout duration, a DIFS, and then a random backoff duration following its own transmission before it retransmits. In the second case, the station waits for an extended inter-frame spacing (EIFS) duration following the incorrectly received packet before resuming random backoff. A station retransmits until the packet is received successfully or the number of retransmissions reaches the specified maximum retry limit. Upon a successful transmission or when retry limit is reached, the station resets CW to CW_{\min} .

For an IEEE 802.11b/g network, the standard defines protection mechanisms such as lengthened inter-frame spacing and CTS-to-self for avoiding collisions among b and g stations. However, no such protection mechanisms were observed in our measurements and hence they are not considered in our analysis. IEEE 802.11b allows for data link rates (DLRs) of 1, 2, 5.5, and 11 Mbps. Most IEEE 802.11g devices, including the experimental MT and APs, support all IEEE 802.11b rates as well as 6, 9, 12, 18, 24, 36, 48, and 54 Mbps; of these, all IEEE 802.11b rates as well as 6, 12, and 24 Mbps are defined as mandatory rates. IEEE 802.11g requires that ACKs be sent only at a mandatory DLR less than or equal to that of the preceding data packet. For example, a 54 Mbps data transmission requires a 24 Mbps ACK response.

IV. TRANSMISSION DURATIONS OF AN IEEE 802.11B/G NETWORK

A. Duration of a Transmission Sequence

To assess the network capacity in terms of efficiency of its utilization, we analyze the durations in which a network is engaged in each event sequence related to a transmission.

According to the DCF protocol, to transmit each packet, each station keeps the network "busy" for the duration of a sequence of events. Depending on the station's adopted options (RTS/CTS, CTS-to-self, etc.), and the eventual success/failure of the transmission, the sequence of events varies. In our study, the experimental MT did not adopt RTS/CTS, but other devices and the AP did send out RTS, CTS, and occasionally CTS-to-self packets.

Each transmitted packet incurs physical (PHY) and MAC layer header overhead. While the PHY overhead is a fixed duration, the MAC header is transmitted with the remaining part of a packet at its present DLR. The standard transmission and framing parameters used are as given in Table I.

Depending on the success or failure of a transmission, three event sequences are possible. Let L_{Data} be the number of data and MAC header bits encapsulated inside PHY overheads; these bits are transmitted at the same DLR R_{Data} over a duration of $T_{Tx}(R_{Data}, L_{Data})$. Similarly, L_{ACK} and R_{ACK} are the length of an ACK (a fixed size) and DLR used to send it. The length of a data transmission is given, where LLC , IP , UDP , and MAC represent the corresponding protocol header sizes, as

$$L_{Data} = \text{Payload} + LLC + IP + UDP + MAC \quad (1)$$

For IEEE 802.11g rates, $R \in \{6, 9, 12, 18, 24, 36, 48, 54\}$:

$$T_{TX}(R, L) = T_{PHY} + T_{SYM} \cdot \lceil (22 + L) / R \cdot 4 \rceil \quad (2)$$

For 802.11b rates, $R \in \{1, 2, 5.5, 11\}$:

$$T_{TX}(R, L) = T_{PHY} + \lceil L / R \rceil \quad (3)$$

Hence,

i. duration for a successful transmission sequence is:

$$T_{Seq} = T_{TX}(R_{Data}, L_{Data}) + SIFS + T_{TX}(R_{ACK}, L_{ACK}) + DIFS \quad (4)$$

ii. duration for a failure due to failed data packet reception is:

$$T_{Seq} = T_{TX}(R_{Data}, L_{Data}) + T_{ACKTimeout} \quad (5)$$

$$T_{ACKTimeout} = SIFS + SlotTime + T_{PHY} \quad (6)$$

iii. duration for a failure due to failed ACK reception is:

$$T_{Seq} = T_{TX}(R_{Data}, L_{Data}) + SIFS + T_{TX}(R_{ACK}, L_{ACK}) + T_{EIFS} \quad (7)$$

$$T_{EIFS} = SIFS + DIFS + T_{TX}(R_{ACK}, L_{ACK}) \quad (8)$$

For each transmission, the random backoff duration is randomly chosen as

$$T_{BO}(CW) = SlotTime \cdot rand([0..CW]) \quad (9)$$

where $rand(\cdot)$ represents a uniform random function for the specified range.

Table I. Transmission and Frame Parameter Values [6]

| | 802.11g | 802.11b |
|--------------------------|------------------------|----------------------------|
| PHY overhead, T_{PHY} | 26 μ s | 192 μ s*, 96 μ s** |
| Symbol Length, T_{SYM} | 4 μ s | - |
| MAC header | 288 bits*** | 288 bits*** |
| UDP/IP/LLC | 288 bits | 288 bits |
| Payload | 3760, 7760, 11760 bits | 3760, 7760, 11760 bits |
| RTS | 160 bits | 160 bits |
| CTS | 112 bits | 112 bits |
| ACK | 112 bits | 112 bits |
| SlotTime | 9 μ s | 20 μ s |
| DIFS | 28 μ s | 50 μ s |
| SIFS | 10 μ s | 10 μ s |
| CW_{min} | 15 | 31 |
| CW_{max} | 1023 | 1023 |
| Retransmission limit | 15**** | 15**** |

*long preamble **short preamble ***includes WEP
****value used by MT, not a fixed standard parameter

B. Achieved Throughput and Throughput Upper Bound

As the experimental MT saturates the network in the presence of traffic from other devices, from the perspective of the experimental MT, all experimental time is spent in one of the following mutually exclusive states:

1. Completion of a successful transmission sequence. The time spent in this state is denoted as the experimental transmission time (ETT) found with Eq. 4.
2. Completion of a failed transmission sequence. The time spent in this state is denoted as the failed transmission time (FTT) found with Eq. 5~8.
3. Sensing an idle channel and counting down the random backoff timer. Time spent in this state is denoted as the backoff duration (BO) found with Eq. 9.
4. Sensing a busy channel during others' transmission. Time spent in this state is denoted as others' transmission time (OTT) found with Eq. 4. Note that certain transmissions can be sensed by the MT but not decoded by the MT if their signal strengths are weak;

such transmissions are not accounted for in our analysis and, as shown in Section VI, can be significant.

Had the MT been able to accurately determine its time spent in each of these four states during any time duration T , $T = ETT + FTT + BO + OTT$ must hold true. Other stations' backoff durations are not counted separately as they coincide with the experimental MT's BO. Since the random backoff period is not determined in the open source network adapter driver but rather in the accompanying binary microcode, we are unable to directly examine BO time. Instead, the expected BO is calculated based on history of MT retransmissions by applying the standard's binary exponential backoff algorithm to determine CW. From the following results, as well as findings in [2], it was found that the actual BO distribution obtained from vendor devices could skew towards lesser values.

Since ETT measures the duration of successful transmissions, the achieved throughput (achieved TPT) can be found as the number of bits sent per unit time during ETT. Since the ultimate purpose of the study is to exploit methods to improve the throughput performance in such a high stress network, it is meaningful to assess the upper bound throughput an MT can ever possibly achieve with any ideal protocol enhancements. We define the upper bound throughput (UB TPT) as the maximum data throughput achievable during the combined ETT, BO, and FTT durations, assuming no collisions and a random backoff of $CW_{min}/2$ SlotTime prior to each transmission at the chosen DLR.

The recorded packet trace contains all header information of each received packet. Specifically, the packet timestamp, DLR, length, and network allocation vector (NAV) were used to compute the MT's busy periods according to its physical and virtual carrier sensing procedures to compute OTT.

V. EXPERIMENTAL SETUP

A. Measurement Testbed

Figure 1 shows the studied IEEE 802.11 infrastructure mode network covering the West Zone club area of the Clemson University Memorial Stadium. Three IEEE 802.11 b/g APs cover an indoor lobby and an outdoor bleacher seating approximately one thousand fans. Measurements were taken during the October 20, 2007 football game. The three APs shown and several other APs in the south side of the stadium chose dynamically from one of three non-overlapping channels. APs selecting the same channel can potentially interfere with each other.

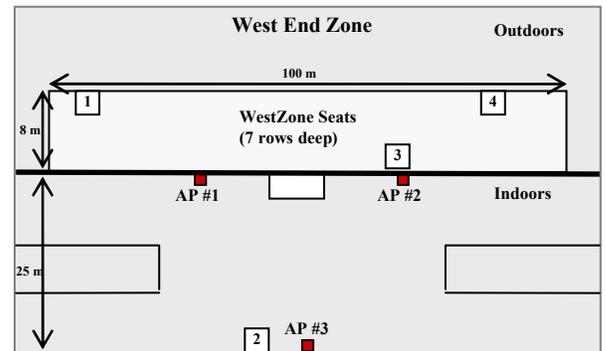


Fig. 1. Network testbed layout. Numbered squares are testing positions.

B. Devices and Software Used

The experimental MT was an IBM Thinkpad T60 laptop running Ubuntu 7.04 Linux. The MT was equipped with a built-in Intel Pro/Wireless 3945 ABG Wi-Fi adapter (driver: ipw3945 v1.2.1), as well as an external Cisco Aironet ABG PCMCIA Wi-Fi adapter (driver: Madwifi v0.9.3.1). The built-in adapter was used by Iperf v.2.0.2 to connect with an AP and saturate the network with data packet transmissions. The MT's ipw3945 driver was modified to provide additional per-packet transmission history and statistics, and to completely lock the DLR (disabling auto rate fallback). To passively monitor other transmissions, the external adapter monitored the channel with Wireshark (formerly Ethereal) v.0.99.5 in promiscuous mode. The MT transmit power was 16 dBm. The APs were Cisco 1200 series transmitting at 20 dBm. All other default settings were used.

C. Experiment Procedure

For each measurement run, the MT performs the following:

1. Set the MT internal adapter to transmit only at one fixed DLR (Table II)
2. Establish a connection to an AP (selected by driver as recorded in Table II)
3. Launch Wireshark to monitor channel
4. Launch Iperf to begin UDP upload with
 - a. Duration: 50 seconds, including 10 sec warm-up
 - b. Injected data rate: 54 Mbps
 - c. Packet size: fixed (Table II)

The purpose of having completely locked DLR during each run is to acquire the link characteristics associated with each DLR, which will form the basis of a multi-rate link model. Table II summarizes the configurations for all measured runs.

D. Limitations and Assumptions

Several limitations were imposed on the experimental methodology. Several assumptions also affected the analysis.

Per-packet received signal strength indication (RSSI) provided by the wireless interface was adopted as an approximate indicator of link quality. Since the AP did not support per packet logging and status inspection, RSSI of received periodic beacons transmitted by the AP was instead monitored. Since received signal includes signal, noise, and interference power, a high RSSI does not necessarily assure correct packet reception; on the other hand, a low RSSI can still be an acceptable indicator of poor channel conditions.

The RSSI values reported in Table III were obtained with the external network adapter rather than the internal one. A receiver sensitivity analysis between the two cards found that for the same transmissions, the external adapter reports a lower RSSI with an average difference of 8 dB.

Since useful measurements could only be taken during the game, only the listed runs in Table II could be carried out (a total of 96 configurations). Trials could not be fully repeated given the time limitations, and channel conditions (crowd movements) varied significantly between measurement runs; thus, the analysis in the next section generally focuses on notable trends within each experimental trial.

E. Empty Stadium Measurements – A Baseline

To establish a baseline, similar measurements were taken

Table II. Test Scenarios

| Position | AP | DLRs tested | MT Packet Payload | Time of game |
|----------|----|-------------------------------|------------------------|--------------|
| 1 | 1 | All 802.11b and 802.11g, auto | 11760, 7760, 3760 bits | Q1-2 |
| 2 | 1 | 24 and 54 Mbps, auto | 11760, 7760, 3760 bits | Halftime |
| 3 | 2 | All 802.11b and 802.11g, auto | 11760, 7760, 3760 bits | Q3-4 |
| 4 | 2 | 24 and 54 Mbps, auto | 11760, 7760, 3760 bits | Q4 |

Table III. Summary of Empty Stadium Results

| Position | AP | Throughput (Mbps) | | RSSI (dBm) | |
|----------|----|-------------------|-----------|------------|-----------|
| | | Mean | Std. Dev. | Mean | Std. Dev. |
| 1 | 1 | 23.21 | 2.14 | -63.60 | 1.90 |
| 2 | 1 | 24.32 | 4.20 | -64.62 | 2.57 |
| 3 | 2 | 26.37 | 0.32 | -55.30 | 2.03 |
| 4 | 2 | 22.53 | 2.75 | -63.45 | 2.54 |

when the stadium was empty during the off season. The achieved throughput, due to lack of other users, was very stable. Table III summarizes throughput and RSSI measured in MT UDP upload experiments with DLR 54 Mbps and payload size 11760 bits at the same four locations.

VI. MEASUREMENT RESULTS AND DISCUSSIONS

In this section, selected metrics of the wireless environment are presented to discuss their relation to the link capacity. All figures shown are for data measured at position 1, where physical channel effects were seen to be greatest.

A. Received Signal Strength Indicator (RSSI)

Fig. 2 shows the RSSI measured at position 1 over the course of the experiments. In two episodes lasting 18 and 25 seconds, respectively, the RSSI dropped well below the usual mean of -57.4 dBm to a lower mean of -76.9 dBm; at both occasions it was noted that the crowd stood to cheer around the MT, causing substantial obstruction between the MT and the AP. The same pattern of RSSI drop was not observed in other positions presumably because fewer people were in the area that could have blocked the signal.

Fig. 3 shows the achieved throughput during a period of fluctuating RSSIs. In Fig. 3-5 and 7-8, time is given relative to the start of the experimental trial in question. As defined in Section VI.B, the UB TPT reflects the maximum achievable throughput by an MT if all its transmissions were successful. In the figure, UB TPT fluctuates slightly over time, but the actual achieved throughput varies substantially. The steep drop in throughput indicates a substantial increase in BO time, which is directly relevant to the corresponding RSSI drops. The finding suggests the extent of human obstruction between an MT and an AP critically affects their link capacity and extent of performance variations.

B. Data Link Rate

Lower DLRs utilize more robust modulation schemes and are more tolerant to severe channel conditions. Fig. 4 and 5 show the effect of two severe RSSI drops (circled in Fig. 2) upon achievable TPT and UB TPT. While 9 Mbps DLR is already more robust than 54 Mbps DLR, its achieved TPT still dropped from 6.56 Mbps to 1.11 Mbps as RSSI dropped from

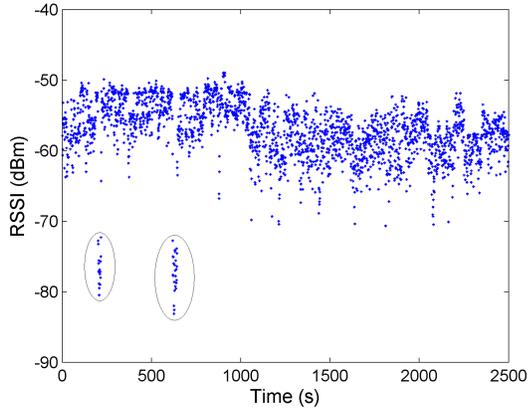


Fig. 2. RSSI over all experiments at position 1.

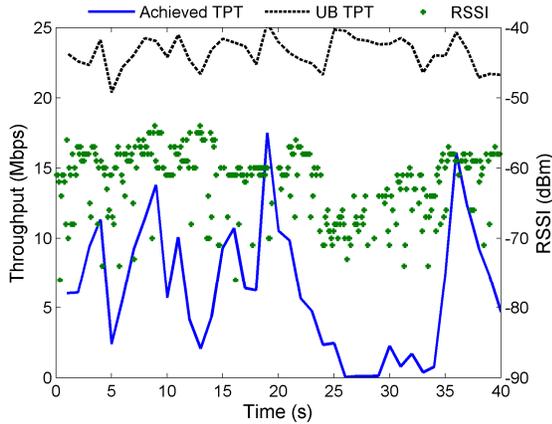


Fig. 3. RSSI with achieved TPT and UB TPT; DLR=54 Mbps, L=11760 bits.

-52 to -83 dBm. With 2 Mbps DLR, the achieved TPT only exhibited a modest decline from 1.73 to 1.41 Mbps. Nevertheless, the overall mean achieved TPT for 9 Mbps DLR was 5.13 Mbps with a packet loss rate of 26.3%, compared to 1.49 Mbps mean achieved TPT and 7.5% loss rate for 2 Mbps DLR under very similar OTT and RSSI conditions. The finding suggests the choice of DLRs must be evaluated with their respective loss rates for any given channel conditions.

C. Others' Transmission Time (OTT) and MT Backoff (BO)

Fig. 6 shows the OTT recorded at position 1 over the course of all experiments, reflecting all transmissions that occurred from other devices in surrounding networks that utilized the same channel as the MT. Fig. 7 shows the per-second breakdown of ETT, FTT, BO, and OTT for one particular experimental trial at 24 Mbps DLR. Fig. 8 shows the same for another trial at 54 Mbps DLR (same trial as Fig. 3).

During the trial, 41 stations other than the MT were seen to have transmitted or received packets, though one particular station was responsible for 85% of the non-experimental traffic. It was noticed that the majority of non-experimental traffic was sent at lower DLRs (mostly 1 Mbps) and hence consumed more time per transmission than the MT's transmissions (e.g., during 20~40 sec).

Fig. 7 and 8 show the breakdown of time into four distinct states of the network. Ideally, the time breakdown should add up to one. In reality, various factors would cause this estimated sum to deviate from one. During 50% of all trials

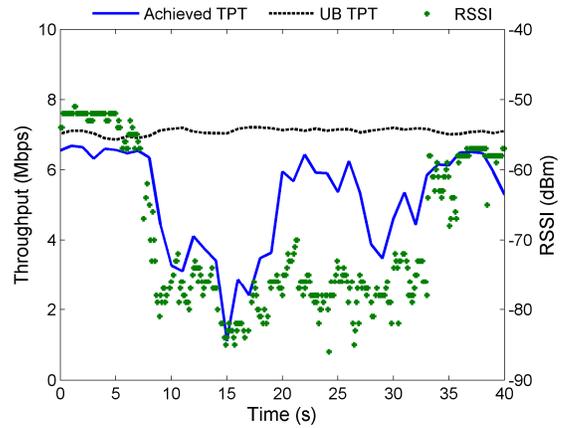


Fig. 4. RSSI with achieved TPT and UB TPT; DLR=9 Mbps, L=11760 bits.

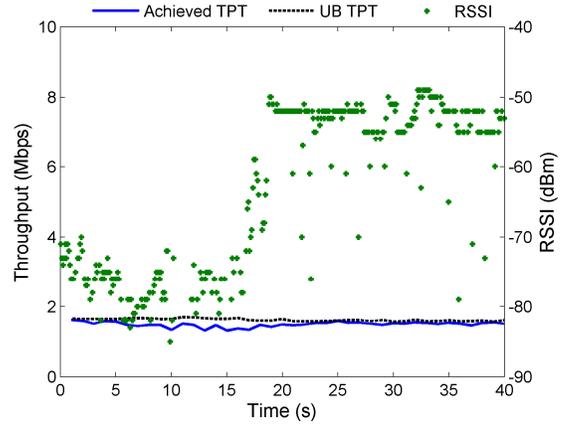


Fig. 5. RSSI with achieved TPT and UB TPT; DLR=2 Mbps, L=7760 bits.

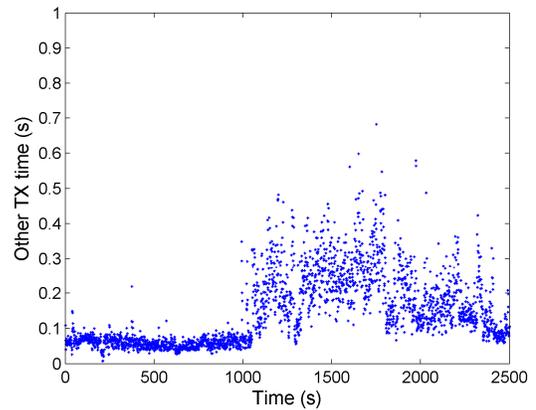


Fig. 6. OTT duration per second over all experiments at position 1.

the total time utilization as seen by the MT was calculated to be within 1 ± 0.02 seconds. However, in the other 50% errors arose from unaccounted elements in the time analysis.

Fig. 7 shows a test run in which the sum is far less than 1, while Fig. 8 shows a contrary case where the sum exceeds 1. The lesser sum can be attributed to interfering transmissions that caused the MT to defer its transmission but could not be decoded and recorded by Wireshark due to insufficient energy or collision with others. Fig. 8 shows relatively steady competing traffic (OTT), while significant MT BO time was seen. Since all time components other than BO were precisely found based on packet trace and only BO was *estimated* based

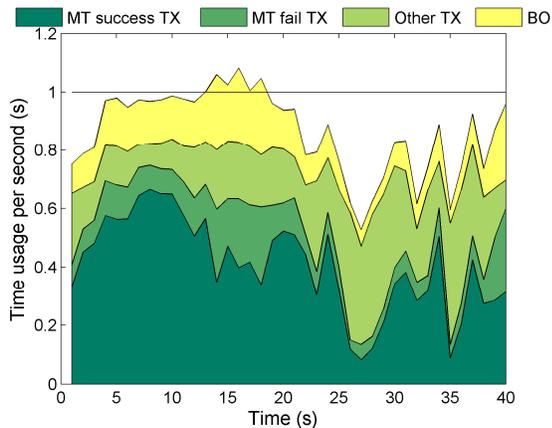


Fig. 7. Time usage breakdown; DLR = 24 Mbps, L = 7760 bits.

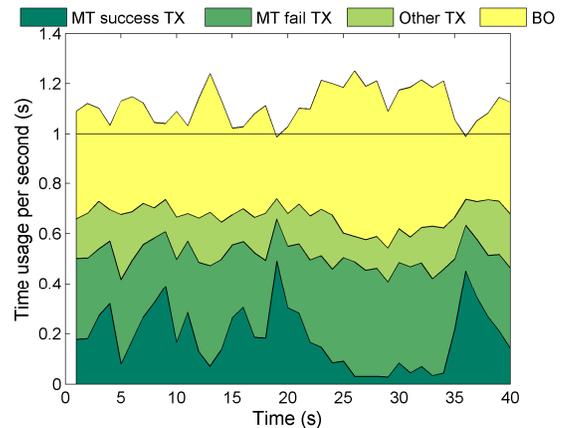


Fig. 8. Time usage breakdown; DLR = 54 Mbps, L = 11760 bits.

on transmission success history and an average contention window duration, BO was the only source of error in time estimation (other than unrecorded interfering transmissions as discussed for Fig. 7). Note that Fig. 8 was obtained with MT at 54 Mbps; at higher DLRs, the MT made more transmission attempts in any given time and potentially caused more collisions or failures due to bad channel (e.g., 20~35 sec). The more BO occurrence, the more potential error in the estimation can be expected. The exceeding sum indicates the actual BO duration being less than the expected duration.

VII. POSSIBLE SOURCES OF ERROR AND METHODOLOGY IMPROVEMENTS

The long term goal of our measurement study is to develop a methodology that adequately captures the essential features of a high stress network, such that an accurate model can be derived to facilitate capacity assessment and protocol development. The experiments and new metrics presented in this paper were a first step towards this goal; more importantly, they allow us to identify potential factors that can be overlooked in measurement studies, and useful information that can be derived in future experiment designs.

Summarizing discussions in previous sections, it was concluded that, while the packet trace recorded at the MT accurately captured its achievable throughput and failures, lack of MT BO information, lack of an AP packet trace, and lack of accurate information of transmissions from other devices have limited our model analysis. Obtaining greater AP control and measuring AP transmissions would resolve questions of channel symmetry with regards to interference and contention. The two adapters used for the experiments had different reception sensitivity – potential discrepancy may occur between their recorded events. Gathering all statistics with one interface is desirable but challenging. One important limitation regards the lack of information of packets that were not received correctly and hence not recorded in the trace, leading to unaccounted time utilization. Such information can potentially be made available through driver modifications and inference based on AP and additional monitoring stations' transmission history.

VIII. SUMMARY AND CONCLUSION

The paper presents a measurement methodology for an

IEEE 802.11b/g network inside a football stadium in game day conditions. The method utilizes saturating transmissions from a MT that also recorded the transmission history of its own and others and analyzed the achieved and upper bound throughput, network utilization, and RSSI. Up to 25 dB RSSI drop and 20 Mbps throughput drop was caused by crowd movements; 2 Mbps DLR was able to tolerate the crowd impairments, yet higher DLRs can still lead to higher long term throughput. The measurement method can be further enhanced with better control of AP, MT's driver, and additional network monitoring stations.

ACKNOWLEDGEMENTS

We acknowledge the support of the Clemson iTiger project team and the Clemson Computing & Information Technology.

REFERENCES

- [1] D. Aguayo, et al., "Link-level Measurements from an 802.11b Mesh Network," in *Proc. ACM SIGCOMM'04*, pp. 121-132.
- [2] M. Arranz, R. Agüero, L. Muñoz, and P. Mähönen, "Behavior of UDP-Based Applications over IEEE 802.11 Wireless Networks", In *Proc. IEEE PIMRC 2001*, Vol. 2, pp.F-72 – F-77.
- [3] G. Bianchi, "Performance Analysis of the IEEE 802.11 Distributed Coordination Function." *IEEE JSAC*, 18(3):535-547, Mar. 2000.
- [4] M. Boulmalf, H. El-Sayed, A. Soufyane, "Measured Throughput and SNR of IEEE 802.11g in a Small Enterprise Environment", In *Proc. IEEE VTC 2005*, Vol. 2, pp. 1333-1337.
- [5] F. Cali, M. Conti, and E. Gregori, "IEEE 802.11 Wireless LAN: Capacity Analysis and Protocol Enhancement." In *Proc. INFOCOM'98*, pp. 142-149.
- [6] IEEE, *IEEE STD 802.11-2007*, 8 March 2007.
- [7] W. Liu, L. Wang, and Y. Lin, "Physical Layer Effects on the MAC Goodput Performance of the Rate Adaptive IEEE 802.11a/g WLAN", In *WCNC 2004*, Vol. 3, pp. 1873-1878.
- [8] L. Metreud and K. Pahlavan, "RF Isolated Real-Time Multipath Testbed for Performance Analysis of WLANs", In *Proc. 40th Ann. Conf. Info. Sci. and Sys. 2006*, pp. 277-282.
- [9] S. Wang, Y. Chen, T. Lee, and A. Helmy, "Performance Evaluations for Hybrid IEEE 802.11b and 802.11g Wireless Networks", In *Proc. 24th IEEE Inter. Conf. Performance, Computing, and Communications 2005*, pp. 111-118.
- [10] X. Zhong, H. Chan, T. Rogers, C. Rosenberg, and E. Coyle, "The Development and eStadium Testbeds for Research and Development of Wireless Services for Large-scale Sports Venues", In *Proc. TRIDENTCOM 2006*, pp. 1-9.