Voice Performance in WLAN Networks—An Experimental Study

Farooq Anjum, Moncef Elaoud, David Famolari, Abhrajit Ghosh, Ravichander Vaidyanathan,
Ashutosh Dutta, Prathima Agrawal

Telcordia Technologies, 445 South St. Morristown, NJ 07960
{fanjum, moncef, fam, aghosh, vravi, adutta, pagrawal}@research.telcordia.com

Toshikazu Kodama, Yasuhiro Katsube

Toshiba America Research Inc., 445 South St. Morristown, NJ 07960
{tkodama, ykatsube}@tari.toshiba.com

Abstract - In this work, we measure Wireless Local Area Network (WLAN) voice performance and capacity. While most WLAN applications today are data centric, the growing popularity of Voice over IP (VoIP) applications and the trend towards convergence with cellular networks will catalyze increased voice traffic. Since voice applications compete not only with each other, but also with data applications for WLAN bandwidth, quantifying voice performance and capacity in the presence of simultaneous data traffic is an important issue. We offer a practical investigation of the 802.11b MAC layer's ability to support simultaneous voice and data applications. We quantify VoIP capacity for standard WLAN networks, indicative of those already in the field, as well as evaluate the practical benefits of implementing backoff control and priority queuing at the access point. Conclusions are drawn based on an extensive set of real-world measurements conducted using off-the-shelf equipment in an experimental testbed.

I INTRODUCTION

Once only seen within the enterprise, Wireless Local Area Networks (WLANs) are increasingly making their way into residential, commercial, industrial and public areas. Examples of such environments are hotels, airports and coffee shops, which typically have a floating end user population. University campuses and conference settings also benefit from WLANs since they provide flexible connectivity and network access at reduced costs.

While the majority of traffic in WLAN deployments is data, we expect that voice will be an increasingly important application and a significant driver for WLAN adoption and integration, particularly as voice over IP (VoIP) applications flourish. Additionally, voice will be especially important in vertical industries such as construction, healthcare, and banking, etc. Therefore it is crucial to understand voice performance in WLANs. Furthermore, since WLAN endpoints share a common transmission medium, voice applications must compete with data applications for access and bandwidth. As such, voice quality and capacity can be significantly affected by the simultaneous transmission of data traffic in these networks. So it is also critical to understand the effects of data transmissions on voice performance and capacity.

We focus exclusively on IEEE 802.11b [1], the most popular and prominently deployed WLAN standard. We measure the achievable voice performance and capacity using an experimental testbed consisting of commercially available, off-the-shelf components indicative of those that have already been deployed. With such a large legacy

base for 802.11b equipment, especially among residential and enterprise customers, we believe that this approach provides the most immediately relevant results.

In addition to standard 802.11b, we investigate MAC-layer and queuing mechanisms, which can be easily implemented and can improve voice performance. Specifically, we measure the effects of backoff control and priority queuing (BC-PQ), as provided by [2]. Using both the standard and additional techniques, we determine the aggregate voice capacity by examining multiple scenarios involving various combinations of simultaneous voice and data traffic. We determine that implementing a simple BC-PQ mechanism at the access point (AP) can provide improvement in both voice performance and capacity.

The rest of the paper is organized as follows. Section II details related work. Section III presents the mechanisms under study. Section IV describes the experimental testbed. Section V discusses the experiment design. We present and analyze the measured voice performance results for each of the studied mechanisms in Section VI. The paper concludes in Section VII.

II RELATED WORK

There has been considerable previous and related work detailing the throughput performance of 802.11 networks. Most of the focus is either on analytical modeling or simulations for data transfers [3][4]. The literature describing voice application performance is relatively sparse in comparison.

The design of MAC protocols for supporting voice traffic, however, has drawn some attention [5]. The 802.11b standard provides two modes of MAC operation, namely the mandatory Distributed Coordination Function (DCF) mode and the optional Point Coordination Function (PCF) mode. The former is designed for data services and the latter for real-time services like voice. Much of the previous work with respect to voice in 802.11 networks considers simulation of the PCF mode, as in [6] and [7]. However, many manufacturers are choosing not to implement the optional PCF mode, claiming that it inhibits interoperability with other access points and does not, in fact, always allocate bandwidth better than DCF [8][9]. Moreover, previous literature is relatively silent

regarding experimental aggregate capacity, preferring to address the performance of individual voice links.

This current work is distinguished from a majority of the previous literature in four respects. First, we examine the performance of the DCF mode, as opposed to the PCF mode, to deliver voice with varying background data loads. Second, the work is entirely experimental in nature, i.e., all results are obtained from an experimental testbed using common off-the-shelf components. Third, we investigate aggregate voice capacity, in addition to performance characteristics of individual links. Finally, we measure the voice capacity resulting from the use of a modified 802.11b MAC supporting backoff control and priority queueing (BC-PQ) at the access point. As a result, our work represents more accurately current field deployments and the real-world performance attributable to BC-PQ.

III MECHANISMS

In this section we describe the two MAC-level mechanisms used in our study. The first mechanism is the standard 802.11b DCF mode. The second mechanism uses SpectraLink² NetLink SVP enhancements to provide backoff control and priority queueing (BC-PQ) at the access point. BC-PQ addresses the two shortcomings of the DCF mode in terms of real-time delivery, namely excessive backoff and content-neutral treatment of information.

A Standard 802.11b DCF

The IEEE 802.11b DCF mode is based on a “listen-before-talk” procedure, where terminals first determine if the medium is free before attempting to transmit. The DCF mode specifies two types of *Inter Frame Spacing* (IFS), including the Distributed IFS (DIFS) and the Short-IFS (SIFS). SIFS is the shortest defined interval with a value of 10 microseconds while DIFS is defined as 50 microseconds. Before transmitting a data packet, each station must wait at least a DIFS. Acknowledgement packets, however, are transmitted following a SIFS period.

A station that has a packet to transmit first senses the medium. If the medium is determined to be free for a duration of a DIFS, the station transmits the packet. Otherwise, the station enters the backoff phase in which it chooses a random backoff timer uniformly from a collection of values known as the *contention window*. The standard specifies the minimum contention window to be 32 time slots and the maximum to be 1024, where a timeslot is defined to be 20 microseconds. After a backoff time has been chosen, the station continues to monitor the medium until it observes an idle period equal to a DIFS, after which, it decrements the backoff timer after every idle timeslot. If the medium becomes busy during the countdown, the station suspends the decrement operation

until the channel becomes idle again for a period of DIFS. When the backoff timer reaches zero, the station transmits the packet.

Clearly, when multiple stations contend for the medium at the same time, the station that picks the lowest random backoff timer wins and will send its packet first.

After every unsuccessful packet transmission the size of the contention window doubles until it reaches its maximum value. Following a failed transmission the sender may attempt to retransmit the packet up to a maximum number of times before it is dropped. Following a successful transmission the contention window range is reset to its minimum value.

Negotiating medium access in this fashion can cause unbounded packet delay. Furthermore, the DCF mode provides no mechanisms for prioritizing voice traffic. In this sense it is “content-neutral” and treats all information, voice or data, equally. These two drawbacks significantly degrade voice performance.

B Backoff Control with Prioritized Queueing (BC-PQ)

BC-PQ addresses the two shortcomings of the DCF mode with respect to voice. First, it distinguishes voice packets from data packets and provides a higher priority to the voice traffic. Prioritized queueing at the AP provides preferential treatment to voice traffic without requiring client-side modifications. This may be achieved by either allocating separate prioritized queues for voice and non-voice traffic or by using a single queue with incoming voice packets being moved to the head of the queue. We provide the desired service differentiation by marking voice packets and using an AP enhanced with SpectraLink-compliant firmware. This firmware allows the AP to preferentially treat voice packets by implementing prioritized queueing on their behalf.

Secondly, in addition to priority queueing, the enhanced AP transmits voice packets using zero backoff instead of the random backoff as required by the 802.11b standard. Random backoffs introduce transmission delays and thus have an adverse effect on voice traffic. Zero backoff eliminates these delays and ensures that voice traffic will never wait needlessly in the queue.

IV EXPERIMENTAL TESTBED

We created an experimental testbed of common off-the-shelf equipment to measure the achievable performance of voice over 802.11b. This testbed is illustrated in Fig. 1 and contains two networks; a wireless network using addresses in the range 192.168.1.101/110, and a wireline network using addresses in the range 192.168.2.1/4. A machine with two IP addresses, 192.168.1.111 and 192.168.2.5, is used to route between the wired and wireless networks.

Figure 1: Experimental 802.11b testbed

C Testbed hardware

The 192.168.1.0/24 network consists of ten laptops acting as mobile hosts. All the machines use the Wi-Fi certified Orinoco Gold PCMCIA card that implements the IEEE 802.11b standard using the Direct Sequence Spread Spectrum (DSSS) physical layer. The AP is an Agere Orinoco AP-1000 with firmware version 3.95. This particular firmware version is compliant with the SpectraLink Voice Priority (SVP) protocol [2]. This system provides mechanisms by which backoff control and priority queuing are enforced at the AP.

The 192.168.2.0/24 network consists of four hosts connected via wireline Ethernet. In addition, a router is used to connect the wireless and wireline subnets. Lastly, we utilize a Toshiba server machine (Observer) to coordinate the execution of experiment runs across the testbed as well as to collect measurement information during these runs. The observer is responsible for creating data logs, and parsing those data logs to obtain experimental results.

D Testbed software configuration

All machines in the testbed have the RedHat 7.2 version of Linux running kernel version 2.4.7-10. Tcpdump, a common packet capture tool included in most Linux distributions, was used to collect and analyze packet traces. In addition to the standard Linux packages, the testbed machines have two extra applications installed: MGEN, and VGEN. MGEN is an open-source software package from the Naval Research Laboratory [10] that is used to generate simulated UDP data traffic. VGEN is an in-house tool developed to artificially generate network traffic corresponding to conversational voice as specified in [11].

V EXPERIMENT DESIGN

Our experimental goal is to isolate the MAC-layer and investigate the DCF's ability to resolve contention and packet collision to deliver quality real-time voice service.

We are more concerned with the effects of medium contention than with signal degradation due to radio propagation. As such, we operate our experiments in favorable radio environments where signal strengths are uniformly high and propagation delays are uniformly low for each terminal. This was accomplished by equidistantly placing the terminals within a 2-meter radius of the access point to ensure parity in terms of received signal strength and propagation delay.

All generated traffic involves a wireless and a wired host so that no traffic is generated between wireless hosts. In the remainder, we refer to data packets as those packets generated by MGEN, and voice packets as those packets generated by VGEN.

VGEN parameters were chosen to emulate the G.711 codec with 10ms packetization intervals, 80 byte packets and silence suppression. Each voice "call" lasted for 3 minutes and consisted of bidirectional traffic. Data packets were 512 bytes and always originated in the wireline network and thus flowed downlink from the AP. The total downlink data load was varied between 0, 1, 2, and 4 Mbps by changing the downlink packet rate. For example, to achieve a total downlink data rate of 4.0 Mbps the packet rate is adjusted to roughly 1000 data packets per second. The number of voice users in the experiments was varied from 1 to 10.

VI PERFORMANCE RESULTS

The key parameter that we use to quantify voice performance is packet loss. Studies have shown that voice conversations can tolerate up to 2% packet loss and 200ms one-way delay while still delivering acceptable voice quality [12]. We note that for our one-hop testbed network, observed end-to-end delays were well below the acceptable bounds and that the variance of inter-packet arrival times (jitter) could easily be accommodated by appropriate playout buffering. Furthermore, we found that packet dropping, most likely due to input queue size limitations, was far more prominent than excessively, or intermittently, delayed packet arrivals. Therefore, we chose packet loss as the primary voice performance metric.

The packet loss parameter is represented as a percentage of packet loss computed as:

$$PL = 100 * (1 - [\text{pkts_rcvd} / \text{pkts_sent}]),$$

where the fraction represents the ratio of received packets to total number of packets sent.

Each voice conversation contains both an uplink and a downlink component. Each voice terminal contends to transmit the uplink portion of a single voice conversation while the AP contends on behalf of all downlink voice components. Therefore, the AP contends for a significant share of the wireless medium compared to a single terminal. For m two-way voice conversations, each voice terminal contends for $1/2m$ share of the wireless medium, while the AP contends for $m/2m$, or one-half. For voice

quality to be acceptable to both participants in a voice call, the percent PL must be upper-bounded by 2% for both the uplink and downlink. Hence, in our experimental evaluation, we consider the uplink and downlink percent PLs individually.

The background data traffic contributes significantly to the downlink traffic. This is consistent with an asymmetric access model, where downlink traffic usually dominates.

E Light background traffic

Fig. 2 depicts the average uplink percent packet loss experienced by the two approaches when the background data traffic was zero and 1 Mbps. As expected, percent packet loss increases with the number of active voice terminals. For uplink traffic, we note that the standard 802.11b implementation generally outperforms the BC-PQ approach. This is anticipated, since applying zero backoff at the AP for voice traffic allows it to grab a higher share of the wireless medium, at the expense of the uplink voice terminals.

Figure 2: Uplink percent packet loss for 0 and 1 Mbps

Figure 3: Downlink percent packet loss for 0 and 1 Mbps

Fig. 3 depicts the downlink percent packet loss experienced for both mechanisms when the background traffic was zero and 1 Mbps. Unlike the uplink case, the BC-PQ scheme outperforms 802.11b. Employing priority queuing and zero-backoff control for voice traffic allows it to compete favorably with data traffic at the AP. Since the BC-PQ scheme effectively prioritizes voice traffic, no

benefits are derived when the entire offered load is from voice terminals.

Figure 4: Supported calls for 0 and 1 Mbps background traffic

We now examine the number of voice calls that can be supported by 802.11b and BC-PQ for these background traffic rates. The number of voice calls that can be supported is determined by the number of two-way voice conversations with less than 2% packet loss for both uplink and downlink voice traffic. This number is determined individually for each trial, and then averaged over all the trials.

Fig. 4 shows the number of supported voice calls for the 802.11b and BC-PQ schemes. From the graphic we note that both 802.11b and the BC-PQ scheme can support up to 10 voice users in the absence of data traffic.

For the 1 Mbps case, the BC-PQ scheme gains a slight edge over 802.11b; being able to support 8 voice users over 7 for standard 802.11b. Even at a relatively light load of 1 Mbps, downlink voice suffers in 802.11b and becomes the limiting factor in supporting voice conversations. In the BC-PQ scheme, the zero backoff and the priority queueing aid voice conversations, accounting for improved downlink performance over 802.11b.

F Heavy background traffic

We now examine increasing background traffic on voice performance. Fig. 5 shows the uplink percent packet loss for 802.11b and BC-PQ with background data rates of 2 and 4 Mbps. As expected, 802.11b outperforms the BC-PQ scheme in the uplink direction. However, the average percent packet loss for BC-PQ stays under the 2% cutoff for up to 8 voice terminals.

Fig. 6 depicts the downlink percent packet loss for 2 and 4 Mbps data traffic. The BC-PQ scheme significantly outperforms standard 802.11b on the downlink with increasing data traffic. Voice traffic suffers high percent packet losses in 802.11b since the voice packets compete unfavorably with more, larger data packets at the AP. As the AP receives packets at a higher rate, its queue fills more quickly increasing downlink packet losses. Priority queuing for voice traffic at the access point ensures that voice traffic is not adversely affected on the downlink.

Thus, even for high offered loads of 4 Mbps in the downlink direction, the BC-PQ scheme allows for excellent voice performance.

Figure 5: Uplink percent packet loss for 2 and 4 Mbps background traffic

Figure 6: Downlink percent packet loss for 2 and 4 Mbps background traffic

Figure 7: Supported calls for 2 and 4 Mbps background traffic

Fig. 7 shows the number of voice calls with acceptable percent packet loss for the 802.11b and BC-PQ schemes. For 2 Mbps background data traffic, the two schemes perform comparably in the presence of 5 voice users. However, beyond this point, 802.11b performance degrades considerably, due to the strain on downlink voice conversations. Further, we note that in the 4 Mbps case, the high data rate precludes the possibility of any voice conversations in the downlink direction for 802.11b. The

BC-PQ scheme, however, continues to support between 5 and 6 voice calls, even under these heavy load conditions.

VII CONCLUSIONS

We described an experimental 802.11b testbed comprising currently available hardware and software that was used to investigate voice performance in WLANs. We measured the link performance of high-quality voice service in standard 802.11b in the presence of varying simultaneous data loads. These results were used to determine the experimental AP voice capacity afforded by the DCF mode. Furthermore, we implemented a Backoff Control and Priority Queuing (BC-PQ) mechanism at the AP and measured the resulting voice performance and capacity. BC-PQ distinguishes delay-sensitive packets and provides them priority treatment by improving their queue position and allowing them to use a zero backoff value during contention. Measurements show that these two effects combine to provide noticeable improvements to voice performance and capacity.

REFERENCES

- [1] IEEE standard for Wireless LAN-Medium Access Control and Physical layer Specification, P802.11, Nov. 1999.
- [2] SpectraLink Voice Priority: <http://www.spectralink.com/products/> (visited October 14, 2002).
- [3] G. Bianchi, "Performance analysis of the IEEE 802.11 Distributed Coordination Function", IEEE Journal on Selected Areas in Communications, Vol. 18. No.3 March 2000
- [4] Y. C. Tay, K. C. Chua, "A capacity analysis for the IEEE 802.11 MAC protocol", Wireless Networks March 2001 Volume 7 Issue 2
- [5] I. Akyildiz et al, "Medium access control protocols for multimedia traffic in wireless networks", IEEE Net. Mag., vol. 13, no 4, pp39-47, Jul/Aug 1999.
- [6] M. Veeraraghavan, N. Cocker and T. Moors, "Support of voice services in IEEE 802.11 wireless LANs", IEEE Infocom 2000.
- [7] A. Koepsel and A. Wolisz, "Voice transmission in an IEEE 802.11 WLAN Based Access Network", In Proc. of WoWMoM 2001, pp. 24-33, Rom, Italy, July 2001
- [8] A. Koepsel, J.-P. Ebert, and A. Wolisz, "A Performance Comparison of Point and Distributed Coordination Function of an IEEE 802.11 WLAN in the Presence of Real-Time Requirements", In Proc. of 7th Intl. Workshop on Mobile Multimedia Communications (MoMuC2000), Tokio, October 2000
- [9] Orinoco Product FAQ, "Does Agere ORINOCO Wireless Network support PCF(Point Coordination Function)?" <http://www.orinocowireless.com/template.html?section=m63&envelope=131&page=2034> (visited October 14, 2002).
- [10] MGEN The Multi-Generator Toolset: <http://manimac.itd.nrl.navy.mil/MGEN/> (visited October 14, 2002).
- [11] ITU-T Recommendation P.59 Artificial Conversational Speech
- [12] ETSI TR 101 329-6 V2.1.1, "Telecommunications and Internet Protocol Harmonization Over Networks (TIPHON) Release 3; End-to-end Quality of Service in TIPHON systems; Part 1: General aspects of Quality of Service (QoS)