

Outdoor IEEE 802.11 Cellular Networks: MAC Protocol Design and Performance

Kin K. Leung, Bruce McNair, Leonard J. Cimini, Jr., and Jack H. Winters

AT&T Labs - Research
Middletown, NJ 07748

Abstract — We explore the feasibility of designing an outdoor cellular network based on the IEEE 802.11 specification. Since the standard is intended for wireless local-area networks (WLAN), there are many technical challenges when applying the air interface to the outdoor environment. We study here how the 802.11 medium access control (MAC) protocol can be applied and how it performs in the outdoor network. By exploiting the fact that timeout intervals are not explicitly specified, without modifying the standard, we propose a new timing structure for the distribution coordination function (DCF) and the handshake of request-to-send (RTS) and clear-to-send (CTS) to handle increased signal propagation delay in the outdoor network. We find that the DCF and RTS/CTS protocols as specified in the standard continue to work properly for a link distance up to 6 km. Our analysis reveals that the DCF performance degrades slightly in the 802.11 network with cell size of 6 km when compared with the 600 m WLAN. Thus, as far as the MAC protocol is concerned, the 802.11 outdoor, cellular network with 6 km cell size is feasible.

I. INTRODUCTION

While the wireless industry is actively developing, testing and deploying third generation (3G) wireless networks, customers are expecting services with data rate higher than that to be provided by 3G networks. To meet such demand for better services, many companies have started to provide high-speed data services using wireless local-area-networks (WLAN) in places such as airports and hotels. Such an approach is particularly attractive due to the maturity and low cost of the IEEE 802.11b technology [I99b, VAM99]. The 802.11b network provides data rates up to 11 Mbps, far exceeding that to be offered by, for example, EDGE [SAE98, CQW99] and W-CDMA networks [HT00].

Besides high data rates, 802.11b networks offer several advantages over 3G networks. First, the cost of 802.11b equipment is much lower than that for 3G equipment because of the simple design of the former networks, coupled with competition among WLAN vendors. Second, 802.11b networks operate in the 2.4 GHz ISM band, which is free spectrum. In contrast, the 3G spectrum is licensed and very expensive. Thus, both reasons make the operating cost of the 3G network higher than that for the WLAN.

On the other hand, each WLAN can serve only a small area, up to a few hundred meters, where a cell radius of ten kilometers is supported in the 3G networks. In addition, future 3G networks are expected to provide ubiquitous coverage and availability. In contrast, public WLAN service is available only in isolated places such as airports and hotels. Users will

use both types of networks, one for excellent coverage while the other for enhanced data rates.

In this research, we explore the following question: Is it possible to design an outdoor, cellular network based on the existing 802.11 air-interface standard for wireless data services? If the answer is affirmative, then users can use the same air interface mechanism to obtain wireless services from indoor WLAN and outdoor 802.11 networks. There are many technical issues pertinent to the design of an 802.11 cellular network. Recall that 802.11 as well as its extension 802.11b [I99b] and 802.11a [I99a] standards were developed specifically for WLAN with the transmission range up to a few hundred meters in indoor environment. First, the signal propagation delay increases when applying the 802.11 to outdoor networks relative to the indoor WLAN, which in turn may affect the applicability of the medium access control (MAC) protocol. Second, the outdoor environment has increased delay spread that causes intersymbol interference. Further, Doppler effects due to mobility may require sophisticated processing for channel estimation.

The focus of this paper is on the MAC protocol design and performance when using the 802.11 specification for outdoor, cellular networks, while radio issues will be addressed in our subsequent papers. Much work related to the 802.11 MAC protocol has been published; see e.g., [B00], [CCG00] and [VCM01]. The organization of the rest of this paper is as follows. We provide a brief description of the 802.11 MAC protocols in Section II. In Section III, we discuss how the protocols may or may not work properly in the outdoor networks. In addition, we estimate the maximum cell radius in outdoor networks due to the consideration of MAC protocols. Section IV analyzes the MAC performance for outdoor networks. Finally, our conclusion is in Section V.

II. IEEE 802.11 MAC PROTOCOLS

The IEEE 802.11 specification [I97] allows three kinds of physical layer: direct sequence spread spectrum (DSSS), frequency hopping spread spectrum (FHSS) and infrared (IR). In particular, the DSSS design supports data rates of 1 and 2 Mbps. Subsequently, while maintaining backward compatibility to the DSSS 802.11 specification, the 802.11b was adopted to support data rates of 5.5 and 11 Mbps, operating in the 2.4 GHz band (the ISM band). As a result, the 802.11b network can support 1, 2, 5.5 and 11 Mbps, depending on radio conditions. Another extension is 802.11a, which uses an entirely different physical layer known as orthogonal frequency division multiplexing (OFDM). 802.11a

can support data rates ranging from 6 to 54 Mbps, operating in the 5.5 GHz band (the U-NII band). It is important to note that it is the 802.11b networks that have been widely used recently. For this reason, we focus on 802.11b networks here. We also note that although data rates have been increased, 802.11b networks continue to use the original MAC protocol in the 802.11 specification. Furthermore, the MAC protocol supports the independent basic service set (BSS), which has no connection to wired networks (i.e., an ad-hoc wireless network), as well as an infrastructure BSS, which includes an access point (AP) connecting to a wired network. The latter is similar to cellular networks with base stations replaced by AP's. We consider only the infrastructure BSS in this paper.

We provide a brief description of the 802.11 MAC protocol here [I97, OP99]. The 802.11 specification defines five timing intervals for the MAC protocol. Two of them are considered to be basic ones that are determined by the physical layer: the short interframe space (SIFS) and the slot time. The other three intervals are defined based on the two basic intervals: the priority interframe space (PIFS) and the distributed interframe space (DIFS), and the extended interframe space (EIFS). The SIFS is the shortest interval, followed by the slot time. The latter can be viewed as a time unit for the MAC protocol operations, although the 802.11 channel as a whole does not operate on a slotted-time basis. For 802.11b networks (i.e., with a DSSS physical layer), the SIFS and slot time are 10 and 20 μ s, respectively. The slot time of 20 μ s is chosen to account for the signal propagation and processing delays. The PIFS is equal to SIFS plus one slot time, while the DIFS is the SIFS plus two slot times. The EIFS is much longer than the other four intervals, and is used if a data frame is received in error.

The 802.11 MAC supports two modes of operation: the Point Coordination Function (PCF) and the Distributed Coordination Function (DCF). The PCF provides contention-free access, while the DCF uses the carrier sense multiple access with collision avoidance (CSMA/CA) mechanism for contention based access. The two modes are used alternately in time. That is, a contention-free period by the PCF is followed by a contention period of the DCF.

A. The PCF Protocol

In the PCF protocol, an AP polls its associated mobile stations one after another by sending polling messages. If the AP has data to send to a mobile station being polled, the data can be included in the polling message. If the polled station has data for the AP, it is sent in the response message. When applicable, an acknowledgment (which acknowledges receipt of a previous data frame from the AP) can also be included in the response message.

As an illustrative example in Figure 1, the AP first sends the polling message and data, if any, to mobile station 1 (denoted by S1). Station 1 should immediately send an acknowledgment or a data frame, if any, to the AP within the SIFS interval. After receiving an ACK or data from station 1, the AP polls mobile station 2 within the SIFS interval. In this

illustration, station 2 does not respond, either because the polling message is lost or station 2 has no data to send to the AP. In this case, as a response is not received from station 2 before the SIFS expires, the AP moves on to poll station 3 within the PIFS interval, which starts from the end of the last polling message for station 2.




Figure 1. The PCF of the 802.11 MAC Protocol

B. The DCF Protocol

The DCF employs the CSMA/CA mechanism and works as follows. A station (including the AP) with a new packet ready for transmission senses whether or not the channel is busy. If the channel is detected idle for a DIFS interval (i.e., 50 μ s for 802.11b networks), the station starts packet transmission. Otherwise, the station continues to monitor the channel busy or idle status. After finding the channel idle for a DIFS interval, the station: a) starts to treat channel time in units of slot time, b) generates a random backoff interval in units of slot time, and c) continues to monitor whether the channel is busy or idle. In the latter step, for each slot time where the channel remains idle, the backoff interval is decremented by one. When the interval value reaches zero, the station starts packet transmission. During this backoff period, if the channel is sensed busy in a slot time, the decrement of the backoff interval stops (i.e., is frozen) and resumes only after the channel is detected idle continuously for the DIFS interval and the following one slot time. Again, packet transmission is started when the backoff interval reaches zero. The backoff mechanism helps avoid collision since the channel has been detected to be busy recently. Further, to avoid channel capture, a station must wait for a backoff interval between two consecutive new packet transmissions, even if the channel is sensed idle in the DIFS interval. This is depicted in Figure 2.




Figure 2. The DCF of the 802.11 MAC Protocol

The backoff mechanism for the DCF is an exponential one. For each packet transmission, the backoff time in units of slot time (i.e., an integer) is uniformly chosen from 0 to $n-1$, where n depends on the number of failed transmissions for the packet. At the first transmission attempt, n is set to a value of $CW_{min}=32$, the so-called minimum contention window. After

each unsuccessful transmission, n is doubled, up to a maximum value of $CW_{\max}=1024$.

The 802.11 specification requires a receiver to send an ACK for each packet that is successfully received. Furthermore, to simplify the protocol header, an ACK contains no sequence number, and is used to acknowledge receipt of the immediately previous packet sent. That is, stations exchange data based on a stop-and-go protocol. As shown in Figure 2, the sending station is expected to receive the ACK within the $10 \mu\text{s}$ SIFS interval after the packet transmission is completed. If the ACK does not arrive at the sending station within a specified ACK_timeout period, or it detects transmission of a different packet on the channel, the original transmission is considered to have failed and is subject to retransmission by the backoff mechanism.

In addition to the physical channel sensing, the 802.11 MAC protocol implements a network allocation vector (NAV), whose value indicates to each station the amount of time that remains before the channel will become idle. All packets contain a duration field and the NAV is updated according to the field value in each packet transmitted. The NAV is thus referred to as a virtual carrier sensing mechanism. The MAC uses the combined physical and virtual sensing to avoid collision.

The protocol described above is called the two-way handshaking mechanism. In addition, the MAC also contains a four-way frame exchange protocol. Essentially, the four-way protocol requires that a station send to the AP a special, Request-to-Send (RTS) message, instead of the actual data packet, after gaining channel access through the contention process described above. In response, if the AP sees that it is appropriate, it sends a Clear-to-Send (CTS) message within the SIFS interval to instruct the requesting station to start the packet transmission immediately. The main purpose of the RTS/CTS handshake is to resolve the so-called hidden terminal problem.

III. MAC PROTOCOLS IN OUTDOOR NETWORKS

A. The PCF Protocol Infeasible

It is important to emphasize that the SIFS and PIFS timing requirements for the PCF in Figure 1 are clearly defined in the standard. In particular, the most stringent requirement is that the ACK has to be received from the polled station to the AP within the SIFS interval, which is $10 \mu\text{s}$ for 802.11b networks. When the standard is used for outdoor, cellular networks, the distance between a mobile station and its AP is expected to be longer than that in the WLAN. Consider a link distance of 1.5 km as an example. The round-trip signal propagation delay for the 1.5 km distance requires $10 \mu\text{s}$. Since at least several μs are needed for signal processing at the receiver, the link distance is likely to be limited to hundreds of meters, as in WLAN environments. In fact, this is the intention of the 802.11 specification. Thus, it is unrealistic to expect that the PCF can be supported for 802.11 outdoor networks with cell radius of several km.

B. Applicability of the DCF Protocol

Let us consider the DCF in the outdoor networks. It is worth noting that as far as the MAC protocol is concerned, the major difference between 802.11 outdoor networks and their WLAN counterparts is increased signal propagation delay. As shown in Figure 2, the major constraint for the applicability of the DCF in outdoor networks is that the ACK is expected to be received within the SIFS interval ($10 \mu\text{s}$) after packet transmission. That is, the $10 \mu\text{s}$ includes the round-trip signal propagation and processing at the receiver. However, in order to be useful, we aim at having an outdoor cell size of several km. Thus, the one-way signal propagation delay can be more than $10 \mu\text{s}$, even neglecting the return propagation and processing time. Evidently, this would not be practical without violating the protocol specification. Our solution is based on the following key observation: Typically, there is no consequence if the ACK is received later than the SIFS interval. This is because, after a station transmits a packet, it starts an ACK_timeout period, which is not specified in the standard and is usually chosen to be a value much larger than $10 \mu\text{s}$ by vendors. Thus, as long as the ACK is received before the timeout expires, the MAC protocol responds properly.

As in typical implementations, we assume that the ACK_timeout period is longer than the DIFS interval of $50 \mu\text{s}$. Then, we argue that as long as the ACK arrives at the sending station within the DIFS interval following a packet transmission, the DCF operates properly in the outdoor network environment where the link distance can reach as much as several km. The reasoning is as follows. First, because the ACK is received within the DIFS interval, the ACK_timeout has not expired so that the protocol can respond upon receipt of the ACK as if it were received within the SIFS interval, as originally specified in the protocol standard. Second, since the DCF protocol requires any station to sense the channel being idle for at least the DIFS interval before transmitting, the return of the ACK within the DIFS interval following the previous packet transmission by the sending station prevents any stations other than the receiving one from gaining access to the channel. Consequently, the channel is implicitly “reserved” for the receiving station to send the ACK. In addition, the pairing of a packet transmission and its ACK transmitted in sequence for any pair of sending and receiving stations remains intact, as required by the specification.

Extending the arrival delay of ACK from the SIFS to the DIFS interval comes with a penalty. That is, the computation of the NAV assumes that the ACK returns within the SIFS interval. So, the delay extension causes an erroneous determination of the NAV, thus incorrect virtual sensing. Nevertheless, since protocol operations are based on both physical and virtual channel sensing, as long as the former works properly, the malfunctioning of the virtual sensing due to incorrect NAV value causes no apparent, negative impacts.

Actually, the extension of the ACK arrival delay from the SIFS interval to the DIFS interval can also be applied to the RTS and CTS handshake for resolving the hidden terminal

problem. Specifically, a sending station starts a CTS_timeout period after sending an RTS. The MAC protocol specifies that the CTS, if any, is supposed to arrive from the receiving station within the SIFS interval (10 μ s). However, similar to the ACK_timeout, the CTS_timeout period is typically chosen to be much longer than 10 μ s by equipment manufacturers. Therefore, by the same arguments discussed above, the arrival delay for the CTS can be extended to the DIFS interval.

A. Maximum Cell Size for the DCF Protocol

With the arrival delay for the ACK and CTS extended to the DIFS interval, let us consider its limit on the maximum cell size (i.e., link distance) in outdoor 802.11 networks.




Figure 3. Allocation of ACK/CTS delay

Recall that the ACK and CTS arrival delay consists of a round-trip signal propagation delay and signal processing time. As shown in Figure 3, one reasonable allocation of the 50 μ s DIFS delay is: a one-way signal propagation delay of 20 μ s and a processing time of 10 μ s at the receiving station. The latter should not cause a processing burden for the receiver because the original delay of the SIFS interval is 10 μ s. For the 20 μ s propagation delay, the maximum cell size is about 6 km. In other words, with the cell size of 6 km or less, the DCF protocol operates properly in 802.11 cellular networks.

IV. DCF PERFORMANCE IN 802.11 OUTDOOR NETWORKS

We present an approximate analysis of the DCF throughput for outdoor networks and WLAN. As shown in Figure 3, if a station with a packet for transmission senses the channel idle for the DIFS interval (denoted by d in μ s in the following), it starts to transmit. Following the packet transmission, the channel remains idle for the DIFS interval and then the ACK is transmitted by the receiver. If the sending station senses the channel busy, it goes through the backoff mechanism discussed above. For simplicity, we do not model the details of the backoff algorithm. Instead, it is assumed that the aggregated traffic, which includes new packets and transmission reattempts, from all stations forms a Poisson process with an intensity of G packets/ μ s. This assumption is reasonable if the backoff period is sufficiently long so that new transmission and reattempts become independent sources.

For simplicity, assume that the signal propagation delay a in μ s is identical between any pair of stations. Thus, the *vulnerable period* is also given by a , during which a new packet transmission cannot be sensed by other stations. As a result, these stations under the CSMA protocol can possibly start their own transmissions and cause collisions. Each station senses the channel idle for d μ s (DIFS interval) before

transmitting. The packet transmission time is assumed to be constant L μ s. Consider the channel activity for a successful packet transmission. The channel is idle for d μ s and followed by packet transmission of L μ s. As Figure 3 shows, the transmitter waits for d μ s (DIFS interval) for the ACK. Let the ACK transmission time be c μ s. The channel is sensed idle again by all stations a μ s after the ACK transmission.




Figure 4. Busy period with Collided Transmissions

Figure 4 shows a typical busy period with collided transmissions due to the vulnerable period for the CSMA protocol, where Y denotes the time span between the first and the last packet transmissions in the busy period. Using the result in [K76], the average duration of Y is given by

$$\bar{Y} = a - \frac{1 - e^{-aG}}{G}. \quad (1)$$

The average length of a busy period (which contains a successful transmission or collisions) is given by

$$\bar{B} = d + \bar{Y} + L + a + (d + c)e^{-aG} \quad (2)$$

where the last term accounts for the waiting and transmission time of the ACK for successful transmission with probability e^{-aG} , based on the Poisson assumption of aggregated traffic. By the same assumption, the average cycle time, consisting of a busy period and the following idle period, is given by

$$\bar{T} = d + \bar{Y} + L + a + (d + c)e^{-aG} + \frac{1}{G} \quad (3)$$

The channel throughput S is defined as the fraction of time at which data is successfully transmitted. Thus, we have

$$S = \frac{Le^{-aG}}{\bar{T}} \quad (4)$$

where the numerator is the average amount of time when data is transmitted without collision and \bar{T} is obtained from (3).

Three common packet sizes of 60 bytes (e.g., TCP ACK), 576 bytes (typical size for web browsing) and 1500 bytes (the maximum size for Ethernet) plus a 34 byte 802.11 MAC header are considered. For an 802.11 network with a 1 Mbps data rate, the corresponding transmission time L is 0.75, 4.88 and 12.27 msec, respectively. The sensing idle time of the DIFS interval of 50 μ s and the transmission time c for the 112-bit ACK is 0.112 μ s. Based on our discussions above, the link distance is assumed to be 6 km, and thus the one-way

propagation delay a is 20 μ s. For comparison, we also consider a WLAN with a service radius of 600 m with a signal propagation delay of 2 μ s. In this WLAN, after packet transmission, a station waits for the SIFS interval of 10 μ s as in the standard, instead of the DIFS interval as shown in Figure 3, for the arrival of the associated ACK.

Applying these parameters to (1) to (4), we obtain in Figure 5 the MAC throughput as a function of the aggregated traffic load for selected packet lengths. As expected, when the link distance increases from 600 m to 6 km for a given packet length, the maximum throughput decreases because of the increased signal propagation delay and thus the vulnerable period. For the 576-byte packet size, the maximum throughput drops from 92.9% to 84.8%, when the link distance increases from 600 m to 6 km. Nevertheless, since a 576-byte size is typical for popular web applications, the throughput of 84.8% is still satisfactory. For 1500-byte packets, the channel throughput for the 6 km cell can reach a maximum of 90.8%. Even for the short TCP ACKs of 60 bytes long, the channel throughput is about 60%. In summary, the MAC throughput is still satisfactory despite the increase of cell size to 6 km.




Figure 5. MAC Throughput Comparison.

V. CONCLUSION AND FUTURE WORK

We have studied how the 802.11 MAC can be applied and how it performs in outdoor networks. By exploiting the fact that timeout intervals are not explicitly specified, without modifying the standard, we have proposed a new timing structure for the distribution coordination function (DCF) and the handshake of request-to-send (RTS) and clear-to-send (CTS) to handle increased signal propagation delay in the 802.11 outdoor network. It was found that the DCF and RTS/CTS protocols as specified in the standard continues to work properly if the cell radius is less than 6 km. Our analysis reveals that the DCF performance degrades slightly for a cell size of 6 km when compared with the 600 m WLAN. Thus, as far as the MAC protocol is concerned, the 802.11 cellular network with a cell size of 6 km is feasible.

In terms of future work, a major issue is to examine and enhance the 802.11 radio design so that it performs properly in the cellular environment. In a companion paper [CLMK01], we shall address the issue of radio link performance in the

802.11 cellular network. We also plan to investigate techniques such as advanced equalizers, smart antennas and call admission control to further improve the performance of the outdoor 802.11 cellular networks.

ACKNOWLEDGMENTS

We would like to thank Martin Clark, Peter Driessen, Zoran Kostic and Stefan Muller-Weinfurter for their discussion.

REFERENCES

- [B00] G. Bianchi, "Performance Analysis of the IEEE 802.11 Distributed Coordination Function," *IEEE J. Selected Areas in Commun.*, Vol. 18, March 2000, pp. 535-547.
- [CCG00] F. Cali, M. Conti and E. Gregori, "Dynamic Tuning of the IEEE 802.11 Protocol to Achieve a Theoretical Throughput Limit," *IEEE/ACM Trans. on Networking*, Vol. 8, Dec. 2000, pp. 785-799.
- [CLMK01] M.V. Clark, K.K. Leung, B. McNair and Z. Kostic, "Outdoor IEEE 802.11 Cellular Networks: Radio Link Performance," *Proc. IEEE ICC 2002*.
- [CQW99] J. Chuang, X. Qiu and J. Whitehead, "Data Throughput Enhancement in Wireless Packet Systems by Improved Link Adaptation with Application to the EDGE System," *Proc. of IEEE VTC'99*, Sept. 1999.
- [HT00] H. Holma and A. Toskala (Ed.), *WCDMA for UMTS*, John Wiley & Sons, New York, 2000.
- [I97] IEEE 802.11, "Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specification," 1997.
- [I99a] IEEE 802.11a, "Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specification: High-Speed Physical Layer Extension in the 5 GHz Band," 1999.
- [I99b] IEEE 802.11b, "Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specification: High-Speed Physical Layer Extension in the 2.4 GHz Band," 1999.
- [K76] L. Kleinrock, *Queueing Systems, Vol. 2: Computer Applications*, John Wiley & Sons, 1976, p. 398.
- [OP99] B. O'Hara and A. Petrick, *IEEE 802.11 Handbook*, IEEE Press, New York, 1999.
- [SAE98] P. Schramm, et. al., "Radio Interference Performance of EDGE, a Proposal for Enhanced Data Rates in Existing Digital Cellular Systems," *Proc. of IEEE VTC*, Ottawa, Canada, May 1998, pp. 1064-1068.
- [VAM99] R. van Nee, et. al., "New High-Rate Wireless LAN Standard," *IEEE Commun. Mag.*, Dec. 1999, pp. 82-88.
- [VCM01] M. Veeraraghavan, N. Cocker and T. Moors, "Support of Voice Services in IEEE 802.11 Wireless LANs," *Proc. of IEEE Infocom 2001*, Anchorage, Alaska, April 2001, pp. 488-497.