The Impact of Channel Usage Information on the Throughput Achieved by 802.11-Style MACs in Urban Mesh Networks

Jonghyun Kim Stephan Bohacek
kim@eecis.udel.edu bohacek@udel.edu
University of Delaware

Abstract— **It is well known that in dense mesh networks, CSMA-based MACs such as IEEE 802.11 achieve lower throughput than optimal spatial TDMA. This paper explores the degree to which the difference in throughput is partly due to a lack of channel usage information and the degree to which the difference is due to packet ordering. To this end, the throughput achieved by a large number of hypothetical MAC algorithms is studied via simulation of a dense urban mesh network. The MACs algorithms explored range from simple modifications of IEEE 802.11 to schemes where global information is assumed to be known.**

I. INTRODUCTION

The principal function of a MAC protocol is to determine when a node can transmit. The ways in which different MAC algorithms make this decision ranges from using no information and coordination (e.g., Aloha) to fully scheduled spatial TDMA (STDMA). Recently, an algorithm has been developed that can tractably compute optimal STDMA for large mesh networks (e.g., 500 nodes) even when co-channel interference arises [2]. It is expected that optimal STDMA will achieve higher throughput than CSMA-based MACs such as 802.11. But the size and the reasons for the difference in throughput are not known.

Perhaps the most obvious reason that optimal STDMA achieves higher throughput than 802.11 is that packets transmitted with 802.11 might experience collisions. These collisions are a result from 802.11's limited channel knowledge. Another possible reason that STDMA achieves higher capacity is that optimal STDMA carefully selects the order that packets are transmitted. For example, consider the topology shown in Figure 1. Assume that nodes cannot receive or transmit when a neighbor is sending or receiving. Furthermore, assume that the data flows are from the base station (BS) to destinations A, B, C, and D. Without loss of generality, assume that the base station first transmits a packet with destination A. Now, the base station has the choice of whether the next packet should be the one with destination B, C, or D. If B is selected, then the base station cannot transmit until the packet with destination A has reached A (node X cannot send or receive while W is transmitting). On the other hand, if the next packet has destination C or D, then the transmission can start as soon as the first packet reaches node W. Thus, the base station is able to transmit 50% sooner if the destination of the next packet is C or D, rather than B. Thus, besides eliminating collisions, optimal STDMA improves throughput by ordering the packets in an optimal way.

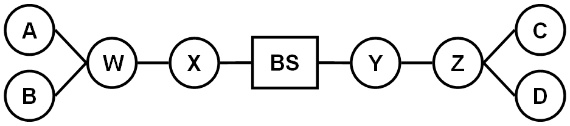

Fig. 1. An exmaple topology where the order of packet transmissions impacts throughput.

If the main reason that optimal STDMA achieves higher throughput than 802.11 is due to a lack of channel usage information, MAC protocols such as TRAMA [5] could be developed to avoid or even eliminate collisions. Moreover, there are a wide range of techniques that can be employed to reduce collisions. On the other hand, if the main reason that optimal STDMA achieves higher throughput than 802.11 is that optimal STDMA carefully selects the ordering of packets, then substantial improvements in capacity will require computationally complex algorithms. Interestingly, while there has been extensive work focused on collisions (e.g., [5]), there has been little effort focused on ordering packets.

In order to provide insight into why optimal STDMA provides higher throughput than 802.11, this paper studies variants of 802.11 along with 802.11-style MACs that have extra channel usage information. The types of channel usage information considered include global channel usage information (i.e., information as to which nodes are transmitting and receiving in the entire network) and regional channel usage information, which indicates whether the nodes in a specific region are transmitting and, perhaps, receiving. This paper does not focus on the way in which nodes obtain the channel usage information. In fact, in some cases, the channel usage information is simply assumed to be from a genie, i.e., perfect information is achieved without any overhead. Thus, the results of this study are a guide as to what types of MAC algorithms might improve the capacity.

This study focuses on the urban mesh network setting. Specifically, it is assumed that there are a set of wireless routers and a set of gateways. The wireless routers have single wireless interfaces, and the gateways have wireless and wired interfaces; gateways provide connectivity to the Internet. This study uses realistic urban propagation models, and hence, the conclusions are directly applicable to the mesh networks currently being deployed throughout the world. The motivation for focusing on mesh networks is that since mesh nodes are powered, not mobile, and under the control of a single administrator, it is conceivable that mesh nodes could measure and distribute channel gains as well as share channel usage information.

Some of the conclusions of this study are as follows.

Optimal STDMA achieves 50-350% higher throughput than standard 802.11 without RTS/CTS, where the amount of the improvement depends on the number of wireless routers and gateways. If 802.11 had global channel usage knowledge, then the throughput is approximately halfway between standard 802.11 and optimal STDMA. Thus, in a sense, half of the difference between 802.11 and optimal STDMA is due to collisions, while the rest is due to ordering of packets. As an alternative to global information, if a node has perfect knowledge of nearby transmissions and receptions, then the throughput is approximately halfway between standard 802.11 without RTS/CTS and 802.11 with global channel usage information. Techniques to improve 802.11 in other ways (e.g., using RTS/CTS and setting the value of CCA) meet with limit success. In fact, we find that aloha provides only slightly lower throughput than 802.11 without RTS/CTS and higher throughput than 802.11 with RTS/CTS.

This rest of the paper has two basic parts, descriptions of the MAC algorithms investigated, and the results of the simulations.

II. THROUGHPUT METRIC

This study focuses on the download throughput, that is, all flows originate at a gateway and terminate at wireless routers. We assume that each wireless router has a path from a single gateway. Each gateway transmits F bps to each of the wireless routers that it has paths to. Let $g_{\phi}(F)$ be the average arrival rate at destination ϕ when the gateway(s) transmit at rate F. The metric of interest in this study is the max-min throughput,

$$\max_{F} \min_{\phi} g_{\phi}(F). \tag{1}$$

Employing the techniques described in [2], it is possible to compute the max-min throughput. For the MAC algorithms described below, the max-min throughput is determined by slowly increasing F, the sending rate to each destination, until the minimum arrival rate $\min_{\phi} g_{\phi}(F)$ reaches its peak. Using bootstrapping, confidence intervals were used to ensure that the $\min_{\phi} g_{\phi}(F)$ was estimated with less than a 10% error.

Note that the sending rate to each destination is the same, so unfairness does not arise. Thus, this study only examines the throughput of 802.11.

III. DATA-RATE SELECTION

As is well known, the bit-rate used for transmission has a significant impact on performance. In the case of mobile wireless networks, selecting a transmission bit-rate is complicated by the variability of the channel that results from the mobility. Thus, several adaptive bit-rates schemes have been developed. One such scheme is automatic rate fallback (ARF), which is a widely deployed scheme [4]. However, in the case of mesh networks, it is possible to make long-term observations and determine a good bit-rate for each link. To mimic this ability, the effective data rate is selected as follows. Let $PSP(SNR, R)$ be the probability of successfully decoding a packet transmitted at bit-rate R when the signal-to-noise ratio at the receiver is SNR, thus the effective data rate is

$$\text{Data-Rate} = \arg\max_{R} R \times PSP(SNR - Guard, R), \tag{2}$$

where $Guard$ is used to reduce sensitivity to interference. The results shown in this paper use $Guard = 3$ dB, however, other simulations indicate that the value of $Guard$ does not have a significant impact on throughput.

IV. IEEE 802.11

Details regarding 802.11 can be found in [3]. Briefly, in 802.11, a node will transmit a frame if the channel is idle and certain timers have expired. Specifically, once the channel becomes *idle*, a node will wait a time period denoted as DIFS. Once the channel remains idle for a continuous period of DIFS, the node will decrement its back-off counter by one every *time slot*. This decrementing continues until the channel is no longer idle or the back-off counter is zero, in which case, the node transmits a packet.

Ideally, a node will declare that its channel is idle if, to the best of its knowledge, its transmission would be successful and its transmission would not cause a failure of an ongoing transmission. In general, it is not possible to make a precise determination without global channel usage information. As an approximation, a node will listen to the channel. If some other node transmission is decoded, then the duration field within the MAC header can be used to determine that the channel is busy. In the case that RTS and CTS frames are used, the duration in the MAC header includes the time required for future frame transmissions. If the MAC header is not decoded, then the channel is declared to be busy if the aggregate interference exceeds the value of the clear channel assessment (CCA) parameter.

Note that 802.11 has several ways to adjust the region from which channel usage is learned. For example, if lower bit-rates for data or control frames are used, then nodes in a larger region are able to decode the MAC header and learn the duration of transmissions. If RTS/CTS is used, then nodes near to the receiver become aware of the transmission. Again, if the bit-rate of the RTS and CTS frames is lowered, then more nodes become aware of the channel usage. Similarly, if the value of CCA is lowered, then the region over which a node gathers channel usage information is increased.

A drawback of lowering the bit-rate for the data frame in order to expand the region over which channel usage information is gathered is that a lower bit-rate increases the time to transmit packets. As an alternative to transmitting data frames at a low bit-rate, a node may transmit a CTS-to-self before transmitting the data frame. This CTS-to-self can be transmitted at a low bit-rate in order to have a large region over which channel usage information is gathered, and yet the data frame can be transmitted at a high bit-rate.

In summary, IEEE 802.11 allows the region over which channel usage information is gathered to be controlled by

- selecting between using RTS/CTS, using CTS-to-self, or just transmitting the data frame,
- adjusting the bit-rates of the various frames,
- adjusting the value of CCA.

A. IEEE 802.11 Variants

1) Stomp: As mentioned above, when a MAC header of a transmission is decoded, the duration of the transmission (and related transmission) is learned by the MAC. However, in this case, the node may still decide that the channel is *not* busy. For example, Atheros chip-sets have the ability to "*stomp*," which allows the channel to be declared to be idle when a MAC header is decoded, but the MAC header indicates that the communication is between nodes in a different BSS. While Atheros chip-sets allow the distinction only by BSSID, this approach could be extended so that a node can declare the channel to be idle depending on the source and destination in the decoded MAC header as well as on the destination that the node desires to transmit to. Specifically, if all channel gains are known, then, once the MAC header is decoded, the node can determine whether its transmission would succeed and whether its transmission would interfere with the transmission that it decoded. If the transmission could proceed, then the channel is declared to be idle. Of course, when the channel is idle, the node may wait DIFS, decrement its back-off timer, or transmit, according to the 802.11 specification.

We define stomp to be the ability to use the channel usage information that includes the source and destination of a transmission and determination of successful transmission without causing a collision as discussed above. A complicating aspect of 802.11 is that the state of the channel is only known if the MAC header is decoded. Thus, if the frame is transmitted at a high data rate, then it might not be possible for a node to decode the MAC header and the detailed channel usage information is unknown to the MAC. Thus, it is assumed that this channel usage information is in the PLCP header (which is transmitted at a low bit-rate).

2) Packet Reordering: Stomp clearly allows the channel to be used more efficiently. However, we have found that in some scenarios, the impact of stomp is reduced by head-of-the-line blocking. Recall that stomp uses the destination of the frame to be transmitted in order to determine whether the channel is declared to be idle. Thus, it is possible that the channel is declared to be busy for one frame, while there is another frame in the buffer, with a different destination, for which the channel is not idle. Thus, we allow packet reordering. Specifically, if head-of-the-line blocking is found to be occurring, then, if possible, packets are reordered so that the channel is declared to be idle. This MAC feature will be referred to as *packet reordering*.

3) Capture of Stronger Signals: Radio implementations such as Atheros 802.11 chip-sets allow the radio to decode a packet even when another transmission is ongoing. For example, if a signal is being decoded and another stronger signal begins, then the radio will stop decoding the first frame and attempt to decode the new one. This decoding will likely be successful if the SINR of this second frame is sufficiently high. The ability to decode stronger signals is important when stomp is used, because a node is likely to transmit when other nodes are transmitting. The QualNet simulator was modified to allow capture of stronger signals.

V. 802.11-STYLE MAC ALGORITHMS WITH EXTRA CHANNEL USAGE INFORMATION

One drawback of the schemes discussed above is that frames need to be decoded in order to gain channel usage information. Hence, it is unclear whether the throughput achieved by these schemes is suboptimal because of the inability to decode frames due to collisions (e.g., a node being unable to decode a CTS frame) and learn channel usage information, or whether knowledge of *local* channel usage (such as the channel usage information transmitted to nearby nodes within MAC headers) is sufficient to determine whether the channel is idle. In this section, MAC algorithms are discussed that are based on different types of channel usage information, but the way in which this information is learned is not addressed. Of course, in practice, the overhead of gaining this information might dramatically limit the performance of the MAC algorithms discussed here. Note that the MACs in this section do not use RTS/CTS and, unless mentioned otherwise, stomp, packet reordering, and capturing stronger signals are used.

A. Global Channel Usage Information

Global channel usage information implies that a node can precisely determine whether a transmission will succeed and whether a transmission will cause a collision with *any* ongoing transmission. Note that ACKs must also be considered. Here we assume that a channel is idle if a transmission will succeed and if the transmission will not interfere with any ongoing transmission. In this case, the MAC behaves as a standard 802.11 MAC with an idle channel (i.e., wait DIFS, decrement back-off timer, or begin transmitting).

B. Regional Channel Usage Information

Global channel usage information might not be required to avoid most collisions. Thus, it might be sufficient to only use regional information. The exact region over which the channel usage is known will impact the capacity. Thus, we explore the throughput for different types of regions. Note that in this case collisions may occur. If a collision does occur, the MAC behaves exactly the same as 802.11.

1) Transmitter and Receiver Regional Channel Usage Information: Transmitter and Receiver Regional Channel Usage Information (TRCUI) assumes that the channel usage around both the transmitter and receiver is known. The region around the transmitter and receiver is defined as follows. Suppose that G_R is the channel gain required to decode a 1500B frame with probability 0.99 that is transmitted at bit-rate R. We assume that the transmitter is aware of the channel usage by any node to which it has a channel gain of at least G_R or to which the receiver has a channel gain of at least G_R. Thus, the region over which channel usage information is known is parameterized by the bit-rate R. Note that in this case, this channel usage information includes the information that a node within this region is receiving a frame, as well as that a node within this region is transmitting a frame.

This scheme is an idealized version of 802.11 with RTS/CTS. Specifically, in 802.11 with RTS/CTS, when a node transmits a CTS, it alerts all nearby nodes that it is receiving a transmission, where "nearby" is parameterized by the bit-rate of the CTS. Similarly, when a node transmits a RTS, it alerts all nearby nodes that it is transmitting. Hence, nodes are aware of nearby transmissions and receptions. Moreover, when an RTS is received, a node will only reply with a CTS if the channel is idle, in the sense that it has not received a RTS or CTS from a nearby node. Thus, when 802.11 with RTS/CTS is used, a node will only transmit if there are no nearby transmissions or receptions and if there are no transmissions or receptions near to the intended receiver.

Note that TRCUI allows the transmitter to have perfect knowledge of receptions near to the intended receiver, thus this scheme allows the transmitter to have knowledge of transmission originating three hops away.

2) Transmitter Regional Channel Usage Information: Transmitter Regional Channel Usage Information (TCUI) implies that the only transmitter is aware of each transmission or reception by a node where the channel gain to the node is at least G_R. Like the transmitter and receiver regional information case, this algorithm is also parametrized by the bit-rate R.

3) Transmitter-Transmitter Regional Channel Usage Information: In the case of Transmitter Regional Channel Usage Information, the transmitter is aware of any nearby node that is transmitting or receiving. Transmitter-Transmitter Regional Channel Usage Information (TTCUI) assumes that a node is aware of only transmissions originating in a nearby region, where, as above, the region is parameterized by the bit-rate R. This channel information includes both the source and destination of the nearby transmissions. Thus, a node can make an accurate determination of whether its transmission will collide with a transmission originating nearby. On the other hand, this scheme will not be able to determine if a node's transmission will collide with a transmission originating from outside of the region. In a sense, this scheme is an idealized version of 802.11 without RTS/CTS.

VI. SIMULATION SET-UP

This investigation employed the realistic urban mesh networks constructed by the UDel Models [1]. The UDel Models propagation simulator is based on ray-tracing and accounts for reflections off of the ground and off of buildings, transmission through building walls, and diffraction around and over buildings. It also accounts for the impact that different materials have on reflections off of walls and transmission through walls. Data sets for several urban areas are available online. For this investigation, a large number of topologies were considered. Each topology was based on a different 6×6 city block region that was randomly located within a 2 km^2 region of downtown Chicago. Various node densities were investigated. Specifically, the number of gateways was 1, 3, and 5, and the number of wireless routers was 18, 36, 54, 72, and 90. The wireless routers and gateways were uniformly distributed throughout the 6×6 city block region. Ten samples of each topology were generated (hence, 150 topologies in total). Since the nodes are all outside, the channels are quite strong, and hence high bit-rate communication is possible.

The simulations used a modified version of QualNet with 802.11a and transmission power of 18 dBm. Furthermore, it was assumed that the relationship between SNR, bit-rate, and bit error probability is the same as the relationship between SINR, bit-rate, and bit error probability.

VII. SUMMARY OF MAC ALGORITHMS

The MAC algorithms examined are listed below. As mentioned above, except for optimal STDMA, all algorithms are based on 802.11. Algorithms G-J use extra channel usage information, however, there is no mechanism for exchanging this information (i.e., the information is provided by a genie). All the algorithms G-J use capture, stomp, and packet reordering. For the algorithms A-F, capture, stomp, and packet reordering are only used as indicated.

Modified Versions of 802.11

A Without RTS/CTS (Baseline)
B With RTS/CTS transmitted at 6Mbps
C With CTS-to-self transmitted at 6Mbps
D Aloha, i.e., the channel is idle unless the node is transmitting a frame or receiving a frame where it is the destination of the frame
E Without RTS/CTS and capture of stronger signals
F Without RTS/CTS and with capturing stronger signal, stomping signals, and packet reordering

Algorithms with Extra Information

G TTCUI with $R = 6$Mbps
H TCUI with $R = 6$Mbps
I TRCUI with $R = 6$Mbps
J Global knowledge
K Optimal STDMA

VIII. COMPARISON OF MAC ALGORITHMS

Figure 2 shows the throughput for 11 MAC types and 150 topologies. The throughputs are normalized by the throughput achieved with MAC type A.

Consider MAC types B, C, and D. These MACs provide considerably lower throughput than MAC type A. The reason that 802.11 with RTS/CTS (MAC type B) achieves such low throughput is that we have found that RTS/CTS did not reduce collisions. Instead, the RTS causes a significant number of collisions, while the transmission of RTS and CTS utilizes precious bandwidth. Similarly, CTS-to-self did not reduce the number of collisions and simply utilized bandwidth. Interestingly, the throughput of MAC type B and C was found to be approximately the same as Aloha (actually, Aloha had a higher throughput, but the difference was less than 10%, which is the estimated accuracy of the throughput estimates). These results show the difficulty in disseminating channel usage information (which is the objective of RTS/CTS and CTS-to-self). Hence, the results based on MAC algorithms with extra channel usage information where the information is provided by a genie are upper bounds on performance and these bounds might be considerably higher than what could be achieved.

While not shown in Figure 2, we performed experiments where the RTS, CTS, or CTS-to-self frames were transmitted with bit-rates ranging from 6Mbps to the bit-rate used for data transmission. We also performed experiments with 802.11 without RTS/CTS but with CCA ranging from -75 dBm to -94 dBm. However, these adjustments had little impact on the throughput.

As mentioned above, MAC type G is an idealized version of 802.11 without RTS/CTS but where channel sensing is able to provide a node perfect knowledge of nearby nodes that are transmitting. There has been some reseach on trying to improve 802.11's knowledge of nearby transmissions. For example, [6] studied the impact of adjusting CCA to get a more accurate knowledge of nearby transmissions. Also,

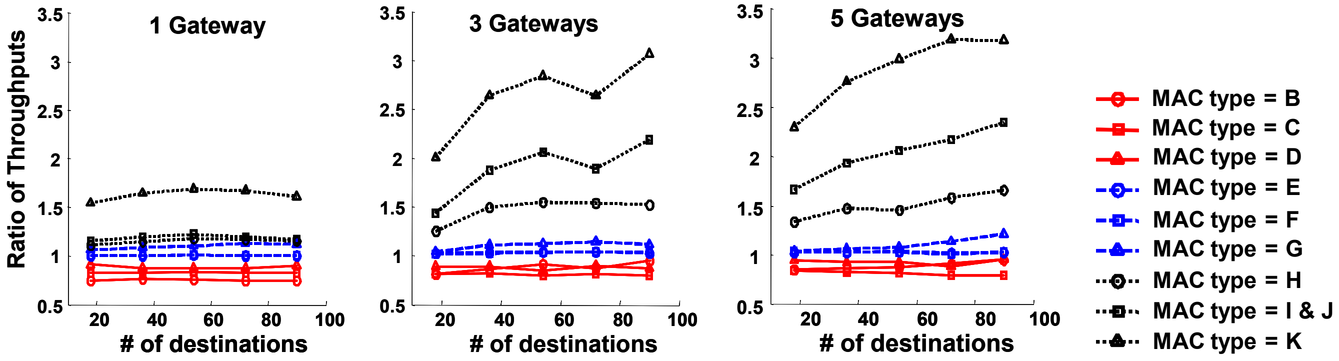

Fig. 2. Ratio of the throughput achieved by the indicated MAC and the throughput achieved by MAC type A, standard 802.11 without RTS/CTS.

Atheros provides the stomp feature for making decisions as to whether the channel is clear. However, Figure 2 shows that while such techniques might be able to increase the throughput, the improvement will be small when compared to MACs that have more detailed channel usage information and optimal STDMA.

Figure 2 shows that when there are 3 or 5 gateways, the knowledge of nearby transmitters and receivers (MAC type H) achieves significantly higher throughput than the knowledge of nearby transmitters alone (MAC type G). Note that MAC type H has prefect knowledge of transmissions and receptions in a one-hop radius. The TRAMA [5] MAC protocol provides a mechanism to exchange this type of information.

MAC types I and J achieve a considerably higher throughput than MAC type H. This indicates that the knowledge of transmissions and receptions within one-hop of the transmitter is not sufficient to eliminate collisions. Specifically, the transmission of an ACK by the receiver might cause a collision with transmissions or receptions near to the receiver. Figure 2 also shows that MAC types I and J achieve the same throughput, which indicates that knowledge of transmissions and receptions near to the transmitter and intended receiver is sufficient.

We should emphasize that the size of the region is critical. For example, in Figure 2, the region was defined by a bit-rate of 6Mbps. On the other hand, if the region is defined by the data-rate given by (2) (which is the data-rate used for data transmissions), then the throughput achieved is nearly the same as that achieved by 802.11 without RTS/CTS. Increasing the region by setting $R = 24$Mbps results in a throughput that is roughly halfway between the one achieved by MAC type H and type A. On the other hand, further increasing the region with $R = 1$Mbps does not further improve the throughput.

Finally, we compare MAC type J to optimal STDMA. Clearly, optimal STDMA achieves significantly higher throughput than global information. Since the dissemination of global information did not require any overhead and was assumed to be error free, the throughput achieved by MAC type J was considerably larger than what could be achieved in practice. As mentioned in the Introduction, this result implies that optimal STDMA not only eliminates collisions, but also orders packets in an optimal way. While collisions have been well studied, there has been little work on ordering packets. Nonetheless, for mesh networks studied here, packet ordering is responsible for at least half of the difference in the throughput between standard 802.11 and optimal STDMA.

IX. CONCLUSIONS

By examining a number of hypothetical MACs, we can make the following conclusions. The throughput gains achievable by adjusting carrier sensing (e.g., adjusting CCA or using Atheros's stomp feature), so that a node has a better knowledge of nearby transmitters will not greatly improve throughput. In fact, even if a node has perfect knowledge of nearby transmitters, the throughput is approximately the same as 802.11 without RTS/CTS. Significant improvement in throughput can be achieved with perfect knowledge of nearby transmitters and receivers, and higher throughput can be achieved if a node is aware of all transmissions and receptions near to itself and the node to which the node would like to transmit. However, even with the highly detailed knowledge, the throughput is considerably less than that achieved by optimal STDMA. Thus, even if all collisions are eliminated, the throughput of greedy 802.11-like MAC protocols will achieve a throughput that is considerably less than optimal STDMA. Specifically, we found that in multihop mesh networks, eliminating collisions is able to achieve a throughput that is approximately halfway between standard 802.11 and optimal STDMA.

REFERENCES

[1] S. Bohacek, V. Sridhara, and J. Kim, "UDel Models," available at: http://udelmodels.eecis.udel.edu/.

[2] S. Bohacek and P. Wang, "Toward tractable computation of the capacity multihop wireless networks," in *INFOCOM*, 2007, available at: http://udelmodels.eecis.udel.edu/publications1.php.

[3] IEEE Std 802.11, *Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications*. IEEE Computer Society, 1999.

[4] A. Kamerman and L. Monteban, "WaveLAN II: A high-performance wireless LAN for the unlicensed band," *Bell Labs Technical Journal*, pp. 118–133, 1997.

[5] V. Rajendran, K. Obraczka, and J. Garcia-Luna-Aceves, "Energy-efficient, collision-free medium access control for wireless sensor networks," in *SensSys*, 2003.

[6] J. Zhu, X. Guo, L. Yang, W. S. Conner, S. Roy, and M. Hazra, "Adapting physical carrier sensing to maximize spatial reuse in 802.11 mesh networks," *Journal of Wireless Communications and Mobile Computing*, vol. 4, pp. 933–946, 2004.