Exploring a New Approach to Collision Avoidance in Wireless Ad Hoc Networks

Jun Peng¹ and Liang Cheng²

¹ Department of Electrical Engineering, University of Texas - Pan American,
Edinburg, TX 78541 USA

pengjun@ieee.org

² Department of Computer Science and Engineering, Lehigh University,
Bethlehem, PA 18015 USA

cheng@cse.lehigh.edu

Abstract. ¹We propose in this paper a new approach of *bit-free* control frames to improving the performance of the IEEE 802.11 DCF. Basically, a *bit-free* control frame does not contain any meaningful bits; instead, its length (i.e., airtime) is encoded with the control information. This new approach has two advantages over the traditional control frames. First, the airtime of a bit-free frame is easy to detect and robust against channel effects. Second, *bit-free* control frames can be very short because no headers or preambles are needed for them. Our investigation demonstrates that the new approach improves the performance of the IEEE 802.11 DCF significantly (network throughput gains from fifteen to more than one hundred percent).

1 Introduction

The hidden terminal phenomenon in wireless packet networks is interesting but problematic. Basically, even if two nodes in a wireless network cannot sense each other, they may still cause collisions at the receiver of each other [1]. If the hidden terminal problem is not well addressed, a wireless network may have a significantly degraded performance in every aspect, since frequent packet collisions consume all types of network resources such as energy, bandwidth, and computing power but generate no useful output.

There are basically two existing approaches to the hidden terminal problem. One is the use of an out-of-band control channel for signaling a busy data channel when a packet is in the air [2,4,3]. This approach is effective in dealing with hidden terminals but requires an additional control channel. The more popular approach to the hidden terminal problem is the use of in-band control frames for reserving the medium before a packet is transmitted [5,7,6]. The popular IEEE 802.11 standard [9] uses this approach in its DCF.

¹ The research was partly supported by the Commonwealth of Pennsylvania, Department of Community and Economic Development, through the Pennsylvania Infrastructure Technology Alliance (PITA), and Lehigh University.

Basically, before an IEEE 802.11 node in the DCF mode transmits a packet to another node, it first sends out a Request to Send (RTS) frame after proper back-offs and deferrals. After receiving the RTS frame, the intended receiver responds with a Clear to Send (CTS) frame, which includes a Duration field informing its neighbors to back off during the specified period. In an ideal case, the hidden terminals of the initiating sender will successfully receive the CTS frame and thus not initiate new transmissions when the packet is being transmitted.

However, control frames have limited effectiveness in dealing with hidden terminals because they may not be able to reach all the intended receivers due to signal attenuation, fading, or interference [8]. In addition, control frames have considerably long airtimes because they are recommended to be transmitted at the basic link rate in both narrow-band and broadband IEEE 802.11 systems for link rate compatibility among nodes. In addition, they also usually carry long physical layer preambles and headers. Therefore, in-band control frames still introduce significant network overhead, even though they do not require an out-of-band control channel.

This paper explores a new approach of *bit-free* control frames to addressing the disadvantages of the traditional control frames. Basically, with the new approach, control information is carried by the *airtimes* instead of the *bits* of control frames. The airtime of a frame is easy to detect and robust against interference and channel effects. In addition, a bit-free control frame carries no meaningful bits so that no preamble or headers are needed for it (in-band bursts of variable lengths were used for priority scheduling in [10]).

To investigate the potentials of the new approach, we have modified the IEEE 802.11 DCF by replacing the traditional control frames with bit-free control frames and have done extensive simulations with the modified protocol. Our investigation has shown that the modified protocol improves the average throughput of a wireless network by from fifteen percent to more than one hundred percent.

The rest of the paper is organized as follows. Section 2 presents our modifications to the IEEE 802.11 DCF. We show in Section 3 the comprehensive simulation results comparing the modified protocol to the original one. Finally, we give our conclusions in Section 4.

2 Applying the New Approach

2.1 Basics

The challenge in applying the new approach to the IEEE 802.11 DCF is the limited capability of the bit-free control frames in carrying control information. Particularly, only the airtime of a control frame can carry control information. To address this issue, we use two basic strategies. One is that the bit-free control frames only carry the *indispensable* information for medium access control, while the other is to use frame *pairs* for backoff duration control.

For sending bit-free control frames, we assume that the IEEE 802.11 hardware has some modification so that it can be commanded to transmit the carrier for a specified amount of time. We also assume that the airtime of a control

frame can be recorded with a degree of accuracy depending on the hardware, bandwidth, and channel conditions. One protocol parameter, the minimum guard gap between the lengths of two control frames, may be adjusted based on the recording accuracy. In fact, with its carrier sense capability, the existing IEEE 802.11 hardware may record the airtime of an incoming frame.

In addition, a bit-free control frame can not be mistaken as a bit-based frame, since a bit-free frame does not include a physical layer preamble and thus the synchronization on the frame can not be done. A bit-based frame, however, may be mistaken as a bit-free frame if the synchronization on the frame fails. This kind of interference is usually filtered out due to the typically long airtime of a bit-based frame and the short airtime of a bit-free control frame.

2.2 Bit-Free Control Frames

The frame type needs to be specified for each frame so that the receiver knows how to interpret the bits in the bit-based frame case or the frame airtime in the bit-free frame case. Bit-free frames carry no meaningful bits so that the frame type information can only be delivered by their airtimes. Particularly, if the airtime of a bit-free frame falls into a specified range or ranges, then the frame belongs to the type of frame denoted by the range or ranges.

Besides the frame type information, the other indispensable information in an RTS frame is the address of the receiver. The length of a bit-free RTS frame needs to fall into the designated range or ranges. We therefore may not be able to encode the address information of each single receiver into the airtime of a bit-free RTS frame. To address this problem, we apply a “Mod- n ” calculation on each receiver address before it is encoded. Basically, we first divide the address by n and then encode the remainder into the frame airtime. Particularly,

$$\text{If } r = \text{Mod}(RA, n), \text{ then } F_L = \text{RTS}(r)$$

where RA is the receiver address, n is an integer, r is the remainder, F_L is the airtime of the bit-free RTS frame to send, and $\text{RTS}(r)$ is an r -indexed element in the set of RTS lengths in *microseconds*.

The Duration field in a bit-based RTS frame is also important because it specifies the period during which a receiver of the frame should back off. A bit-free RTS frame does not have the capacity for the duration information. Instead, a receiver of a bit-free RTS frame starts to back off upon receiving the frame and ends the backoff only after the medium has been sensed idle for a specified amount of time (more details later).

In our proposed design with bit-free frames, all CTS frames have the same fixed length that distinguishes them from other bit-free frames. In addition, we use control frame pairs to communicate the backoff duration information of a traditional CTS frame, which will be introduced later. Similarly, all bit-free ACK frames in our design have the same fixed length that distinguishes them from other types of bit-free frames (the address issue of these frames is discussed in Section 2.5).

In addition to the RTS, CTS, and ACK bit-free frames, we add another type of bit-free control frame named CTS-Fail frame in our design. A CTS-Fail frame has a fixed length and is sent by a *CTS frame sender* in two cases to notify other nodes to end their backoff. The first case is that a CTS frame sender does not receive any packet after SIFS (Short Interframe Space) plus propagation delays after sending the CTS frame. The second case is that a CTS frame sender receives a packet after sending the CTS frame but finds that either the packet is not intended for it or the packet has errors. In this case, the CTS-Fail frame is sent only after the packet is fully received.

2.3 Frames Working Together

To explain how the four types of bit-free control frames work together in the modified IEEE 802.11 DCF, we describe how a node contends for the medium when it has a packet to transmit. The IEEE 802.11 DCF is basically a CSMA/CA protocol, and our modifications to the protocol are only on the CA part.

When a node has a packet to transmit, it starts to listen to the channel. If the channel has been found idle for a period of time longer than the DCF Interframe Space (DIFS), the node starts a random backoff timer whose value is uniformly drawn from the node's contention window (CW). If the node detects no carrier before its backoff timer expires, it proceeds to transmit an RTS frame upon the expiration of its backoff timer. Otherwise, the node backs off.

As soon as the backoff timer of the node expires, the node starts to transmit a bit-free RTS frame. As explained earlier, the airtime of the bit-free frame is determined by the address of the intended receiver. After finishing the transmission, the node waits for a CTS frame, whose airtime is fixed and known.

After a neighbor of the initiating sender receives the bit-free RTS frame, it does the "Mod- n " calculation on its own address and compares the remainder to the length of the received frame in microseconds. If the remainder matches the length, the neighbor sends out a bit-free CTS frame and then waits for a packet. If the CTS frame sender does not receive any packet after a period of SIFS plus propagation delays, it sends out a CTS-Fail frame. On the other hand, if the remainder does not match the length of the received RTS frame, the neighbor will enter backoff and remain in the backoff until the medium has been sensed idle for a period of time that is SIFS plus either the CTS frame length or the ACK frame length, whichever is longer.

After the initiating sender obtains the bit-free CTS frame, it waits for SIFS and then starts to transmit the packet. If for any reason the RTS frame sender fails to obtain the expected CTS frame, the sender starts over to contend for the medium. In such a case, the sender doubles its CW. On the other hand, if a node receives an unexpected bit-free CTS frame (i.e., the node is not an RTS frame sender), the node increases its CTS frame counter Num_{cts} by one, starts a backoff monitor timer, and then enters backoff. Such a node exits the backoff in two cases. One is that its CTS frame counter Num_{cts} reaches zero when the node decrements the counter by one after receiving an ACK or CTS-Fail frame, while the other is that its backoff monitor timer expires (more details later).

After the initiating sender succeeds in contending for the medium, receives the expected CTS frame, and fully transmits the packet, it expects a bit-free ACK frame from the receiver. If the sender does not obtain the expected acknowledgment, it doubles its CW and starts to monitor the channel again for a retransmission.

On the other hand, after a node receives the data packet, it checks if the packet is intended for it and free of error. If so, the node sends back a bit-free ACK frame. If the packet is not intended for it or the packet has errors, the node checks whether it has sent a CTS frame for the packet. If so, the node sends out a CTS-Fail frame to notify its neighbors to exit backoff.

The whole process repeats until the initiating sender obtains an acknowledgment for the packet or the retry limit is reached. The node discards the packet in the latter case and resets its CW to the minimum size in both cases.

2.4 Some Design Considerations

The first design consideration on the modified MAC protocol is the choices of receive power thresholds for its bit-free control frames. Unlike bit-based frames, bit-free control frames can be correctly received as long as they can be sensed. The receive power threshold for a bit-free control frame may thus be adjusted for controlling the transmission range of the frame. As introduced earlier, a bit-based CTS frame may not successfully reach all the hidden terminals of the initiating sender [8]. A node with the modified protocol, therefore, needs a lower receive power threshold for bit-free control frames.

The lowest power threshold that a node may use for receiving a bit-free control frame is the carrier sense power threshold. In such a case, a node decodes a bit-free frame if the frame can be sensed. The implementation in our simulations uses this conservative choice to ensure the coverage of bit-free control frames. However, there is an exception. When a node receives a bit-free RTS frame matching its address, the node responds with a CTS frame only if the received power of the RTS frame is above the receive power threshold for data frames, since the node should not respond if it can not correctly receive a packet from the other node.

Another design consideration on bit-free control frames is the set of lengths in terms of airtimes that the frames should use. The basic rule is that control frames should be easy to detect and distinguish from one another. The shortest control frame in our simulations is $20\text{-}\mu\text{s}$ long and the minimum guard gap between two lengths in the set is $5\mu\text{s}$, which corresponds to 5-bit airtime at the transmission rate of 1Mb/s (even in broadband systems such as 802.11g, the control frames are recommended to be transmitted at a basic link rate). In reality, the minimum guard gap should be set based on the length detection accuracy of bit-free frames, which may be affected by the hardware, bandwidth, and channel conditions.

When choosing the length for a specific control frame that has a fixed length, we need to consider another factor. In particular, when multiple bit-free frames arrive at the same node in the same time segment, they may form a “merged” bit-free frame that has a length denoting another defined bit-free control frame.

Fig. 1. Merging of Control Frames

This kind of false control frame may appear when the merged frame has a longer airtime than any individual merging frame, as demonstrated by Case 3 in Fig. 1.

The possible adverse effects of the merged frame phenomenon are alleviated by the discrete lengths of the defined control frames and the strict timelines for receiving CTS and ACK frames. Particularly, only when a merged frame matches a defined bit-free control frame, would it possibly cause some harm. Moreover, for control frames such as CTS and ACK, a false frame may be harmful only if it emerges in the right timeline and at the right node (if a false ACK appears accidentally at a sender, the lost packet will be recovered by upper layers).

However, we may still further address the merged frame phenomenon by carefully choosing the lengths for the fixed-length control frames. We have three types of fixed-length control frames, which are CTS, ACK, and CTS-Fail. Among them, a false CTS frame would arguably generate the worst scenario, in which the nodes receiving the false frame enter backoff and wait for a non-existing ACK or CTS-Fail frame for exiting the backoff. Therefore, to avoid false CTS frames generated by merging frames, we need to assign a CTS frame the shortest length in the chosen length set for control frames.

What happens if a false CTS frame emerges anyway due to a reason such as environmental noise? A backoff monitor timer is used to address this problem. When a node receives a CTS frame, it starts a backoff monitor timer before it enters backoff. The backoff monitor timer is set to a value T_m that is the transmission time of the largest allowable frame in the network. The node exits the backoff anyway when its backoff monitor timer expires. Additionally, a backoff monitor timer also solves the problem of lost ACK or CTS-Fail frames due to interference or failed nodes.

In addition, it needs some extra caution to receive a CTS frame. A RTS frame may be interpreted by two or more nodes as being intended for them due to the "Mod-n" calculation design and thus two or more bit-free CTS frames may be generated for a single RTS frame. The consequence in such a case is that the received CTS frame may be slightly longer than usual because of the various propagation delays between the RTS frame sender and its receivers (besides, the medium may be reserved in a larger space than necessary in such a case). A degree of tolerance on length variation is therefore needed for decoding a CTS frame. Particularly, if we denote the transmission distance of a node by d_{tx} and

the signal propagation speed by c , then the decoding tolerance δ on the length of a CTS frame should be the maximum possible difference of round trip times for the receivers, i.e.,

$$\delta = 2 \times \frac{d_{tx}}{c}. \quad (1)$$

Finally, a bit-free ACK frame needs to have a longer length than a bit-free CTS-Fail frame. A data frame may have more than one active receivers because the length of an RTS frame is determined by a modulus calculation. In such a case, the false active receivers will respond with CTS-Fail frames after finishing receiving the data frame. To enable the sender to still recognize the ACK frame, the ACK frame must be longer than the CTS-Fail frames.

2.5 More Design Issues

One disadvantage of bit-free control frames is that they carry no specific addresses so that they may be interpreted by any receiver as legitimate. One basic observation, however, is that when an initiating sender is expecting a CTS or ACK frame, it has already notified its neighbors except the intended receiver to back off. Therefore, an initiating sender may only receive a CTS or ACK frame from the intended receiver in a general case. Moreover, an initiating sender sets a strict timeline for receiving a CTS or ACK frame. Therefore, an initiating sender can hardly receive a false and harmful CTS or ACK frame, which makes the lack of address information in the CTS and ACK frames almost harmless.

There is a special case to consider, which is that two senders may start to transmit their RTS frames almost at the same time. If the two nodes can hear each other, there is usually no harm, since in such a case the sender with a shorter RTS frame will usually detect the other sender after it finishes its RTS frame transmission. If the two senders cannot hear each other, there may exist a harmful situation in which one sender overhears the CTS frame intended for the other and mistakenly starts to transmit its packet. This kind of harmful situation occurs, however, with low probabilities because two senders with different RTS frames have different timelines for receiving their CTS frames. If the two RTS frames have the same length, a collision may occur and the following random backoffs of nodes will resolve the issue.

3 Scheme Evaluations

We have done extensive simulations with ns-2 [11] to investigate the performance of the modified IEEE 802.11 DCF and compare it to the original protocol. As mentioned earlier, we only modified the collision avoidance (CA) part of the original protocol, while other parts of the original protocol were kept unchanged. For easy reference, we named the modified MAC protocol as CSMA/FP, which denotes Carrier Sense Multiple Access with Frame Pulses.

Fig. 2. Network Throughput vs. Number of Nodes in the Network

Fig. 3. Average CW Size vs. Number of Nodes in the Network

Fig. 4. Average Medium Access Delay vs. Number of Nodes in the Network

3.1 Configuration Details

We first evaluated CSMA/FP in a wireless LAN with saturation traffic and compared it to the original protocol. We then used a more general scenario of a multi-hop ad hoc network to investigate its performance. Particularly, we evaluated the protocols from the perspective of an individual user in the ad hoc network.

From an individual user's perspective, the network is better if the user can have statistically higher flow throughput. Although a contention-based MAC protocol may not be always fair to contending nodes in terms of one-hop throughput, the statistical rate of a random flow in the network truthfully reflects the throughput of the network, especially when the transport layer does not apply rate control over the flows in the network, as configured in our simulations.

The ad hoc network has 100 nodes in an area of 1000 by 1000 square meters. Each node uses a transmission power of 0.2 watt, which means a carrier sense range of about 500 meters with the default power threshold settings of ns-2. The link rate of each node is 1Mb/s (a higher rate means that more bits will be transmitted in the time saved by CSMA/FP for using more effective and efficient control frames). In addition, there are a maximum of 25 Constant Bit Rate (CBR) background flows. The routing protocol used in the simulations is the Dynamic Source Routing protocol (DSR) [12].

In modifying the IEEE 802.11 DCF with the bit-free control frame approach, we used an n of 20 in the “Mod- n ” calculation over the receiver's address for obtaining the length of an RTS frame. Twenty is the average number of nodes that fall into the transmission range of a node in the ad hoc network (however, we also investigated the impact of a halved n).

The elements in the length set designated for RTS frames fall into two ranges for balancing the average length of an RTS frame with the average length of other control frames. One of the ranges is from 40 to 90 μ s, while the other is from 120 to 170 μ s (with a guard gap of 5 μ s). In addition, a CTS frame, a CTS-Fail frame, and an ACK frame have fixed lengths of 20, 100, and 110 μ s, respectively.

For other parameters, the modified protocol shares the default ns-2 configurations with the original protocol. For example, the minimum and maximum sizes of the CW of a node are 32 and 1024 timeslots, respectively, while a timeslot is 20 μ s.

Fig. 5. Flow Throughput in Percentage, Max. Node Speed 0m/s

Fig. 6. Flow Throughput in Percentage, Max. Node Speed 5m/s

Fig. 7. Flow Throughput in Percentage, Max. Node Speed 10m/s

3.2 Wireless LANs

Fig. 2 shows the wireless LAN throughput versus the number of nodes in the LAN. In the simulations, every node always has packets to send (i.e., saturation traffic) and the destination of each packet is randomly selected (all nodes are in the transmission range of each other). In addition, each packet is 512-byte long. As shown in Fig. 2, the modified protocol has a relative throughput gain of about 15% (an absolute gain of about 100kb/s) when there are 5 nodes in the network. As the number of nodes in the network increases, the throughput gain of the modified protocol increases too. When the number of nodes in the network reaches 25, the relative gain increases to 25% (an absolute gain of 150kb/s).

The average maximum CW size and the average medium access delay for a packet in the network are shown in Fig. 3 and Fig. 4, respectively. As shown in the two figures, a packet experiences less delay when the modified MAC protocol replaces the original one in the network. These results conform to the throughput results shown earlier. For conciseness, we only show the throughput results for ad hoc networks in the following sections.

3.3 Ad Hoc Networks

We first tested the two protocols in the ad hoc network with stationary nodes. In particular, we run a series of simulations in which the rate of the background flows varied from 0.5×512 bytes/second (B/s) to 8×512 B/s with an increase factor of 100%. A test flow, meanwhile, kept its rate *constant* at 4×512 B/s to monitor the actual throughput that it could obtain in various cases of network load. Fig. 5 shows the percentage of the packets in the test flow that are successfully received by the flow receiver as the network load varies.

As shown in Fig. 5, when the network load increases to some degree, more packets of the test flow are delivered by the network if the modified MAC protocol replaces the original IEEE 802.11 DCF. For example, when the rate of the background flows is 4×512 B/s, the throughput of the test flow increases from about 44% to 61% as the modified MAC protocol replaces the original one, which indicates a relative performance gain of about 39%. A similar relative

performance gain is observed for the modified protocol when the rate of the background flows is 8×512 B/s.

Fig. 6 shows the throughput of the test flow when the nodes in the network have random waypoint movement and have a minimum and maximum speed of 1.0 and 5.0m/s, respectively (the average pause time is 0.5 second). Fig. 7 shows the throughput of the test flow in a similar case but the maximum node speed increases to 10.0m/s. As shown in these two figures, the modified protocol, on average, has a relative performance gain of more than 50% in both cases of node mobility. Note that the network is more dynamic when the maximum node speed increases and a more dynamic network is more challenging for medium access control.

3.4 More Hidden Terminals

Using the case of a maximum node speed of 5.0m/s, the rest of the section further investigates the performance of the modified protocol. This section shows how the modified protocol performs when there is a higher probability of hidden terminals for a transmitter in the network. To increase the probability of hidden terminals, we increased the carrier sense (CS) power threshold of a node from less than one twentieth to half of its packet receive power threshold. The increase of the CS power threshold shrinks the carrier sense range of a node in the network.

Fig. 8 shows the throughput of the test flow when the CS power threshold has been increased in the network. By comparing this figure to Fig. 6, we find that the modified protocol has even higher performance gains as the probability of hidden terminals is increased in the network. This is expected because the modified protocol is better in handling hidden terminals than the original protocol.

3.5 Link Losses

We also investigated the impact of link frame losses on the performance of the modified protocol. The bit-free control frames of the modified protocol are robust against channel effects², but a data frame may still experience link losses. We have considered both independent and burst link frame losses.

Fig. 9 shows the results for the case of an independent link loss rate of 5%. As shown by Fig. 9 and Fig. 6, independent losses may actually increase the relative performance gains of the modified protocol over the original protocol. Similar results have been observed for other cases of link frame losses.

3.6 Environmental Noise

We also investigated the impact of environmental noise on the modified protocol. To test the impact of environmental noise, we placed a noise source at the center of the network and let it generate random-length noise signals at an average rate of 100 signals per second. Moreover, we restricted the noise signal lengths to the range from $1\mu\text{s}$ to $200\mu\text{s}$, which were the range designated for the bit-free control

² Note that a receiver is in its sender's packet transmission range while bit-free frames can be received at the CS range of the sender.

Fig. 8. Higher CS Power Threshold Case

Fig. 9. An Average Link Loss Rate of 5% Case

Fig. 10. Environmental Noise Case

frames. The simulation results for this scenario are shown in Fig. 10. As shown by the comparison of Fig. 10 to Fig. 6, the modified protocol is not more sensitive to noise than the original one. In fact, after introducing the noise source in the network, the modified protocol shows higher *relative* performance gains over the original one. Note that a noise signal may not be able to change the length of a bit-free control frame even if when it damages a bit-based frame. Meanwhile, as explained in Section 2, a noise signal must have the right length, arrive at the right node, and possibly arrive at the right time for being harmful.

3.7 Protocol Resilience

The above subsections are about how external factors may impact the performance of the modified protocol. This subsection shows how the protocol's own parameters affect its performance. We have investigated the three most important parameters of the protocol, which are the receive power thresholds for control frames, the length set for control frames, and the base n of the Mod- n calculations for obtaining RTS frame lengths.

Fig. 11 shows how the modified protocol performs when all its control frames use the same receive power threshold as data frames, which deprives the modified protocol of its advantage of better hidden terminal handling. As shown in the figure, the performance of the protocol does degrade but still maintains significant gains over the original protocol.

Fig. 12 shows the performance of the modified protocol as the average length of its control frames becomes similar to the average length of the bit-based control frames of the original protocol. As shown in this figure, the performance of the modified protocol degrades gracefully in this case.

Fig. 13 shows how the modified protocol performs as the base n of the Mod- n calculation is halved. Halving the n is similar to doubling the node density of the network in terms of investigating how the redundant CTS frames for an RTS frame may affect the protocol performance. As shown in Fig. 13, the performance of the modified protocol has a graceful degradation when the n is halved.

Fig. 11. Data Receive Power Threshold Case

Fig. 12. Long Bit-Free Control Frames Case

Fig. 13. Mod-n: n Changes from 20 to 10

4 Conclusion

We have proposed in this paper a new *bit-free* control frame method for collision avoidance in wireless packet networks. Bit-free control frames do not need headers or preambles, so they can be short. In addition, bit-free control frames are robust against channel effects due to their simplicity. We have investigated the new approach by applying it to the IEEE 802.11 DCF and conducting extensive simulations. We have tested the new approach in both wireless LANs and ad hoc networks. We have also investigated how hidden terminals, link losses, and environmental noise may impact the new approach. Additionally, we have examined how protocol parameters such as the average length, the receive power thresholds, and the size of the length set of control frames may impact the performance of the new approach. Our conclusion is that the new bit-free control frame method is able to significantly improve the performance of the IEEE 802.11 DCF.

References

1. Tobagi, F. A., Kleinrock, L.: Packet Switching in Radio Channels: Part II - The Hidden Terminal Problem in Carrier Sense Multiple Access and the Busy Tone Solution. *IEEE Transactions on Communications* **23** (1975) 1417–1433
2. Kleinrock, L., Tobagi, F. A.: Packet switching in radio channels: Part I - carrier sense multiple-access modes and their throughput-delay characteristics. *IEEE Transactions on Communications* **23** (1975) 1400–1416
3. Peng, J., Cheng, L., Sikdar, B.: A new MAC protocol for wireless packet networks. *IEEE GLOBECOM*, San Francisco, CA (2006)
4. Haas, Z. J., Deng, J.: Dual Busy Tone Multiple Access (DBTMA) - A Multiple Access Control Scheme for Ad Hoc Networks. *IEEE Transactions on Communications* **50** (2002) 975–985
5. Karn, P.: MACA - A New Channel Access Method for Packet Radio. *Proc. of the 9th ARRL Computer Networking Conference*, Ontario, Canada (1990)
6. Bharghavan, V., Demers, A., Shenker, S., Zhang, L.: MACAW: a medium access protocol for wireless LANs. *ACM SIGCOMM*, London, United Kingdom (1994)

7. Fullmer, C. L., Garcia-Luna-Aceves, J. J.: Floor acquisition multiple access (FAMA) for packet-radio networks. ACM SIGCOMM, Cambridge, Massachusetts (1995)
8. Xu, K., Gerla, M., Bae, S.: How Effective is the IEEE 802.11 RTS/CTS Handshake in Ad Hoc Networks? IEEE GLOBECOM, Taipei, Taiwan (2002)
9. IEEE 802.11 Wireless Local Area Networks <http://grouper.ieee.org/groups/802/11/>
10. Sobrinho, J. L., Krishnakumar, A. S.: Real-time traffic over the IEEE 802.11 medium access control layer. Bell Labs Technical Journal 172-187 (1996)
11. The Network Simulator - ns-2 <http://www.isi.edu/nsnam/ns/>
12. Johnson, D. B., and Maltz, D. A., Hu, Y. C.: The Dynamic Source Routing Protocol for Mobile Ad Hoc Networks (DSR). IETF Interet draft,draft-ietf-manet-dsr-10.txt (2004)