

IMPACT OF TURNAROUND TIME ON WIRELESS MAC PROTOCOLS

Eric E. Johnson*, Manikanden Balakrishnan*, and Zibin Tang*
New Mexico State University
Las Cruces, NM

ABSTRACT

Media access control (MAC) protocols for mobile ad-hoc networks are of interest in applications ranging from wireless local area networks (WLANs) to global skywave high-frequency radio networks. Much of the recent work in wireless MAC protocols has emphasized networks with short turnaround times (e.g., IEEE 802.11 WLANs). However, some wireless networks suffer lengthy turnaround times in the physical layer. This paper explores the interaction of turnaround time with various types of MAC protocols, and offers guidance for selecting an efficient wireless MAC protocol.

1. INTRODUCTION

When a channel (or a pool of channels) is shared for multiple-access communication among a group of nodes, some mechanism is required to arbitrate among competing needs to send data at any instant. This channel allocation function may be performed centrally (as in cellular or trunked land mobile radio systems) or cooperatively among the nodes contending for access. Such media access control (MAC) protocols may be further categorized as either contention-based (e.g., Ethernet) or contention-free (e.g., token passing).

Use of an inappropriate MAC protocol in a wireless network can degrade network performance significantly. One of the characteristics that can interact strongly with a MAC protocol is the link “turnaround” time, the delay between arrival of a packet at a node and the beginning of its response.

In this paper, we will analyze the packet radio physical layer to identify components of link turnaround time. This is followed by analysis of several common wireless MAC protocols to discover the effect that long or short link turnaround times have on their efficiency in allocating time on the channel. The analytical models are then checked via simulation, and some guidance is presented for selecting an appropriate MAC protocol in light of these results.

1.1 Components of Link Turnaround Time

In the analysis that follows, link turnaround time is measured from the instant that the first data bit of a packet arrives at the antenna of a node until the instant that the first data bit of a response leaves the antenna of that node. This is distinct from the propagation delay, which is the time of flight of a signal from a sending antenna to a receiving antenna. The

turnaround time consists of delays in the physical layer for channel-related functions plus processing time for the MAC protocol (Figure 1).

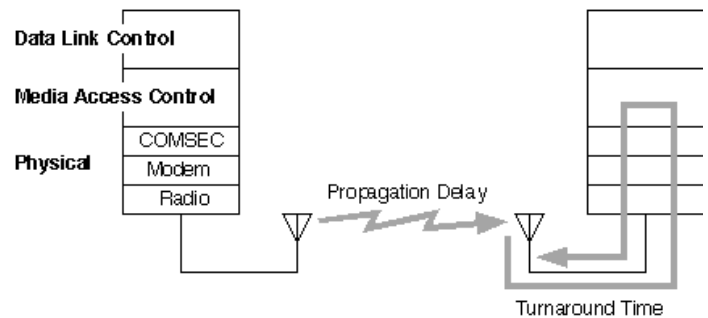


Figure 1. MAC-Layer Turnaround Time

The sublayers shown in the Physical layer represent functional groupings; they do not necessarily indicate implementation boundaries.

- The communications security (COMSEC) function, if present, provides link-by-link encryption of packets. Delays that may be added by the COMSEC function include serial communications time to send the first block of a packet to an external COMSEC device (if using a block algorithm), processing time, and time to send COMSEC preamble bits to the modem before the encrypted data bits are sent.
- The modem function provides forward error correction coding, interleaving to cope with a fading channel, and channel symbol formation. Delays in the modem functions include processing time, time to fill the first interleaver at both the sending and receiving nodes, time for a synchronization preamble at the beginning of the channel packet, and time to send the channel symbols composing the packet.
- The radio function represents modulation of the radio frequency carrier by the channel symbols as well as power amplification as necessary for each application. Settling time delay (e.g., for transmit level control) after the start of a transmission usually overlaps the modem preamble.

A simplified timing model that reflects these delays in the exchange of MAC packets is shown in Figure 2.

Published analyses of MAC protocols (e.g., [1]) usually consider only packet duration (μ), processing time (γ), and propagation delay (τ); the other delays, because they are relatively small in technologies such as IEEE 802.11 [2], have received minimal attention in previous analyses. However, COMSEC and the measures needed to overcome the chal-

* Supported by US Navy SPAWAR Systems Center, San Diego under contract N6601-97-D-5028.

allenges of some less-benign channels (e.g., HF skywave) result in significant values for η , ψ , and δ , resulting in notable impacts on the efficiency of common MAC protocols. In the next two sections, we summarize the relevant characteristics of some MAC protocols, both contention-based and contention-free. We then develop mathematical models of these protocols that accommodate the all of the turnaround time components from Figure 2 and examine their impact on performance.

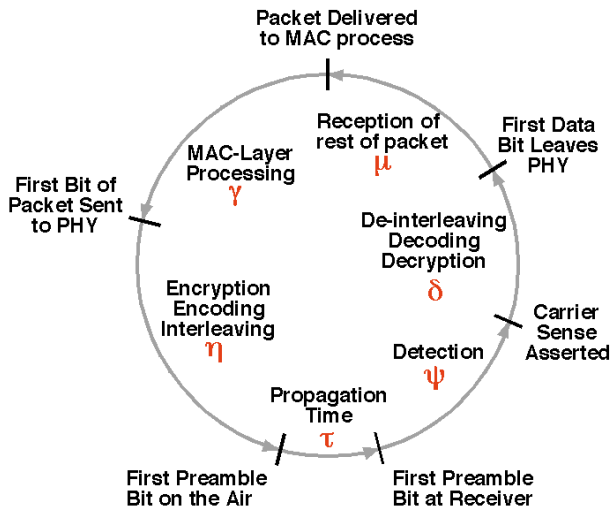


Figure 2. MAC Turnaround Time Model

1.2 Contention-Free MAC Protocols

Contention-free MAC protocols ensure that at most one node transmits at any instant. Collisions between packets do not occur during normal operation because transmission times are agreed in advance among the network nodes. We consider here two common contention-free MAC protocols: token-passing and simple time-division multiple access (TDMA).

TDMA. The simpler of the two contention-free protocols is TDMA:

- All nodes are synchronized to within some small time uncertainty ϵ .
- Each node is assigned one or more slots for transmissions out of repeating cycles of slots.

Positive aspects of simple TDMA:

- Each node is guaranteed a portion of channel time.
- No time is lost to collisions, even under heavy load.
- No overhead packets are necessary in normal operation.

Drawbacks of simple TDMA:

- Unused slots are wasted, resulting in low efficiency under light or asymmetric loading.
- A robust management protocol is required to reallocate slots to accommodate entry and departure of nodes, or to rebalance slot assignments as traffic profiles change.

- A robust synchronization mechanism is needed if an external source of precise time (e.g., GPS) is not available.

Token-Passing. An alternative contention-free protocol that addresses some of the drawbacks of TDMA is called token-passing:

- Permission to transmit (a notional “token”) is explicitly passed among the nodes.
- Upon receipt of the token, a node may transmit for a bounded time, after which it must pass the token to the next node.
- A node that has no traffic to send when it receives the token must immediately pass the token.

Positive aspects of Token Passing:

- Each node is guaranteed a portion of channel time.
- Unneeded portions of slots are not wasted, as the token is immediately passed to another node. Thus, channel capacity is dynamically allocated to nodes that need it.
- No time is lost to collisions, even under heavy load.
- The system is self-synchronizing. No external time base or synchronization protocol is needed.

Drawbacks of Token Passing:

- Channel capacity is consumed in passing the token among all nodes, even if they have no traffic to send.
- A robust management protocol is required to recover from lost or duplicate tokens.

1.3 Contention-Based MAC Protocols

We can avoid the management overhead of the contention-free protocols by employing a MAC protocol that does not guarantee collision-free operation. Of course, it is important to minimize the probability of collisions so that the channel capacity lost when multiple nodes transmit at the same time does not outweigh the management overhead avoided by permitting the occasional collision.

A common thread in current contention-based MAC protocols for wireless networks is avoiding collisions by sensing or predicting when the channel will be busy. By contrast, wired network nodes using the Ethernet protocol are able to detect collisions as they occur and cease transmission, thereby reducing the duration (and impact) of collisions.

Carrier sensing protocols call upon the physical layer to listen just before transmitting, and defer transmission if the channel is found to be busy. Several variants of such Carrier-Sense Multiple Access (CSMA) protocols are discussed in the literature (e.g., [3]). An alternative approach that does not require such close cooperation between the MAC and physical layers includes channel reservation information in MAC headers or control packets. Nodes that receive these reservations honor them by deferring their own transmissions until they believe the channel will be free. The popular IEEE 802.11 MAC protocol [2] uses this “virtual carrier sensing” in addition to physical carrier sensing.

The challenge of avoiding collisions is increased when some nodes cannot “hear” other nodes in the network. This circumstance leads to the infamous “hidden node” problem: consider three nodes, A, B, and C, situated so that node B can hear both A and C, but C and A cannot hear each others’ transmissions. Thus, if A is sending to B, node C cannot detect those transmissions, either with virtual or physical carrier sensing. If C then begins transmitting while B is receiving from A, the transmission from C will collide at B with the transmission from A, corrupting one or both transmissions.

A hidden node (C in this case) is defined with respect to a data transfer as a node that is able to interfere with the node receiving data (B), but that cannot detect the signal from the transmitting node (A). The solution to this problem is for the receiving node (B) to somehow inform the hidden node (C) that a data transfer is underway or will take place, either by sending a “busy tone” on a separate control channel or by broadcasting a channel reservation on the data channel before the transfer takes place.

The latter approach is commonly used in recent wireless MAC protocols: a node with traffic to send will first send a Request to Send (RTS) packet to the intended destination. That destination will reply with a Clear to Send (CTS) that completes the channel acquisition handshake for the announced transfer. Both the RTS and the CTS specify the duration of the data transfer to follow. Assuming that propagation among stations is symmetric, any station that could interfere with reception at the destination will hear the CTS and defer its interfering transmission(s) until the announced data transfer time has elapsed.

We here consider two contention-based wireless MAC protocols: the IEEE 802.11 Distributed Coordination Function (DCF) [2] and a variation on DCF called DCHF that uses some ideas from MACAW [4]. Each of these uses the RTS/CTS solution to the hidden node problem.

- The normal sequence of transmissions for DCF and DCHF is RTS, CTS, Data, and a MAC-layer Acknowledgement.
- The four packets in each sequence are separated by a Short Inter-Frame Space (SIFS) which amounts to a node’s internal turnaround time: $\psi + \delta + \gamma + \eta$.
- A slotted contention window for channel acquisition occurs after each ACK. The slot time allows for sending, receiving, and processing an RTS packet: $\eta + \tau + \psi + \delta + \mu + \gamma$.
- When a node is in contention for the channel, it randomly selects one of the slots in the contention window for sending its RTS, and monitors the channel during the preceding slots.
- Nodes that overhear an RTS in a preceding slot defer long enough to avoid interfering with an answering CTS.
- Nodes that overhear a CTS defer past the announced data packet and the acknowledgement that will follow.
- If a node sends an RTS but does not receive a CTS, it “backs off” before trying again: the contention window grows with each unanswered RTS (with an upper limit),

and shrinks after a successful RTS-CTS handshake. DCF and DCHF differ in their backoff algorithms.

DCF. The IEEE 802.11 Distributed Coordination Function (DCF) has the following additional characteristics:

- Physical-layer carrier sensing is performed before the contention window during an extended DCF Inter-Frame Space (DIFS).
- Binary exponential backoff is used: after a channel acquisition failure, the contention window size *at the requesting nodes* is doubled, limited by a maximum size. After successful channel acquisition, the backoff range at the requesting node is reset to its minimum size.
- When a node defers sending an RTS, its slot countdown is “frozen” while the channel is in use. It resumes counting down when the channel next becomes idle.

Positive aspects of DCF:

- A node with traffic can immediately enter contention for the channel, without waiting for a token or an assigned slot.
- Synchronization is not required.
- Nodes can join and leave the network freely, without management overhead.
- The protocol copes naturally with hidden nodes and partitioned networks.

Drawbacks of DCF:

- Nodes are not guaranteed a share of channel bandwidth.
- RTS and CTS packets consume channel capacity.
- Channel capacity is also lost to collisions, especially under heavy load.
- Nodes in the backed-off state are idle despite having traffic to send. Under heavy traffic loads, the channel can sit idle despite traffic waiting at most nodes.
- Binary exponential backoff can result in unfair channel sharing, with recent senders having the advantage due to their smaller contention windows.

DCHF. The distributed coordination for HF radio (DCHF) protocol is similar to the 802.11 DCF, but adds some features from MACAW [4] to address DCF shortcomings:

- Physical carrier sensing and the DIFS delay are not used.
- After successful channel acquisition, the backoff range is halved rather than reset to the minimum.
- Data packets carry the size of the sending node’s contention window, which is then used by all nodes.

Positive aspects of DCHF: same as DCF, plus

- Nodes are aware of localized congestion and fairly share the remaining capacity, reducing the unfairness of DCF.
- The contention window remains larger during heavy traffic, reducing collisions (compared to DCF).
- Eliminates need for reliable physical-layer carrier sensing.

Drawbacks of DCHF: like DCF (with less unfairness), but

- Lack of physical carrier sensing can increase collisions.

2. ANALYSIS OF MAC PROTOCOLS

Two canonical metrics are used to compare MAC protocols: latency under light loading and throughput under heavy loading. In this section, we develop mathematical models of the MAC protocols under consideration, and use these models to compare their latency and throughput.

- *Light-Load Latency* is measured in a situation in which a single node in the network has one packet to send. The latency is defined as the average delay from the instant that the MAC layer entity receives a packet to be sent until that data packet begins transmission. This lower-is-better metric measures the responsiveness of the MAC protocol in the absence of interference from other traffic.
- *Saturation Throughput* is measured under the opposite condition: every node in the network constantly tries to send traffic. This higher-is-better metric measures how well the MAC protocol handles heavy traffic.

For each protocol, we consider a fixed number of nodes in the network (N), with no arrivals or departures (although a background rate of management overhead is included). After acquiring the channel, a node is allowed to transmit the necessary modem preamble followed by up to T_{data} of data. Overhead bits (e.g., a COMSEC preamble, a MAC-layer acknowledgement, or a token) share this time with upper-layer (client) data. The number of overhead bits for a COMSEC preamble and a MAC-layer ack is denoted $b_{overhead}$. Bits for a token are indicated separately as b_{token} . In this analysis, the MAC layer communicates with the PHY layer via a synchronous serial interface at a rate of r bits per second.

2.1 TDMA Model

Our simple TDMA protocol allocates one fixed-size slot to each node in turn. Local timebases are assumed to vary from each other by no more than ϵ . The slot is made long enough for receiving *and processing* the packet in the current slot, so that a node can prepare and send a response in the following slot: $T_{slot} = \epsilon + T_{data} + \tau + \psi + \delta + \gamma + \eta$. The cycle time, including management overhead, is $T_{cycle} = N T_{slot} + T_{mgmt}$.

Latency. In this simple TDMA system, the average latency is simply half of the cycle time.

Throughput. When the TDMA system is saturated, every slot carries a data packet along with, on average, one acknowledgement. Throughput is the total number of user data bits sent during a cycle divided by T_{cycle} :

$$X = \frac{\bar{U}}{T_{cycle}} = \frac{N(rT_{data} - b_{overhead})}{N(\epsilon + T_{data} + \tau + \psi + \delta + \gamma + \eta) + T_{mgmt}}$$

2.2 Token-Passing Model

Latency. When the token-passing network is carrying no traffic, the token circulates at its maximum rate: each node immediately forwards the token upon receipt. The delay per node is $T_{forward} = T_{token} + \tau + \psi + \delta + \gamma + \eta$. The cycle time is

$T_{cycle} = N T_{forward} + T_{mgmt}$. Average latency is half of the cycle time.

Throughput. When the network is saturated, every node sends a full data packet along with the token and, on average, one acknowledgement. Throughput is the total number of user data bits sent during a cycle divided by T_{cycle} :

$$X = \frac{\bar{U}}{T_{cycle}} = \frac{N(rT_{data} - b_{overhead} - b_{token})}{N(T_{data} + \tau + \psi + \delta + \gamma + \eta) + T_{mgmt}}$$

2.3 DCF Model

The DCF protocol employs both physical-layer carrier sensing and MAC-layer collision avoidance. Before each contention window, nodes with traffic to send first listen for transmissions on the channel for a period called the DCF Inter-Frame Space (DIFS). If no traffic is in progress, a node selects a slot uniformly from S slots in the contention window. It then listens during slots preceding the one it selected, and defers its transmission if it receives an RTS (or CTS) in one of those preceding slots. If no other node transmits in an earlier slot, the node sends its RTS in its selected slot, and listens for a CTS. If a CTS is received, channel acquisition was successful, and the node sends its data packet. If no CTS is received, the protocol reacts by doubling S and restarting the channel acquisition process. S is S_{min} until acquisition fails. After each failure, S is doubled, up to a maximum value S_{max} . S is reset to S_{min} after a successful RTS/CTS handshake.

The basic timing parameters for DCF are as follows:

- SIFS time: $T_{SIFS} = \psi + \delta + \gamma + \eta$
- Slot time: $T_{slot} = T_{RTS} + \tau + \psi + \delta + \gamma + \eta$
- DIFS time: $T_{DIFS} = T_{SIFS} + T_{slot}$

Latency. When only one node has data to send, its RTS packets will encounter no collisions. S will be S_{min} every time, so the average waiting time to send RTS will be $S_{min}/2$. The average latency, including DIFS, waiting time, RTS, and CTS, will be $T_{DIFS} + (S_{min}/2 + 2)T_{slot}$.

Throughput. Bianchi [5] has analyzed the throughput of the DCF under saturation conditions. In this model, the probability that a particular node transmits in a particular slot is θ :

$$\theta = \frac{2(1 - 2\kappa)}{(1 - 2\kappa)(S_{min} + 1) + \kappa S_{min}(1 - (2\kappa)^\omega)}$$

where κ is the probability that a packet collides

$$\kappa = 1 - (1 - \theta)^{N-1}$$

and ω is the number of stages of exponential backoff:

$$\omega = \left\lceil \log_2 \left(\frac{S_{max}}{S_{min}} \right) \right\rceil$$

The nonlinear simultaneous equations in θ and κ can be solved numerically.

The probability that there is at least one transmission in a slot is then $P_{tr} = 1 - (1 - \theta)^N$ and the conditional probability that a that a packet sent is successful is:

$$P_{success} = \frac{N\theta(1-\theta)^{N-1}}{P_{tr}}$$

If a node is successful, it will send its data packet and (usually) receive an acknowledgement. Since the time for each overhead packet plus the SIFS equals the slot time, the total time occupied in case of success is

$$T_{success} = 3T_{slot} + T_{data} + T_{DIFS}$$

If a collision occurs, the colliding nodes will listen in the following slot and hear no CTS; other contending nodes detect the collision and make no transmissions. A new contention window begins with a DIFS after the RTS slot, so the channel time occupied for a failed attempt is

$$T_{fail} = T_{slot} + T_{DIFS}$$

Using these values, the heavy-load throughput for DCF is:

$$X = \frac{P_{tr} P_{success} (rT_{data} - b_{overhead})}{(1 - P_{tr})T_{slot} + P_{tr} P_{success} T_{success} + P_{tr} (1 - P_{success}) T_{fail}}$$

2.4 DCHF Model

DCHF, like MACAW, omits physical sensing of the channel, and is slower to reduce the contention window size than is DCF. On channel acquisition failure, the protocol doubles the value of S and tries again; on success it halves S and sends its data packet. The other features and the basic timing parameters of DCHF are the same as for DCF.

Latency. Just as for DCF, the average waiting time to send RTS will be $S_{min}/2$. The average latency, including waiting time, RTS, and CTS, (no DIFS) will be $(2 + S_{min}/2)T_{slot}$.

Throughput. With N nodes competing for the channel, the situation is more complex. The analysis that follows assumes that all nodes in contention for the channel during any contention window adjust their backoff range S in unison (due to the announcement of S values in data packets).

For each possible value for S , we need to compute the probability of that case, the expected slot number of the first slot that will contain an RTS, and the probability that only one node sends an RTS in that slot. The probability that slot i is the first occupied slot of S , when N nodes are contending for the channel is

$$a_{S,N}(i) = \left[\left(\frac{S-i+1}{S} \right)^N \right] - \left[\left(\frac{S-i}{S} \right)^N \right]$$

When N nodes are contending for S slots, we expect the first occupied slot in a contention window to be

$$A_{S,N} = \sum_{i=1}^S i \left(\left[\left(\frac{S-i+1}{S} \right)^N \right] - \left[\left(\frac{S-i}{S} \right)^N \right] \right)$$

The backoff states can be analyzed using a Markov chain (Figure 3), where a success from state S is shown as σ_S .




Figure 3. Backoff Behavior of DCHF

To compute transition rates and state probabilities, we need to know the probability of success for each value of S conditioned on N nodes in contention during a contention window. It can be shown that the probability of a successful acquisition in slot i and the overall probability of success in state S are

$$\sigma_{SIN}(i) = \left(\frac{N}{S} \right) \left(\frac{S-i}{S} \right)^{N-1} \quad \text{and} \quad \sigma_{SIN} = \sum_{i=1}^S \sigma_{SIN}(i)$$

The state probabilities p_S are computed as usual from the σ_S , and are used to compute A_N , the expected value of the first occupied slot in a contention window, and $P_{success}$, the probability of successful acquisition

$$A_N = \sum_S p_S A_{S,N} \quad P_{successN} = \sum_{i=1}^S p_S \sigma_{SIN}$$

If a node is successful, it will send its data packet and (usually) will receive an acknowledgement. Since the time for each overhead packet plus the SIFS equals the slot time, the total time occupied in case of success is

$$T_{success} = (A + 2)T_{slot} + T_{data}$$

If the first occupied slot in a contention window instead contains a collision, the colliding nodes will listen in the following slot and hear no CTS; other contending nodes detect the collision and make no transmissions. A new contention window begins with a DIFS after the slot that would have carried a CTS, so the channel time occupied for a failed attempt is

$$T_{fail} = (A + 1)T_{slot}$$

Using these values, we can compute the expected throughput:

$$X = \frac{\mathcal{U}}{\overline{B}} = \frac{P_{success}(rT_{data})}{P_{success}T_{success} + (1 - P_{success})T_{fail}} = \frac{P_{success}(rT_{data})}{(A + 1)T_{slot} + P_{success}(T_{data} + T_{slot})}$$

2.5 Analytical Results

The latency and throughput results of our protocols are plotted as functions of turnaround time in Figures 4 through 7, with the following system characteristics:

- Network sizes are 5 ($S_{max} = 16$) and 50 nodes ($S_{max} = 128$).
- Management overhead is set to $T_{mgmt} = (\sqrt{N}/10) T_{slot}$.
- The data rate $r = 6400$ bps, typical of a surface-wave LAN.
- The data packet time for each protocol is $T_{data} = 4.32$ s. This data tenure was selected as a compromise between responsiveness and efficiency in the presence of turnaround times on the order of a second, and to match the ‘‘Long’’ in-

terleaver setting of the MIL-STD-188-110B Appendix C HF radio modem.

Note the benefit to DCHF of omitting the DIFS in both large and small networks when traffic is light.

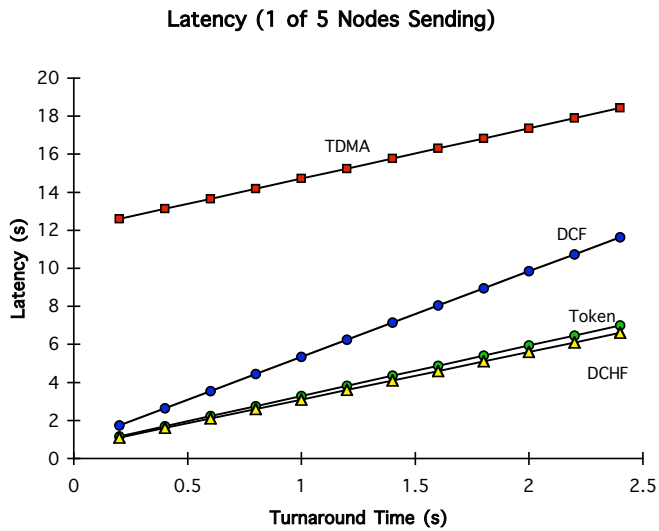


Figure 4. MAC Protocol Latency (5 Nodes)




Figure 6. MAC Protocol Throughput (5 Nodes)

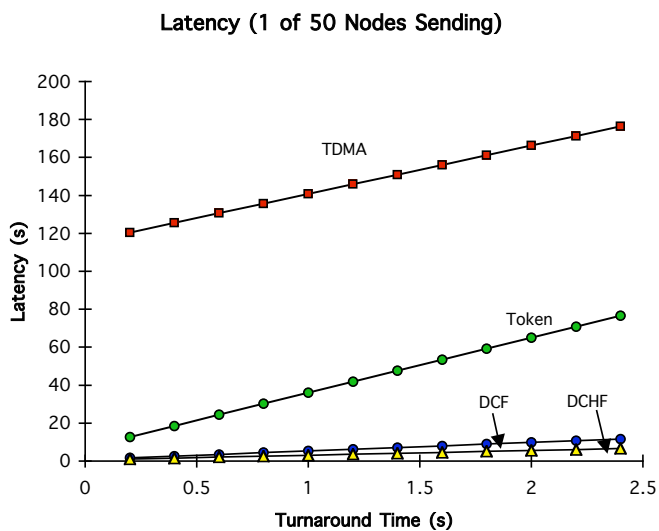


Figure 5. MAC Protocol Latency (50 Nodes)

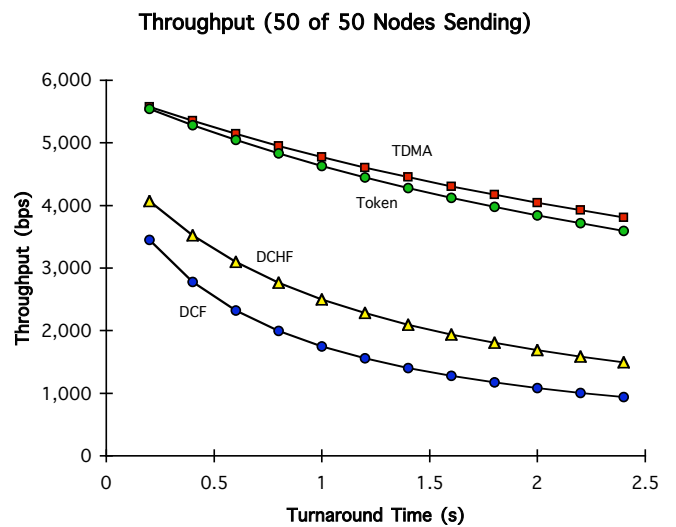


Figure 7. MAC Protocol Throughput (50 Nodes)

As expected, the contention-based protocols provide low latency in both large and small networks when the load is light. The long latencies of TDMA are likewise no surprise. However, the good performance of token passing in small networks may seem anomalous: despite the need to circulate the token among all of the nodes in the network, token passing is faster than DCF. This occurs because light-load latency is dominated by link turnarounds. Passing a token requires only a single turnaround time per node, with on average $N/2 = 2.5$ nodes before the token is received. In contrast, DCF waits through two link turnaround times during its DIFS and requires two more link turnarounds for the RTS/CTS exchange.

Under heavy load, the contention-free protocols (TDMA and token passing) provide the expected high throughput, though with some degradation as turnaround times increase. (Note that TDMA throughput would be insensitive to turnaround time if we did not require that a node be able to respond in its assigned slot to a packet received in the preceding slot.)

The contention-based protocols perform poorly under heavy load. Their saturation throughput approaches that of the contention-free protocols only when turnaround time is short, but does not appear sensitive to network size.

The ability of DCF to hold the contention window near its maximum size evidently overcomes the extra collisions that result from omitting the DIFS when traffic is heavy.

3. DISCUSSION

Contention-based MAC protocols are quite attractive compared to contention-free protocols because of their relative freedom from network management overhead. Under light loading, contention-based protocols also offer superior channel access latency (Figures 4-5). These advantages no doubt contribute to the popularity of contention-based MAC protocols in current wireless networks such as Wi-Fi (IEEE 802.11b). However, this simplicity in network management comes at a cost of reduced efficiency in using the medium, and that cost increases when large turnaround times are imposed by the MAC and PHY technology.

Heavy traffic favors MAC protocols that pre-schedule the channel. Contention-based protocols under heavy load can approach the throughput of TDMA and token passing when turnaround times are short, but longer turnaround times severely reduce their effectiveness. This is a direct result of the need for RTS/CTS handshakes to address the hidden node problem, with their additional link turnarounds.

Dynamic allocation TDMA protocols were not examined here, but their timing structure is similar to token passing, and they may be expected to offer similar performance.

4. SIMULATION OF MAC PROTOCOLS

Simulation models of the MAC protocols were developed independently of the analytical models, both to provide a check on the analytical models and for future investigations of protocol behavior in more detail than is possible using analytical models.

Simulation results for DCF throughput are compared to corresponding results from analysis in Figure 8. The results indicate good agreement: the analytical model appears to capture nearly all of the turnaround time effects exhibited by the simulation model.

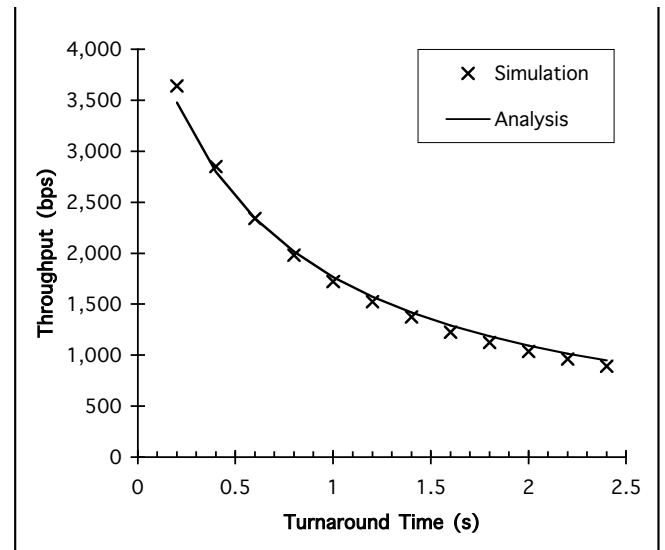


Figure 8. Simulation vs Analysis: DCF

5. CONCLUSIONS

This investigation has explored the impact of link turnaround time on the performance of various channel access protocols. Under light traffic loading, contention-based protocols such as the IEEE 802.11 (WiFi) DCF were found to offer lower latency than contention-free protocols such as TDMA.

Under heavy traffic loads, however, network throughput is severely degraded if turnaround times are long and a contention-based MAC protocol is used for channel access control. TDMA provides efficient channel access control under heavy load, but it has relatively long light-load latency. TDMA requires network synchronization as well as management intervention to assign and re-assign slots to network members, but it can be insensitive to long turnaround times.

Token passing also requires some overhead, but its performance is attractive under both light and heavy loading. The exception is large, lightly loaded networks with long turnaround times, where its performance suffers relative to contention-based protocols.

5.1 Guidance for Selecting an Efficient MAC Protocol

Factors that favor use of a contention-based MAC protocol include light traffic (most important), short turnaround times, and dynamic network membership. Contention-based protocols are also the least sensitive to network size.

When a wireless network must operate near its saturation capacity, however, a contention-free MAC protocol will provide higher throughput, especially when link turnarounds are long. In small networks, token passing excelled for both light and heavy traffic, and with fast or slow link turnarounds.

5.2 Future Work

This investigation was limited to an initial exploration of the interaction of turnaround times with MAC protocols. Additional work is planned in several areas, including MAC interactions with link-layer ARQ protocols and the effects of realistic (lossy) channels, as well as the impact of frequent changes in network membership.

REFERENCES

1. C.L. Fullmer and J.J. Garcia-Luna-Aceves, "Floor Acquisition Multiple Access for Packet-Radio Networks," *Proceedings of SIGCOMM '95*, ACM, 1995.
2. IEEE 802.11, "IEEE Standard for Information technology–Telecommunications and information exchange between systems–Local and metropolitan area networks–Specific requirements. Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications," IEEE, 1999.
3. A.S. Tanenbaum, *Computer Networks*, Prentice-Hall, 2002.
4. V. Bharghavan, *et al*, "MACAW, A Media-Access Protocol for Wireless LANs," *Proceedings of SIGCOMM '94*, ACM, 1994.
5. G. Bianchi, "Performance Analysis of the IEEE 802.11 Distributed Coordination Function," *IEEE Journal on Selected Areas in Communications* **18** (3), pp. 535-547, March 2000.