

12-2006

A cross-layer approach to increase spatial reuse and throughput for ad hoc networks

Steven Boyd

Clemson University, boyds@clemson.edu

Follow this and additional works at: http://tigerprints.clemson.edu/all_theses

 Part of the [Electrical and Computer Engineering Commons](#)

Recommended Citation

Boyd, Steven, "A cross-layer approach to increase spatial reuse and throughput for ad hoc networks" (2006). *All Theses*. Paper 7.

This Thesis is brought to you for free and open access by the Theses at TigerPrints. It has been accepted for inclusion in All Theses by an authorized administrator of TigerPrints. For more information, please contact awesole@clemson.edu.

A CROSS-LAYER APPROACH TO INCREASE SPATIAL REUSE
AND THROUGHPUT FOR AD HOC NETWORKS

A Thesis
Presented to
the Graduate School of
Clemson University

In Partial Fulfillment
of the Requirements for the Degree
Master of Science
Electrical Engineering

by
Steven Wesley Boyd
December 2006

Accepted by:
Dr. Harlan Russell, Committee Chair
Dr. Michael Pursley
Dr. Daniel Noneaker

ABSTRACT

Ad hoc networks employing adaptive-transmission protocols can alter transmission parameters over a particular link to suit the perceived channel environment. Such capabilities can lead to significant performance improvements over statically configured protocols. Ad hoc networks rely on channel-access mechanisms to govern temporal use of the transmission medium amongst the individual network nodes. Effective operation of a channel-access mechanism can improve the ability of a physical-layer adaptive-transmission protocol to accommodate changing channel conditions. Joint consideration of the interoperability of these two mechanisms motivates cross-layer design of adaptive-transmission protocols for ad hoc networks.

In this thesis we examine the integration of a new channel-access mechanism with a physical-layer adaptive-transmission protocol to create a cross-layer protocol with enhanced capabilities. We derive specific physical-layer measurements which are used by each node in the network to control channel-access behavior in a distributed manner. We examine the design tradeoffs associated with various protocol strategies and compare the performance of several protocols. Effective ways to integrate spatial reuse capabilities without introducing excessive interference into the network are investigated. We propose a distributed heuristic using cross-layer information to drive a channel-access protocol which works in conjunction with an adaptive-transmission protocol. We show that the new protocol outperforms statically configured transmission protocols as well as protocols which act independently of cross-layer enhancements.

ACKNOWLEDGEMENTS

I would like to thank my advisor, Dr. Harlan Russell, for his insight, dedication, and guidance in helping me prepare this thesis. I would also like to thank Dr. Michael Pursley and Dr. Daniel Noneaker for serving on my thesis committee.

I also greatly appreciate the opportunities MIT Lincoln Laboratory has provided for me to collaborate with them on research ideas as well as present my research findings.

TABLE OF CONTENTS

	Page
TITLE PAGE	i
ABSTRACT.....	ii
ACKNOWLEDGEMENTS.....	iii
LIST OF FIGURES	vi
CHAPTER	
1. INTRODUCTION	1
1.1 Ad hoc networks	1
1.2 Protocol design strategies	1
1.3 Brief literature review	3
1.4 Organization of thesis	5
2. SYSTEM DESCRIPTION.....	6
2.1 System model.....	6
2.2 Physical layer model.....	7
2.3 Adaptive-transmission strategy.....	7
2.4 Adaptive-transmission protocol operation.....	10
2.5 MAC-layer functions	13
3. CROSS-LAYER PROTOCOL	17
3.1 Protocol design.....	17
3.2 Transmission strategies.....	17
3.3 PDSQ statistics.....	21
3.4 MAC protocol for selective silencing.....	22
3.5 RTS acceptance criteria	26
3.6 MAC backup process.....	27
3.7 Transmitter-side validation	29

Table of Contents (Continued)

	Page
4. SIMULATION MODEL AND RESULTS	31
4.1 Protocol strategies.....	31
4.2 Simulation parameters	32
4.3 Octagonal network.....	33
4.4 Neighbor-state information.....	39
4.5 Random network topologies	51
4.6 Link utilization.....	59
5. CONCLUSIONS.....	68
APPENDIX	69
REFERENCES	75

LIST OF FIGURES

Figure	Page
2.1 Adaptive-transmission protocol operation when multiple-access interference is detected	12
2.2 MAC layer diagram	16
3.1 Network flow example.....	19
3.2 Data rate vs. D' for one link in example network	20
4.1 Octagonal network scenario 1.....	34
4.2 Throughput curves for octagonal scenario 1.....	36
4.3 Octagonal network scenario 2.....	37
4.4 Throughput curves for octagonal scenario 2.....	39
4.5 Network topology of example scenario 1	43
4.6 Performance curves for network topology 1.....	44
4.7 Network topology of example scenario 2	45
4.8 Performance curves for network scenario 2.....	46
4.9 Network topology of example scenario 3	47
4.10 Performance curves for network scenario 3.....	47
4.11 Network topology of example scenario 4	49
4.12 Performance curves for network scenario 4.....	49
4.13 Performance curves for scenario 4 with RTS broadcasting.....	50
4.14 Throughput performance for fully-connected networks of 10 nodes with $\alpha = 10$	55

List of Figures (Continued)

Figure	Page
4.15 Throughput performance for fully-connected networks of 20 nodes with $\alpha = 10$	55
4.16 Total throughput for 2000 by 2000 network topologies with $\alpha = 3$	57
4.17 Total throughput for 4000 by 4000 network topologies with $\alpha = 3$	58
4.18 Total throughput for 6000 by 6000 network topologies with $\alpha = 3$	58
4.19 Total throughput for 8000 by 8000 network topologies with $\alpha = 3$	59
4.20 Performance of 2000 by 2000 network with $Q = 20$ and $\alpha = 3$	63
4.21 Performance of 4000 by 4000 network with $Q = 20$ and $\alpha = 3$	64
4.22 Performance of 6000 by 6000 network with $Q = 20$ and $\alpha = 3$	65
4.23 Performance of 8000 by 8000 network with $Q = 20$ and $\alpha = 3$	66

CHAPTER 1

INTRODUCTION

1.1 Ad hoc networks

An ad hoc network is a self-organizing wireless network that operates without a base terminal or any other centralized network infrastructure. The nodes control network communications in a distributed manner with minimal communications overhead. Distributed mobile networks are advantageous in that they provide quickly deployable means of communications in situations where permanent networks have been disabled or are not feasible to establish. Examples of such situations include a broad range of military and disaster relief applications. Distributed protocols for ad hoc networks must be able to accommodate changing channel conditions [1]. Adaptive-transmission protocols which can sense and adapt to such changes are required to achieve reliable and efficient network communications. Protocols must provide a channel-access mechanism to allow all nodes in the network to have sufficient access to the channel without introducing unnecessary interference to the surrounding network. Interference from a neighboring node that is not the intended source of a transmission, which is referred to as *multiple-access interference*, poses a fundamental challenge to designers of ad hoc networking protocols.

1.2 Protocol design strategies

There are several strategies for dealing with interference caused by other transmissions in the network. In one strategy, similar to that used in the IEEE 802.11

protocol, the network attempts to prevent such interference by allowing only one transmission at a time [2]. However, this method restricts performance gains that occur when simultaneous transmissions can coexist. Further complications arise when the network is not fully connected and a node might not be aware of transmissions occurring just outside of the node's transmission range. This common problem is referred to as the *hidden-terminal* problem [3]. Another strategy for dealing with interference caused by other transmissions is to take advantage of the multiple-access capability of direct-sequence spread-spectrum signaling [4] by adapting the coding or spreading for a particular link to mitigate the multiple-access interference. However, this method requires additional bandwidth resources or decreases the information rate for the link. Coding and spreading adaptation is also limited by the system capabilities and cannot always overcome extreme levels of multiple-access interference. One such case is the so called *near-far* interference problem [5] in which a receiver node attempts to receive a transmission when the received power from the desired transmitting node is significantly less than the received power from an interfering transmission. Although each channel-access strategy works well in some situations, neither strategy works well in all situations.

We combine spatial reuse techniques with spreading adaptation to derive efficient cross-layer protocols for direct-sequence spread-spectrum packet radio networks. We extend previous adaptive link-layer capabilities [6] to take advantage of the ability of the medium access control (MAC) protocol to mitigate near-far interference and allow lower-layer mechanisms to be more efficient at link adaptation. Previous work [7] has shown that significant performance gains can be obtained by designing MAC protocols to work

in conjunction with adaptive-transmission protocols. We suggest use of specific physical-layer measurements to control a distributed MAC protocol. We show that using certain MAC methods in conjunction with an adaptive-transmission protocol outperforms statically configured physical layer protocols as well as adaptive physical-layer protocols that operate independently of the MAC protocol.

1.3 Brief literature review

The various strategies for mitigating multiple-access interference can be broadly categorized by power control, rate control, or transmission scheduling. Energy is a valuable resource in distributed networks and efficient expenditure of power is needed to prolong battery life. Several protocols that use power control to provide efficient network operation are given in [8]. For one such protocol it is claimed that the same minimum power level that preserves network connectivity should be used for each link in the network to maximize throughput [9]. Another power-control strategy [10] uses a two-phase, centralized algorithm for dealing with multiple-access interference. A scheduling phase is used to eliminate interference via temporal separation of significant interfering transmissions and a power-control phase then operates in a distributed manner to satisfy the interference tolerances of the remaining transmissions. A centralized method for determining optimal parameters for power, coding rate, and scheduling is given in [11] for a small number of nodes only. However, the method is not tractable for larger networks nor does it suggest distributed mechanisms to select the optimal operational parameters.

Distributed protocols rely on overhead traffic to coordinate information about the network environment and to control network operation. One distributed power-

controlled protocol [12] uses an interference margin to determine which subsequent transmissions can occur without disrupting already scheduled transmissions. This protocol, however, relies on the assumption that an additional transceiver is available at each node to correctly estimate interference tolerances to schedule future transmissions. A single-channel protocol for distributed transmission scheduling that employs session synchronization prior to each slot is described in [13]. Each synchronization period is used to determine which subset of transmissions can occur in the following session and to set the transmission powers of the nodes. Hence, spatial reuse is governed solely by transmission power restrictions. This protocol, however, becomes very difficult to implement as the number of nodes to coordinate becomes significant and the network is not fully connected. Protocols that employ spatial reuse are proposed in [14] for heterogeneous networks with both directional and omnidirectional antennas. The protocols allow reuse of multiple traffic channels based on estimates of the total interference power at the receiver and the multiple-access interference generated by the transmitter. Like the MAC protocols defined in [12-14], the protocol presented in this thesis employs an exchange of control packets to coordinate channel access. However, the decision mechanism which governs channel reuse for our protocol differs considerably from the previous protocols and is driven by different assumptions concerning the capabilities of the physical layer and the mechanisms for controlling transmission parameters.

Several protocols exist which employ adaptive-transmission protocols to react to dynamic channel conditions. The protocol investigated in [15] develops specific channel quality estimates for direct-sequence spread-spectrum systems that operate in multipath

environments to drive an energy-efficient adaptive-transmission protocol. A similar protocol for adaptive transmission in frequency-hop systems is given in [16]. The latter work uses counts of errors and erasures which are used to drive transmission parameter adaptation and route selection. The work described in [6] employs adaptive spreading at the physical layer to dynamically adapt packet transmissions to accommodate levels of interference based on channel quality estimates from previous packet transmissions. This work provides a basis for the work presented in this thesis. The work in this thesis differs in that it focuses on integration of lower-layer measurements in conjunction with operation of higher-layer protocols. We use information from the physical layer to drive a distributed MAC protocol. This protocol, in turn, allows more efficient adaptive capabilities to exist at the physical layer.

1.4 Organization of thesis

In this thesis we consider the design, implementation, and performance of a distributed, cross-layer protocol. Chapter 2 provides a description of the system model used in the design of the new protocol. Here we describe the physical-layer model as well as the basic operations of the adaptive-transmission protocol and the MAC protocol. Chapter 3 describes the operation and implementation of our cross-layer protocol. New cross-layer metrics are derived which are used to drive the distributed protocol. Specific information on the cross-layer requirements and description of protocol initialization and implementation are included. Chapter 4 defines the simulation model and the metrics used for performance evaluation. Our new protocol is compared to several other protocols in a variety of scenarios. Chapter 5 presents the conclusions of this investigation as well as interpretation of significant findings.

CHAPTER 2

SYSTEM DESCRIPTION

2.1 System model

In this thesis we examine a packet radio network employing a single transceiver with two frequency channels and adaptive-transmission capabilities at the physical layer. One of the channels is used exclusively for data packet transmissions while the other is reserved for control packet transmissions. Prior to a data packet transmission, a node wishing to transmit must coordinate via exchange of control packets with a receiver node to initiate a point-to-point transmission. It is the full responsibility of the nodes to coordinate transmissions in a distributed manner in this asynchronous environment. We employ a system model that operates when significant contention for channel access exists and operation of a particular node is highly dependent on the operation of the neighboring nodes. Our system incorporates a standard network protocol stack; however, this investigation is focused on physical and MAC-layer protocols. Implementation of a routing protocol is not considered in the scope of this study and remains as an area of future research. Transmission parameters are set according to the adaptive-transmission protocol rules described in Section 2.4. Strategies for selecting transmission parameters are discussed in Section 2.3. Channel access is the responsibility of the MAC-layer model which is described in Section 2.5.

2.2 Physical-layer model

We consider a direct-sequence spread-spectrum packet radio network. Data packets consist of L encoded binary symbols which must be transmitted over the data channel. Each symbol is spread by a spread-spectrum waveform consisting of N chips, so the data portion of the packet comprises NL chips. We refer to N as the *spreading factor* for the packet and N is restricted to be one of M values. The number of coded symbols per packet and the chip transmission rate are fixed. Increasing the spreading factor also increases the number of chips in the packet and hence increases the packet transmission time. Each packet is encoded with a rate $1/2$ convolutional code with constraint length 7 and binary phase shift keying with coherent demodulation and hard-decision decoding is used at each node.

Our channel model is characterized by propagation loss, thermal noise, and multiple-access interference. The acquisition model follows the work of [17] and [18]. As in [17], the acquisition header length and threshold are selected to make the acquisition probability much larger than the probability of successful decoding. The details of the simulation model for decoding are given in [19]. The model is based on the first-event error probability results from [5].

2.3 Adaptive-transmission strategy

The responsibility of the adaptive-transmission protocol is to adjust physical-layer transmission parameters to accommodate the channel conditions as efficiently as possible. This adaptation operates independently at each link and is not designed to make joint decisions about transmission parameters for multiple links. For the purposes of describing the adaptive-transmission protocol, we define a symbol-energy to interference

plus noise ratio (EINR) for each link. The adaptive-transmission protocol does not rely solely on the EINR estimate to control transmission parameters. However, such a definition provides insight into the operation and behavior of the adaptive-transmission protocol. The received energy per symbol E_s depends on the power level, spreading factor, and propagation loss. We establish a threshold β for which $\text{EINR} > \beta$ for a link implies satisfactory communication. The value of β depends on the target packet error rate and the coding and modulation parameters. The EINR for a packet is calculated as follows. Let $P_{i,n}$ denote the power received at node i from a transmitting node n and $N_{i,n}$ is the spreading factor used for that transmission. The average spectral density for a transmission is given by $P_{i,n}/W$ where W is the receiver bandwidth which is assumed to be equal to the inverse of the chip duration T_c . The one-sided noise power spectral density is denoted by N_0 . Different portions of a packet reception can be subjected to different levels of multiple-access interference depending on which transmissions are active. Let the packet reception time be partitioned into D intervals with each interval corresponding to a set of transmitting nodes. Let \mathbf{T}_k denote the set of nodes transmitting in interval k (i.e., $P_{i,n} > 0$ for all $n \in \mathbf{T}_k$). We define the EINR for the packet transmission from node j to node i as

$$\text{EINR} = \frac{N_{i,j} P_{i,j} T_c}{N_0 + \max \left\{ \sum_{n \in \mathbf{T}_k, n \neq j} P_{i,n} T_c : 1 \leq k \leq D \right\}}. \quad (2.1)$$

For purposes of determining packet decoding error probability, our model assumes that the largest sum of the interference powers experienced in any interval k is present for the entire packet transmission.

The adaptive-transmission protocol [6] attempts to maintain the constraint that $\text{EINR} > \beta$ for every transmitted packet. To track the signal quality of previous transmissions, postdetection signal quality (PDSQ) statistics, automatic gain control (AGC) statistics, and past link performance are employed to estimate E_s / N_0 , the ratio of the received energy per symbol to the one-sided spectral density of the thermal noise. This approximation of E_s / N_0 , which is a function of a given PDSQ statistic q and the spreading factor N , is denoted by $f(q, N)$. Our investigation employs a previously developed model for the PDSQ statistic and an associated approximation for E_s / N_0 [20-21], and the details of the model and derivation of the approximation are given in Appendix A.

Adaptive-transmission protocols are capable of adjusting many transmission parameters for a particular link including the transmission power and the spreading factor. The particular studies we consider for this investigation involve dense networks with significant demand for channel access that can result in environments with high levels of multiple-access interference. For dense networks, alteration of transmission parameters on one link has a significant impact on nearby links. An increase in transmission power on a link increases the multiple-access interference to the rest of the network. Adaptation of the spreading factor increases the link EINR without increasing the interference. The protocol investigated in this thesis adapts the spreading factor only. The adaptive-transmission protocol operates by selecting one of the M available spreading factors prior to each packet transmission. It should be noted that increasing the spreading factor will increase the transmission time of the packet but will also increase the received symbol energy E_s . In a sparse network, MAC-layer improvements, such as

the proposed mechanisms described in this thesis, have little impact on throughput performance, and the operation of the adaptive-transmission protocol follows the previously described implementation [6].

2.4 Adaptive-transmission protocol operation

A PDSQ statistic is calculated for each of n intervals of a packet where each interval contains L/n channel symbols. As described in Appendix A, the physical layer provides this vector of PDSQ statistics, denoted by \mathbf{Q} , for each packet reception. Each component of \mathbf{Q} corresponds to the data symbols in a particular portion of the data packet. Let Q_{\min} denote the smallest value in the vector \mathbf{Q} . The channel symbols corresponding to the Q_{\min} statistic are the symbols most detrimentally affected by multiple-access interference. The adaptive-transmission protocol uses Q_{\min} to form a pessimistic estimate of E_s/N_0 denoted by $f(Q_{\min}, N)$ which is compared to several thresholds to determine parameters for the next packet transmission. The adaptive-transmission protocol adopts the same rules as described in [6] with the exception that transmitter power levels are not adapted. Our investigations consider the section of the adaptive-transmission protocol in which multiple-access interference has been detected and flagged by the receiver statistics. This section of the adaptive-transmission protocol is illustrated in Figure 2.1. A new *MAC-layer backup process* is introduced to trigger the MAC protocol. The MAC-layer backup process is described in detail in Section 3.6 and does not affect the operation of the adaptive-transmission protocol.

For each packet that is acquired, the adaptive-transmission protocol determines transmission parameters for the next packet. If a packet does not decode correctly, the adaptive-transmission protocol resets the decrementing counter for consecutive

receptions, which is denoted by φ , and increases the spreading level to be used on the next transmission over the link. If such an increase is not possible because the maximum spreading level N_{\max} is already being used, then the MAC-layer backup process is triggered and no additional changes are needed. If a packet is decoded successfully, the adaptive-transmission protocol compares $f(Q_{\min}, N)$ against two thresholds. If $f(Q_{\min}, N)$ is greater than the desired minimum signal quality threshold β but less than some higher threshold σ , then the current spreading factor is ideal. If the estimate is greater than σ and has been for a specified number of consecutive receptions (counted by φ) then the adaptive-transmission protocol reduces the spreading factor over that particular link by one level, if possible, to increase efficiency.




Figure 2.1 Adaptive-transmission protocol operation when multiple-access interference is detected

2.5 MAC-layer functions

The goal of the MAC layer is to control access of nodes to a shared transmission channel. Control packets are transmitted on the control channel and used to request and allow access to the data channel. The control packets employed by our system include request-to-send (RTS), clear-to-send (CTS), acknowledgement (ACK), and negative-acknowledgement (NACK) packets. All control packets are transmitted at the same fixed power level P_i and use the maximum spreading factor N_{\max} . To initiate a transmission to a neighboring node, a source node first transmits an RTS packet to the neighboring node using a receiver-directed spreading code. The receiving node responds by transmitting a CTS packet if it successfully acquires and decodes the RTS packet and determines that the requested transmission can occur. We have developed new criteria that must be met for a node to determine that a requested transmission can occur, and the RTS acceptance criteria are described in Section 3.4. The CTS packet is transmitted using a common spreading code so that any nodes in communications range that are in the acquisition mode on the control channel can acquire and decode it. Following transmission of a CTS packet, a node switches to the data channel and attempts to acquire the data packet. If the original source of the RTS packet acquires and decodes the CTS packet from the desired destination node, then it proceeds with transmission of the data packet on the data channel. Upon acquisition of the data packet, the destination node sends an ACK or NACK packet using a receiver-directed spreading code. Both nodes return to the idle state and resume monitoring the control channel after completion of the transmission.

A *failed forwarding attempt* occurs if a node transmits an RTS but does not receive the corresponding CTS or if a node transmits the data packet but does not receive an

ACK. After a failed forwarding attempt, the node returns the packet to the queue. If the node has failed to forward the packet after ψ attempts, it discards the packet. During an RTS, CTS, data, and ACK packet exchange, a node that does not receive an expected transmission returns to the idle state and resumes monitoring of the control channel.

A random *pacing delay* [19] is imposed on any node that returns to the idle state. The node is prevented from initiating transmission of an RTS packet for the duration of the imposed delay. Such a delay helps reduce the number of collisions of control packets. While in this pacing state, a node remains in the acquisition mode on the control channel until it acquires a control packet. Once the pacing delay has expired, the node can initiate a packet transmission or remain in the acquisition mode on the control channel.

All nodes not currently transmitting attempt to acquire control packets on the control channel. Unlike most approaches to MAC layer operation, the reception of a CTS packet by a node that is not the destination does not necessarily prohibit the node from initiating its own transmission. Instead, we implement a selective node silencing strategy that triggers a silencing mechanism in nodes that are listed in the data portion of the CTS packet. Specifically, a CTS packet identifies the transmitter, receiver, and spreading factor for the data packet. Additionally, the CTS packet lists which neighbor nodes are prohibited from initiating transmissions during the subsequent data packet transmission. Any node that is able to receive the CTS packet stores this information in its network allocation vector (NAV). The number of information bits in a data packet is fixed so the specification of the spreading factor is sufficient to determine the packet transmission duration. The process for determining which subset of neighbor nodes to silence is described in Chapter 3, where we introduce a new metric for selective silencing.

If a node is silenced by a CTS packet, it sets a timer to enter a silenced state for the duration of the packet. A silenced node is prohibited from transmitting RTS packets, but it can accept RTS requests. Hence, the silenced node monitors the control channel for additional RTS or CTS packets. Following reception of a packet, a node checks its internal silencing timer to determine if a silencing restriction is in effect. If the silencing timer has expired, the node is eligible to initiate a transmission or continue monitoring the control channel. To prevent multiple silenced nodes from attempting a transmission immediately after their respective silencing timers expire, the random pacing delay is also imposed on nodes that leave the silenced state.

Figure 2.2 shows the general MAC layer state diagram. The silencing criteria verification process is described in Section 3.5 where the MAC layer implementation details are provided. The parameter selection process corresponds to the adaptive-transmission protocol evaluation of transmission parameters as described in Section 2.4 and shown in Figure 2.1.




Figure 2.2 MAC layer diagram

CHAPTER 3

CROSS-LAYER PROTOCOL

3.1 Protocol design

We propose a MAC protocol to operate in conjunction with an adaptive-transmission physical-layer protocol to improve operation of the underlying adaptation capabilities as well as improve overall network performance. Our MAC protocol takes advantage of the spatial-reuse capabilities of direct-sequence spread-spectrum modulation, but it attempts to avoid excessive multiple-access interference. The protocol described in this thesis is fully distributed and does not require excessive communications overhead. The goal of the MAC protocol is to govern channel access restrictions while allowing the physical-layer adaptive-transmission protocol to alter the transmission parameters in response to perceived changes in channel conditions. We formulate a cross-layer metric based on information from the adaptive-transmission protocol to control our distributed channel-access protocol.

3.2 Transmission strategies

The primary objective of a MAC protocol is to select the set of transmissions that should be active at any given time. The spreading factor of each link must be sufficient to mitigate the multiple-access interference caused by the other active transmissions. A centralized algorithm, such as [11], can determine the optimal sets of active transmissions and the required spreading for each link. However, the optimization problem is intractable for networks with more than a few nodes.

Permitting only adaptation of the spreading or coding to overcome excessive multiple-access interference requires no transmission scheduling. As can be seen from (2.1), a large spreading factor ($N_{i,j}$) increases the EINR for a link. However, the gain from increased spreading gives a proportional decrease in the symbol rate. Figure 3.1 illustrates a network with two competing traffic flows for which adaptive transmission can benefit greatly from transmission scheduling. It is assumed that each source always has packets to send. Node A transmits packets to node B which also receives multiple-access interference from node C . The same situation is present on the link from node C to node D where node A is the interfering node. Let $D' = d_1 / d_2$. If D' is large, the multiple-access interference is also large because the interfering node is closer to the receiver than the source transmitter. Rather than force spatial reuse, each link can achieve a higher data rate if the network activates only one link at a time. Let N_1 denote the spreading factor needed for each link to satisfy (2.1) if there is no multiple-access interference. For this network, we scale the data rate so that a value of 1.0 corresponds to a packet transmission with the spreading factor N_1 .

Figure 3.2 provides the average link data rate achievable for the network of Figure 3.1 for various transmission strategies. The *perfect-scheduling* protocol uses a spreading factor of N_1 for each link but alternates between transmissions to avoid multiple-access interference. Hence, the data rate of each link is one half that of what is possible if no multiple-access interference is present. For this example, it is assumed that the perfect-scheduling protocol is able to utilize exactly half of the transmission time for each link and that the remaining protocols continually transmit packets over the channel. Also shown in Figure 3.2 are the achievable data rates if both links are always active and the

spreading factor is fixed and the same for both links. We examine two fixed spreading schemes with spreading factors of $1.25N_1$ and $8N_1$. The data rates corresponding to the two spreading factors are 0.8 and 0.125 respectively. The curves in Figure 3.2 show the region of D' values that satisfy (2.1) for the two spreading factors. For this analytical example, if $\text{EINR} < \beta$ all transmissions fail and the achievable data rate is zero. The adaptive-spreading protocol selects the spreading factor that satisfies the EINR requirement for both of the links. Because this network is symmetric, the spreading factor is identical for both links. The data rate corresponding to this required spreading factor is plotted in Figure 3.2. For all curves in Figure 3.2, β equals 6.7 dB, d_1 is fixed such that N_1 equals 16, and the propagation loss is inversely proportional to the cubed distance.




Figure 3.1 Network flow example




Figure 3.2: Data rate vs. D' for one link in example network

As can be seen from Figure 3.2, adaptive spreading performs well if the multiple-access interference is not excessive. The perfect-scheduling protocol is unaffected by the near-far interference due to its temporal separation of link transmissions. However, the perfect-scheduling protocol fails to take advantage of channel-reuse capabilities for small D' . Using a fixed spreading of $8N_1$ provides more protection against multiple-access interference than the fixed spreading of $1.25N_1$; however, the increased spreading also reduces the data rate. Ideal operation for this network is to select the perfect-scheduling protocol if D' is greater than a specific cut-off point and, for values of D' that are less than this cut-off point, allow simultaneous transmissions to occur and adapt the spreading factors. For this particular network of two flows, derivation of the cut-off point as a function of D' that optimizes the data rate that can be achieved over both flows is

straightforward. For larger networks, creation of a multi-dimensional capacity region is infeasible. Hence, it is clear from this simple network that ideal operation of the MAC protocol is topology dependent. In order for effective channel access to occur in a distributed network, nodes must be able to derive a metric for operation based on available neighbor-state information. Description of the information required and how that information is to be used to drive the MAC protocol is given in the subsequent sections.

3.3 PDSQ statistics

A crucial component of a cross-layer protocol is the ability to obtain accurate and relevant information from other layers. Our MAC protocol silencing operation requires that each node is able to identify which neighboring nodes are significant sources of multiple-access interference. In the next section, we provide a manner to establish this information for use in the MAC protocol. These metrics are derived from E_s/N_0 estimates from the physical-layer adaptive-transmission protocol. For that reason we must ensure that the adaptive-transmission protocol is able to provide accurate estimates to the MAC protocol and that any inaccuracies do not degrade network performance.

The work in [15] first provided estimates of E_s/N_0 using receiver statistics derived from physical-layer side information as part of the adaptive-transmission protocol. The MAC protocol proposed in this thesis uses E_s/N_0 estimates to build neighbor-state information at each node to ultimately govern channel access. The nature of the MAC protocol does not require that E_s/N_0 estimates are obtained for every packet transmitted over a link. Rather, reliable estimates of E_s/N_0 from neighboring nodes can be established over several packet transmissions to form more accurate neighbor-state

information. The previous work described in [6] uses several criteria to declare if interference is present or not for a particular packet reception. If the adaptive-transmission protocol determines that interference is not present, then reliable estimates of E_s/N_0 can be obtained. However, in dense networks in which the traffic levels are high, it is likely that most transmissions will be subjected to significant levels of multiple-access interference. For this environment, we have extended the previous approach to provide a new estimate to approximate E_s/N_0 . The simulation model for the PDSQ statistics and the original estimate for E_s/N_0 were developed by Pursley [20] and Block [21], and are provided in Section A.1 of the Appendix. Our new estimate of E_s/N_0 utilizes the same model for the PDSQ and AGC statistics, and the derivation of the estimate is given in Section A.2 of the Appendix. For the simulation results reported in this thesis, we use the latter derivations when referring to E_s/N_0 estimates.

3.4 MAC protocol for selective silencing

The goal of the MAC protocol is to target the largest sources of interference for an intended receiver and prohibit those nodes from transmitting while allowing the physical-layer spreading adaptation to adjust to smaller levels of interference from the remaining nodes that are transmitting. Each receiver determines which nodes are prohibited from transmitting before it can begin receiving a packet from a particular transmitter. In particular, each receiver node i maintains a set of nodes $S_{i,j}$ for each transmitter node j such that a transmission of a data packet from node j to node i should occur only if all the nodes in the set $S_{i,j}$ are not transmitting. We refer to $S_{i,j}$ as the *silencing set* for the transmission from node j to node i . The worst possible multiple-access interference environment for a link occurs when all nodes which are not in the silencing set $S_{i,j}$ are

transmitting. The objective of the silencing set is to limit the worst possible multiple-access interference environment to a level that can be mitigated by the adaptive-transmission protocol. Although multiple-access interference levels may be substantially less than the worst possible, the silencing set restrictions ensure that the adaptive-transmission protocol is able to adapt to the level of multiple-access interference experienced.

Consider the calculation of the silencing sets at a receiver node i . Let q_j be the estimate of E_s/N_0 from node j and let $N_{i,j}$ be the spreading factor for the link associated with that estimate. We define the *scaled energy ratio* q_j^* for each node j as $q_j^* = q_j / N_{i,j}$. Values for E_s/N_0 and the corresponding q_j and q_j^* measures are given by

$$q_j \approx \frac{E_s}{N_0} = \frac{P_{i,j} N_{i,j} T_c}{N_0} \quad (3.1)$$

and

$$q_j^* = \frac{q_j}{N_{i,j}} \approx \frac{P_{i,j} T_c}{N_0} \quad (3.2)$$

and are dependent on the received power from node j , the chip duration T_c , the spreading factor $N_{i,j}$, and the one-sided spectral density of the thermal noise N_0 .

Let EIR denote the ratio of energy per symbol in the desired signal to the sum of the energies from the interfering transmissions. If a node knows values for q_n for all nodes in the network then EIR for the link from node j to node i can be calculated with

$$\text{EIR} = \frac{P_{i,j} N_{i,j} T_c}{\sum_{\forall n, n \neq j} P_{i,n} T_c} = \frac{q_j}{\sum_{\forall n, n \neq j} q_n^*}. \quad (3.3)$$

The interference summation includes all nodes n that cause some level of interference at receiver node i . An alternative representation of the EIR as defined in (2.1) is given by

$$\text{EINR} = \frac{1}{\frac{1}{E_s/N_0} + \frac{1}{\text{EIR}}} = \frac{1}{\frac{1}{q_j} + \frac{\sum_{\forall n, n \neq j} q_n^*}{q_j}} = \frac{q_j}{1 + \sum_{\forall n, n \neq j} q_n^*} > \beta. \quad (3.4)$$

From the requirement that $\text{EINR} > \beta$ during reception of a packet on a link, the maximum interference that a link can sustain for a given spreading factor can be calculated. Our heuristic for choosing nodes to silence stems from allowing as many nodes to transmit as possible. This is equivalent to restricting our silencing set to have as few nodes as possible while limiting the worst multiple-access interference to a level that can be mitigated by the adaptive-transmission protocol. The maximum spreading factor N_{\max} is used to derive the maximum interference possible such that successful transmission can still occur. Let $q_{\max} = N_{\max} * q_j^*$ which represents the value of E_s/N_0 for the link from node j to node i assuming that the maximum spreading factor is used. Replacing q_j with q_{\max} in (3.4) and solving for the interference summation, we establish the interference summation bound as

$$\sum_{\forall n, n \neq j} q_n^* < \frac{q_{\max}}{\beta} - 1. \quad (3.5)$$

The node with the largest q_n^* value represents the largest possible source of multiple-access interference. Hence the receiver node i begins by finding the node with the largest q_n^* value and adding it to its silence set $S_{i,j}$, thus removing it from the interference sum in (3.5). This node selection process continues until the interference summation constraint is satisfied.

Complete knowledge of the vector of q_n^* values from all other nodes in the network allows straightforward calculation of the silencing sets such that (3.5) is satisfied. If a node is not within communication range of all other nodes or if external sources of

interference are present, then the node cannot account for all components of the interference summation in (3.5). However, the node will have an ordered vector of q_n^* values corresponding to the neighboring nodes that are potential sources of multiple-access interference. The significant sources of multiple-access interference are *typically* from nodes that are within communication range of the receiver. So, the silencing set can be initialized using just the q_n^* values from the neighboring nodes. In Section 3.6 we describe the MAC backup process which permits additional nodes to be placed in the silencing set to account for interference that is from distant nodes or external sources of interference. The silencing set initialization procedure is repeated for each link so that every receiver node i has a silencing set $S_{i,j}$ for each possible transmitter node j .

Each node is responsible for maintaining a vector of q_n^* values for nodes from which packet transmissions are received. A particular q_n^* measure is available for every point-to-point reception. Similarly, a q_n^* measure is also available from a node that broadcasts a periodic control packet that is transmitted with a specified power level and spreading factor. Each node maintains a weighted average of the q_n^* values for each neighbor node from which it obtains packet transmissions for use in silencing set formulations. A key feature of the silencing protocol is that the only information used to determine silencing sets is self-contained at each node in this vector of q_n^* values. Maintaining the silencing set requirements for a particular link is the responsibility of the receiving node. Hence, reevaluation of the silencing set at a particular node does not require any additional communications overhead and is a function of the vector of q_n^* values at a particular node.

For static networks, which are considered in this investigation, the weighted average of q_n^* values for each node does not change significantly as new q_n^* measures are obtained. In highly mobile environments, however, significant changes to the q_n^* vector occur and the spatial dependencies between nodes is altered. Reevaluation of the silencing sets might be required to account for different sources of multiple-access interference. The application of our cross-layer protocol to highly mobile networks remains an area of future research. Within a static network, however, silencing sets are formed from a given vector of q_n^* values and necessary additions to the silencing set are permitted as determined by the MAC layer backup process described in Section 3.6.

3.5 RTS acceptance criteria

Upon reception of an RTS packet, a node determines whether the requested transmission should be accepted. In order for a requested transmission from node a to node b to occur, the receiver b must first ensure that its silencing set criteria corresponding to a transmission from node a are satisfied. This silencing set $S_{b,a}$, which is maintained at the receiver b , contains the list of nodes that must not transmit during the reception from transmitter a . Node b compares this list of nodes with its local NAV to determine which nodes in the set $S_{b,a}$ are transmitting. If any of the nodes in the set $S_{b,a}$ are currently transmitting, then the requested transmission from node a is not accepted due to the imposed silencing restrictions and node b does not transmit a CTS packet. However, if no nodes in the set $S_{b,a}$ are currently transmitting according to the local NAV at node b , then the requested transmission is accepted by node b and a CTS packet containing the list of nodes in $S_{b,a}$ is transmitted. It should be noted that the silencing sets, which are maintained at each receiver node, are specific to the node requesting the

transmission. Thus, depending on the neighbor-state information in the local NAV, a receiver might reject an RTS packet from one transmitter but allow a transmission from a different transmitter which has less stringent silencing set requirements.

3.6 MAC backup process

The overall goal of the selective silencing is to govern channel access mechanisms so that for a given link, the adaptive-transmission protocol is able to mitigate the multiple-access interference. As described in Section 3.4, the silencing sets are calculated at node i after the q_j^* estimates are obtained from all neighboring nodes j . If the vector of q_j^* values accounts for all the multiple-access interference that the particular node is exposed to, then the silencing set formulation is straightforward. However, the adaptive-transmission protocol provides an estimate for E_s/N_0 (and hence q_j^*) for node j only if a transmission from node j has been acquired. Thus, it is possible that multiple-access interference from unaccounted nodes exceeds the allowable interference margin after silencing sets have been initialized. In this case, the adaptive-transmission protocol is unable to overcome the multiple-access interference with the current set of nodes being silenced.

We declare that a *MAC failure event* has occurred if the adaptive-transmission protocol fails to decode a packet when the maximum spreading factor N_{\max} is used for a packet transmission. A MAC failure event can occur for two reasons. First, the silencing operation may not function correctly if one or more nodes fail to receive a CTS packet. A node may not be aware that it has been silenced, or a node accepting a RTS may not be aware that a node in the silencing set is transmitting. In this situation a change to the silencing set is not warranted because the MAC failure event is due to a control channel

failure rather than multiple-access interference by nodes not included in the silencing set. To prevent an unnecessary alteration to the silencing set from occurring for every instance of this type of failure, the MAC backup process requires that α consecutive transmissions for a particular link result in a MAC failure event before additional nodes are placed in the silencing set.

The second reason for a MAC failure event occurs when the current set of nodes in the silencing set for a particular link does not suffice to limit the multiple-access interference experienced at the receiver. This situation arises if there is interference that has not been accounted for in the original silencing set calculations and it is necessary that a more restrictive silencing set be employed for this link. The silencing set is increased by adding the node j with the largest q_j^* value that is not already in the silencing set. If the link continues to experience significant multiple-access interference, the process of placing an additional node in the silencing set continues until the link can achieve successful packet transmissions without triggering α consecutive MAC failure events.

It should be noted that it is possible for the number of nodes that are in a particular link's silencing set to be equal to the number of nodes in range of the receiver for the particular link. For example, this is possible on a link for which EINR is only slightly larger than β and the maximum spreading factor N_{\max} is being utilized. In this example, even a small amount of additional multiple-access interference is likely to prevent successful reception of the packet transmission. It is not possible to add nodes to the silencing set to mitigate the multiple-access interference because the remaining sources of interference are out-of-range. The MAC layer silencing mechanism is able to identify

such links by the number of nodes that are contained in the corresponding link's silencing set. If a receiver has a large number of neighbors and the silencing set for a link contains a large fraction of the neighbors, the link will be difficult to utilize because it is unlikely that all the nodes in the silencing set are not transmitting data packets when an RTS packet is received. Furthermore, with a large fraction of nodes in the silencing set to consider, it is more likely that the status of the nodes in the silencing set may be inconsistent with the information in the NAV at the transmitter or receiver, leading to unexpected multiple-access interference. Thus, link selection plays a critical role in this type of network. Section 4.6 discusses link selection strategies and its implications on network performance.

3.7 Transmitter-side validation

The proposed silencing mechanism as described thus far is implemented at a receiving node based on information derived from the adaptive-transmission protocol. One of the underlying motives of this protocol design is to avoid introducing excessive communication overhead that could drive down network performance. As described in the previous sections, the silencing set initialization uses side-information already available from the adaptive-transmission protocol. Operation of the protocol requires no additional communications overhead other than embedding silencing set information into the CTS packets.

An additional mechanism which provides a layer of added protection to the silencing operation is a transmitter-side validation check that improves the ability of the silencing mechanism to operate efficiently. If a node wishing to initiate a transmission has knowledge of the silencing requirements at the receiver, it can avoid sending RTS

packets that would be deemed infeasible at the receiver. Because the transmitter also maintains neighbor-state information in its NAV, an additional NAV validation can occur at the transmitter prior to sending an RTS packet. However, the silencing set for each link is maintained at the receiver. Thus, in order for a transmitter-side check to occur, the transmitter must obtain the silencing set for a receiver before initiating a transmission to that receiver.

One byproduct of our silencing mechanism is that the silencing set is embedded in the CTS packet. Thus after node a transmits an RTS packet to node b and receives the corresponding CTS packet, node a learns $S_{b,a}$. Prior to sending subsequent RTS packets, node a compares the nodes specified by $S_{b,a}$ with its local NAV to determine if any of the nodes are currently transmitting. If node a determines that any node specified in $S_{b,a}$ is transmitting then it will refrain from sending node b an RTS packet until the interfering transmission is complete. Thus, the responsibility of NAV validation occurs at both the transmitter and receiver nodes. This dual-validation model increases the effectiveness of the silencing mechanism against missed control packets. After the silencing set is altered for a receiver, the transmitter updates its copy of the receiver's silencing set after the next exchange of RTS and CTS packets. It should be noted that the transmitting node a might not be within communications range of a node specified in the set $S_{b,a}$. In this situation, the responsibility of validating that this node is not transmitting is the responsibility of the receiver. Operation of this transmitter and receiver validation mechanism is evaluated in Section 4.4.

CHAPTER 4

SIMULATION MODEL AND RESULTS

4.1 Protocol strategies

We establish a set of base protocols that employ various transmission strategies so that we can evaluate the performance of our adaptive-transmission with selective-reuse protocol in a variety of network scenarios. We create a simulation framework so that the individual components of the cross-layer protocol can be studied independently as well as jointly to determine the effect that various protocol strategies have on overall network performance. We first consider a *single-transmission* protocol that attempts to silence all possible sources of multiple-access interference prior to transmission of a data packet. All data packets are sent with a fixed spreading factor of N_{\max} in this protocol. The single-transmission protocol is a special case of the adaptive-transmission with selective-reuse protocol in which the number of spreading levels is set as $M = 1$ and the silencing set for each link consists of all neighbors of the receiver. In addition to the single-transmission protocol, we consider a *fixed with reuse* protocol which implements the MAC silencing as described in Chapter 3 without spreading adaptation. The spreading factor is fixed at N_{\max} for this protocol so that the focus is on the MAC enhancements only. Additionally, we consider the performance of an *adaptive-spreading* protocol that does not use any form of MAC silencing. This protocol is restricted to using spreading adaptation only to mitigate multiple-access interference. This protocol is capable of selecting one of M spreading levels in steps of 3.0 dB with a maximum spreading factor

of N_{\max} according to the adaptive-transmission protocol description given in Chapter 2. The *adaptive with reuse* protocol implements the MAC silencing while also allowing spreading adaptation.

In addition to evaluating performance of our cross-layer protocol against a spectrum of approaches, we consider specific evaluation of transmission strategies within the framework our own protocol. Specifically, we address the performance tradeoffs associated with link selection strategies and the effect that link utilization has on overall network performance. We use throughput as the primary metric for performance evaluation. In later sections we provide multiple criteria for throughput measures to provide additional insight into protocol performance.

4.2 Simulation parameters

An OPNET simulation which makes use of detailed physical and link-layer models is used to test the performance of the various protocols. As described previously, the physical-layer implementation follows the work of [17], [18], and [19] to model acquisition and packet decoding error probabilities. Custom modules implement the distributed MAC operations. Parameters that are common to all simulation scenarios are given in Table 4.1. The maximum spreading factor (N_{\max}), number of spreading levels (M), number of network nodes (Q), and consecutive MAC failure threshold (α) depend on the details of the scenarios, and they are specified in the subsequent sections.

Table 4.1 Parameters used by the simulation model

Simulation Parameter	Parameter Symbol	Parameter Value
Lower adaptation threshold	β	6.5 dB
Thermal noise	N_0	2.0×10^{-21} W/Hz
Transmit power	P_i	1.0 W
Path-loss index	ξ	3.0
Carrier frequency	f_c	1.0 GHz
Receiver bandwidth	W	4 MHz
Data packet length	L	10000 bits
Control packet length	L_c	100 bits
Chip duration	T_c	2.5×10^{-7} sec
Upper adaptation threshold	σ	9.7 dB
Consecutive receptions counter	ϕ_{\max}	3
Forwarding attempts	ψ	3

Each node that generates data packets uses an independent Poisson process. The destination for a packet depends upon the scenario under investigation and is defined in the subsequent sections. However, forwarding is not considered in this investigation. We first examine several small network topologies to highlight the performance of the protocols under specific conditions. For these investigations we consider the throughput for each link, measured in packets per second, as the packet generation rate is varied. We also examine randomly generated topologies with 10 and 20 nodes with a very heavy traffic demand. For these simulations, a node always has packets waiting to be transmitted to each possible destination. The performance metrics considered for these simulations are described in Section 4.5.

4.3 Octagonal network

We begin evaluation of our cross-layer protocol by considering a simple network of eight nodes oriented in an octagonal topology with four traffic demands competing for access to the channel. Consider octagonal network scenario 1 with the four traffic

demands shown in Figure 4.1 where $d_1 = 250$ meters and $d_2 = 1000$ meters. All of the nodes in this network are located so that a transmission by one node can be received by all other nodes. For this scenario, we observe the average throughput for each link as the packet generation rate for each link is increased. The average throughput of each link in the network as a function of the packet generation rate is given in Figure 4.2 for the four protocols under consideration.




Figure 4.1 Octagonal network scenario 1

Both the single-transmission and the fixed with reuse protocols employ $M = 1$ spreading levels and a fixed spreading factor of $N_{\max} = 64$ for this network scenario. The difference between these two protocols is that the fixed with reuse protocol allows concurrent transmissions to occur whereas the single-transmission protocol uses the RTS/CTS exchange to silence all other nodes. It is assumed that the network has been initialized and that silencing sets have been formed at each node as described in Chapter 3. Consider the MAC silencing operation for the transmission from node E to node A . Node E begins by transmitting an RTS packet to node A . Upon reception of the RTS

packet from node E , node A checks its local silencing set to determine the silencing requirements that must be satisfied to allow transmission from node E to occur. For this scenario, the silencing set $S_{A,E}$ is empty because the energy in the desired signal from node E can more than offset the multiple-access interference from the remaining nodes. In other words, the interference summation as defined by (3.5) for this link is satisfied without adding any neighbor nodes to the silencing set. The same operation occurs for all four links considered in this scenario. Thus, the immediate benefit of spatial reuse is seen in this example because the four transmissions can occur simultaneously. In addition to the performance gain induced by spatial reuse is the benefit that can be obtained by allowing adaptation of the spreading factor. The adaptive-transmission protocol employs the same value for $N_{\max} = 64$, and allows adaptation among $M = 3$ spreading levels (i.e., 64, 32, and 16). For the scenario illustrated in Figure 4.1, it is possible for the spreading factor to be reduced by two levels and still maintain an acceptable signal quality. Thus, a performance gain over the fixed with reuse protocol is possible because the spreading factor is reduced by the adaptive-transmission protocol. For this scenario, the MAC silencing ability does not provide a throughput performance gain over what is achievable without a MAC silencing mechanism. This is simply because this particular set of traffic demands does not warrant any silencing of neighboring transmissions.




Figure 4.2 Throughput curves for octagonal scenario 1

Using the same network topology as Figure 4.1, consider a different set of four traffic demands as shown in Figure 4.3. This scenario depicts a case of the near-far interference problem common to ad hoc networks. Specifically, a transmission from node B to node A experiences significant multiple-access interference from node E if there is a concurrent transmission from node E to node F . Hence, each silencing set specifies that the closest interfering transmitter be silenced for this scenario. That is, the silencing sets are $S_{A,B} = E$, $S_{F,E} = C$, $S_{D,C} = H$, and $S_{G,H} = B$.




Figure 4.3 Octagonal network scenario 2

The same four transmission strategies are considered for this scenario. The average link throughput as a function of the packet generation rate is shown in Figure 4.4 for each of the transmission strategies. The generation rate is the same for each of the four links. This scenario employs N_{\max} equal to 64 for all protocols and allows three spreading levels for the adaptive protocols.

The single-transmission protocol operates in a similar manner to the previous network scenario and achieves the same throughput performance. However, the silencing requirements for this scenario are different than the previous network due to the different traffic demands. Before node A can accept an RTS packet from node B , it must determine that its silencing set criteria have been met according to its neighbor-state information about current transmissions stored in its local NAV. If node A determines that node E is not currently transmitting, then it can proceed with reception from node B and broadcasts a CTS packet with the information that node E is silenced for the duration of the pending packet transmission. This exchange notifies node E that it should prohibit packet transmissions to node F during this time. After the data transmission from node B to node A has started, assume that node H sends an RTS packet to node G . Node G

determines that the transmission cannot occur because its silencing set restrictions $S_{G,H}$ are violated (i.e., node B is already transmitting). Hence, the silencing restrictions imposed at node G prevent G from accepting a request to receive from node H until node B is finished transmitting. The only other data transmission that is possible is the transmission from node C to node D . The topology and traffic demands for this network are such that the only simultaneous transmissions that can be successful are the transmissions from B to A and C to D or the transmissions from E to F and H to G . The fixed with reuse protocol allows at most two transmissions to occur concurrently, and each transmission employs a spreading factor equal to N_{\max} . The adaptive with reuse protocol allows the same transmissions to occur and the adaptive-transmission protocol reduces the spreading factor by one level. The performance of the adaptive-spreading protocol in octagonal scenario 2 is extremely poor due to the significant multiple-access interference experienced at each receiver. The adaptive-transmission protocol does not have the capability to overcome the multiple-access interference experienced for these links. Hence, the need for a channel-access mechanism to incorporate temporal separation of specific transmissions is evident.




Figure 4.4 Throughput curves for octagonal scenario 2

Both octagonal network scenarios illustrate how the MAC silencing mechanism operates for the particular traffic demands. Significant performance gains are seen for the protocol that employs channel-reuse capabilities and the protocol that employs channel-reuse capabilities integrated with spreading adaptation. Additionally, the performance gains achievable from adaptive spreading over fixed spreading are apparent in both octagonal scenarios. Hence these two examples serve to illustrate the need for both spreading adaptation and selective channel reuse.

4.4 Neighbor-state information

Performance of the distributed protocol depends on the ability of a node to obtain neighbor-state information regarding transmissions by any of its neighbors that are contained in any of its silencing sets. Specifically, CTS packets are transmitted using the

common spreading pattern allowing information regarding impending transmissions to be distributed to neighboring nodes as described in Chapter 3. We consider four topologies shown in the following figures in which a given node is susceptible to significant multiple-access interference from two interfering transmissions. We consider the ability of the *primary link* from node A to node B to sustain throughput as the traffic load is increased over the interfering links from node C to node D and node E to node F . The traffic demand for the primary link is set so that there is always traffic to send from node A to node B . We evaluate the throughput over the primary link as the packet generation rate is increased for the interfering links.

The dotted circular regions surrounding nodes A and B represent the approximate transmission and reception ranges for the control packets sent or received by the nodes. It should be noted, however, that the simulation model does not use a hard threshold for packet decoding success probability and that these regions are shown for instructive purposes only. These scenarios are selected such that the probability of receiving control packets from nodes outside of this region is very small. In particular, node A is unable to receive transmissions from nodes C or E and vice versa. This topology is an example of the hidden terminal problem in which a transmitter is unaware of nodes that are close to a given receiver. The performance of the link from node A to node B is dependent on the MAC silencing operation at node B . The silencing operation is in turn dependent on the ability of node B to recognize and coordinate with the possible interfering transmissions originating from nodes C and E .

We examine the ability of our adaptive with reuse protocol to operate in these various network environments. We first consider the adaptive-spreading protocol in

which the adaptive-transmission protocol sets the spreading factor for a data transmission in response to feedback from previous transmissions. No MAC silencing occurs for this particular protocol. We also consider comparisons with the single-transmission protocol which operates in the same manner as described in Section 4.1 where each CTS packet will silence all nodes that acquire and decode it. In order to evaluate the benefit of the transmitter-side validation described in Section 3.7, we also consider a variation of the adaptive with reuse protocol in which the transmitter-side validation process is omitted.

The importance of the different locations of these nodes becomes apparent when considering the operation of the neighbor-state information acquisition process. Normal operation of the control channel as described in Chapter 3 requires that CTS control packets be transmitted with the common spreading pattern so that any nodes tuned to the control channel can acquire and decode them to obtain the necessary neighbor-state information. The various scenarios represented below depict various orientations of transmissions so that interaction of silencing sets can be analyzed. In all network scenarios, the silencing set requirements established at node B stipulate that both nodes C and E refrain from transmitting on the data channel in order for the transmission of a data packet from node A to occur successfully. The transmissions from node C to node D and node E to node F do not incur any silencing set restrictions at the receiver nodes D and F , respectively. These scenarios evaluate the ability of the transmission from node A to node B to sense the possible interfering transmissions from nodes C and E and obtain access to the channel.

In scenario 1, shown in Figure 4.5, node B is able to receive CTS packets sent from nodes C , D , E , and F . Node B is able to determine when the transmissions from nodes C

and E are occurring from the CTS packets broadcasted from nodes D and F , respectively. Hence, node B will only allow transmissions from node A to occur when it determines that its silencing set requirements are satisfied. After receiving an RTS packet from node A , if node B is not already blocked, it transmits a CTS packet which serves to notify both nodes C and E that they are required to remain silenced for the duration of the packet transmission from node A to node B . As expected, the transmission from node A to node B has a more difficult time accessing the channel than the interfering transmissions due to the presence of silencing restrictions imposed at node B for the protocols which incorporate channel reuse.

Figure 4.6 shows the throughput of the primary link (shown as a solid line) and the average throughput of the interfering links (shown as a dashed line) for the protocols under consideration. For this protocol, the primary link requires a spreading factor of N_{\max} for successful transmission to occur even if no multiple-access interference is present. It should be noted that in this first scenario, the protocol without transmitter-side validation is the same as the default adaptive with reuse protocol. This is due to the fact that node A is unable to obtain any neighbor-state information regarding transmissions that originate at nodes C and E . In this and subsequent scenarios the throughput performance of the primary link is significantly degraded for the adaptive-spreading protocol which does not make use of selective silencing. The primary link is unable to overcome the multiple-access interference with spreading adaptation only, so as the traffic load on the interfering links is increased the probability that a transmission on the primary link is successful decreases. Figure 4.6 shows that the silencing mechanism based on the neighbor-state information at node B achieves better throughput

performance on the primary and interfering links than the single-transmission protocol. It should also be noted that the adaptive-spreading protocol operates with a bias towards the two interfering links in this scenario because there is no need for any silencing requirements. Hence, for the purposes of evaluating overall network performance, it is important to consider the throughput achieved by each link and not just total network throughput. We address this method of performance evaluation in Section 4.7 where we introduce alternative throughput performance measures for larger networks.




Figure 4.5 Network topology of example scenario 1




Figure 4.6 Performance curves for network topology 1

In scenario 2 shown in Figure 4.7, node B is still able to overhear CTS packets from nodes C , D , E , and F . However, node A is also able to overhear CTS packets from nodes D and F . Hence, when node D or F broadcasts a CTS packet declaring that it is preparing to receive, both nodes A and B have an opportunity to obtain the transmission information. The transmitter-side validity check described in Section 3.7 is thus able to provide an additional check to determine if the proposed transmission is likely to be successful prior to transmission of the RTS. The silencing set for a transmission from node A to node B , denoted $S_{B,A}$, is known at receiver node B . During the first transmission from node A to node B , node B notifies node A of the particular silencing set corresponding to node A . For this scenario the set $S_{B,A}$ consists of nodes C and E . For subsequent transmissions from node A to node B , node A first checks its local neighbor-

state information to determine if nodes C or E are currently transmitting. If node A determines that the silencing set criteria specified by $S_{B,A}$ is violated, it will refrain from transmitting an RTS packet. Node A will continue to monitor its local NAV until the silencing set criteria are satisfied at which point it can then transmit an RTS packet after a pacing delay. Figure 4.8 relates the throughput performance curves for this particular scenario. Included in this figure is the throughput curve corresponding to operation of the adaptive with reuse protocol which omits the transmitter-side check of the neighbor-state information contained at node A . As can be seen, a slight reduction in throughput performance arises from omitting this transmitter-side validation check. The dual validation process allows the silencing set responsibility to be distributed partially to the transmitting node for this scenario. A more significant result based on the transmitter-side validation check can be seen in the following scenario.




Figure 4.7 Network topology of example scenario 2

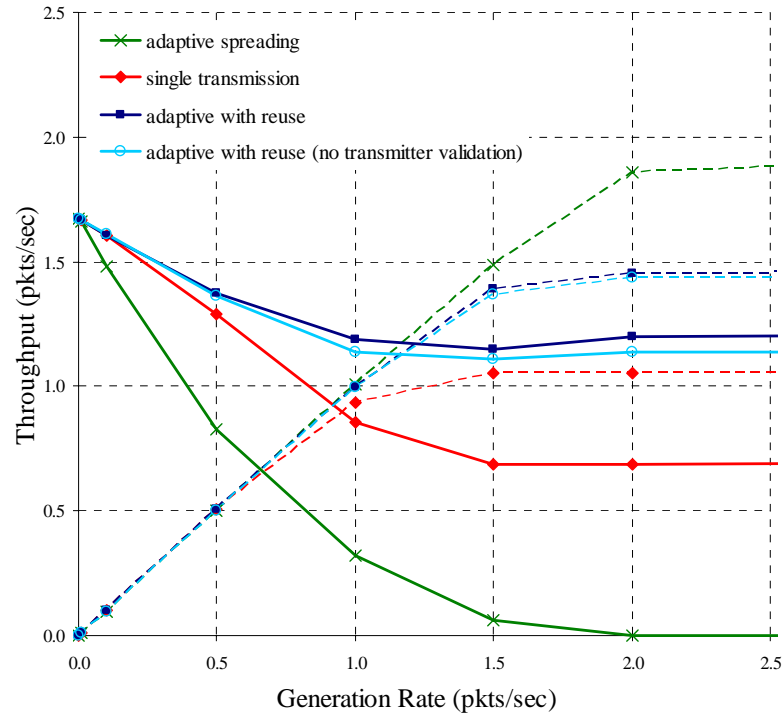


Figure 4.8 Performance curves for network scenario 2

In scenario 3 shown in Figure 4.9, only node A is able to overhear CTS packets transmitted from nodes D and F . In this scenario, the transmitter-side validation process is the only method to verify the silencing set requirements. Because node B is out of range of nodes D and F , it is unable to obtain information about transmissions originating from nodes C and E in this scenario. Figure 4.10 shows the resulting throughput performance for this network. From the graph it can be observed that the primary link experiences significant degradation when the transmitter-side validation is omitted. The transmitter-side neighbor-state information is sufficient to implement the selective silencing restrictions for the links in this network.




Figure 4.9 Network topology of example scenario 3




Figure 4.10 Performance curves for network scenario 3

In scenario 4, shown in Figure 4.11, nodes *A* and *B* are unable to receive the CTS packets transmitted from nodes *D* and *F*. The throughput performance of the primary link and the corresponding interfering links for this scenario is given in Figure 4.12. It is apparent that the normal control channel operation of including neighbor-state information in CTS packets does not suffice for this particular network topology. The problem is that the silencing set restrictions are not able to be effective because nodes *A* and *B* are unable to discern when transmissions occur on the interfering links.

A solution to this problem is to also transmit the RTS packet with the common spreading pattern so that node *B* is able to determine when transmissions are taking place from nodes within its silencing set. One drawback to this approach is that neighbor-state information obtained from RTS packets does not necessarily represent transmissions that have been accepted by the receiving node. CTS packets contain more accurate neighbor-state information because they represent proposed transmissions that have received dual-validation checks at the nodes in question. However, if a node is unable to obtain control packets from the receiving node of an interfering transmission as is the case in example scenario 4, then it must rely on broadcasted RTS packets to obtain neighbor-state information for the purposes of driving the silencing mechanisms. Figure 4.13 shows the resulting improved performance on the primary link when the use of the common spreading pattern for the RTS transmissions is integrated into the protocol.




Figure 4.11 Network topology of example scenario 4




Figure 4.12 Performance curves for network scenario 4




Figure 4.13 Performance curves for scenario 4 with RTS broadcasting

It should be noted that using the common spreading pattern for the RTS transmissions is a viable option to eliminate the hidden terminal problem represented by scenario 4. However, using the common spreading pattern for all RTS packets in a network with a larger number of nodes can cause excessive congestion on the control channel and lead to significant degradation in throughput. The use of the common spreading pattern for RTS packets is not used in the subsequent investigations, which consider networks of considerably larger size.

This study of four network topologies serves to illustrate how the silencing mechanism interacts with the collection of neighbor-state information in specific scenarios. The underlying motivation behind this thesis is to study the interaction of a cross-layer silencing mechanism in the scope of larger networks where significant

contention for channel-access exists. The networks considered in subsequent sections explore a more random traffic generation model in which all nodes operate as both transmitters and receivers. Hence the interdependency of specific links, which governs throughput performance for the previous network examples, is now overshadowed by the operation of the silencing set restrictions on the network as a whole. We evaluate overall network performance while still maintaining a notion of link performance as described in the next section.

4.5 Random network topologies

We extend performance evaluation of the previously described protocols to include randomly distributed networks of Q nodes that are both fully connected and not fully connected. Two nodes are considered to be *connected* (or neighbors) if they are able to receive control packet transmissions from each other. A network is *fully connected* if every node is connected to every other node in the network. We establish a heavy traffic demand across all viable links between neighboring nodes. Our evaluation observes link capacities across all $Q*(Q-1)$ links in the fully-connected network cases. In networks that are not fully connected, traffic demand exists between connected nodes only. The *link connectivity factor*, denoted C , represents the percentage of the $Q*(Q-1)$ links in the network that are between connected nodes and hence used for traffic transmissions. This study considers network topologies that are randomly distributed in a square region of R by R meters. We consider networks with R equal to 2000, 4000, 6000, and 8000 in this study. It should be noted that the topologies created in the 2000 by 2000 meter square region are fully connected. All networks with R equal to 4000, 6000, or 8000 have some nodes that are not in communications range and thus are not fully connected. All

performance results are averaged over 10 randomly generated networks for each size network. The average connectivity factor for the various network sizes are given in Table 4.2 as well as the average number of neighboring nodes for a given node.

Table 4.2 Average link connectivity and number of neighbor nodes for various network topologies

Network Size	Number of Nodes (Q)	Avg. Link Connectivity (C)	Avg. Neighbor Nodes
2000 by 2000 meters	10	100.0%	9.0
	20	100.0%	19.0
4000 by 4000 meters	10	94.0%	8.4
	20	94.4%	17.9
6000 by 6000 meters	10	61.3%	5.5
	20	62.6%	11.9
8000 by 8000 meters	10	40.4%	3.8
	20	40.9%	7.8

Within the cross-layer design framework, we consider different strategies for selecting links at each node. Unlike the previous examples in which each node had traffic destined for only one other node, we now consider a traffic model in which a node must send independent traffic to multiple destinations. A heavy-traffic model is assumed in which each node always has packets to send to each of its neighbors. A node that is not blocked from transmitting uses the following algorithm to select a neighbor. A node first determines which of its neighbors are candidates for packet reception based on its local NAV and the silencing set restrictions of each neighbor. If the node finds that a neighbor is already transmitting or receiving then it is deemed ineligible as a possible receiver. Next, the node eliminates a candidate neighbor if the silencing set restrictions are violated by the current transmissions. Each candidate neighbor that passes both tests is eligible for a packet transmission attempt.

Among eligible receivers, we consider two strategies for selecting the link on which a transmission should be attempted. The first strategy employs a metric for link selection based solely on prior link usage for each of the eligible receivers, and we denote this as the *capacity fairness* strategy. Capacity fairness observes the total link throughput that has been achieved over each currently eligible link and chooses the receiver that has received the fewest packets from this node. Thus, this strategy seeks to transmit on underutilized links that are deemed eligible for transmission. All channel-access protocols evaluated in this section employ this strategy unless noted otherwise.

We also consider a link selection strategy based on the on-air transmission time of each link. The *on-air fairness* strategy is slightly different from capacity fairness in that the total link throughput is scaled by the transmission time of the packets. In other words, this strategy considers the time used to transmit packets to a particular receiver, which is directly proportional to the average spreading factor used over the link, rather than the actual number of packets sent over the link. Hence, the on-air fairness metric is biased towards using links with a smaller spreading factor whereas the capacity fairness protocol emphasizes the requirement that all eligible links from a particular transmitting node achieve equal throughput. Both link selection strategies are evaluated with our adaptive with reuse protocol

We establish two primary criteria to measure the performance of the channel-access protocols for fully-connected networks. We define *total throughput* as the throughput achieved when summed over all links in the network. To investigate fairness, we also determine the throughput for each link and find the minimum throughput T_{\min} over all the links in the network. The *base throughput* is equal to $Q*(Q-1)*T_{\min}$, and is a measure of

fairness that represents the fraction of the total throughput that is equally distributed among all links. Figures 4.14 and 4.15 give the total and base throughput comparisons for each channel-access protocol in the fully-connected networks with 10 and 20 nodes, respectively. All protocols use the capacity fairness strategy with the exception that we now evaluate both link selection strategies for our adaptive with reuse protocol. For all results shown in this section the maximum spreading is 128 for all protocols and $M = 4$ for protocols that utilize the adaptive-transmission protocol. For the protocols which employ MAC silencing, the consecutive MAC failure threshold is set to $\alpha = 10$. The first investigations are focused on the performance for a fully-connected network in which all links should be utilized. This value for α effectively disables the MAC backup mechanism. Further consideration of the MAC backup process and its relation to throughput performance is given in Section 4.6.




Figure 4.14 Throughput performance for fully-connected networks of 10 nodes with $\alpha = 10$




Figure 4.15 Throughput performance for fully-connected networks of 20 nodes with $\alpha = 10$

From the figures it is evident that the protocols which employ spreading adaptation, selective channel reuse, or a combination of both achieve significantly better total throughput than the single-transmission protocol. One drawback of the adaptive-spreading protocol, as mentioned in prior sections, is that some links are unable to overcome the significant interference caused by the near-far interference problem. This is reflected in the performance of the adaptive-spreading protocol which achieves a base throughput that is approximately zero. Although this protocol achieves good total throughput performance, the inability of this protocol to achieve base throughput of any significance makes it unattractive. Both of the adaptive with reuse protocols are able to take advantage of the total throughput performance gain derived from spreading adaptation while maintaining good base throughput performance. On the other hand, the fixed with reuse protocol provides good base throughput but moderate total throughput performance due to its inability to reduce the spreading factor when channel conditions permit such a change. The link selection strategy also has an effect on overall network performance. As expected, the capacity fairness protocol provides greater base throughput because it attempts to transmit on all links. On the other hand, the on-air fairness protocol provides greater total throughput because it takes advantage of attractive links which employ lower spreading factors.

We consider the evaluation of the same protocols when the networks nodes are randomly distributed over a larger region and hence are no longer fully connected. The base throughput measure is no longer applicable in these scenarios because it is not possible to evaluate link throughput between all nodes because some nodes are no longer connected. A different measure of throughput, which considers link fairness for networks

that are not fully connected, is described in the next section. Total throughputs for all network sizes are given below for networks of 10 and 20 nodes and with $\alpha = 3$. Figures 4.16, 4.17, 4.18, and 4.19 show the simulation results for networks with R equal to 2000, 4000, 6000, and 8000 meters, respectively. Total throughput performance in these results is similar to that of the fully-connected network scenarios.




Figure 4.16 Total throughput for 2000 by 2000 network topologies with $\alpha = 3$




Figure 4.17 Total throughput for 4000 by 4000 network topologies with $\alpha = 3$




Figure 4.18 Total throughput for 6000 by 6000 network topologies with $\alpha = 3$




Figure 4.19 Total throughput for 8000 by 8000 network topologies with $\alpha = 3$

4.6 Link utilization

The performance results presented in the previous section show that link selection plays a role in network performance. Because our protocol introduces distributed channel-access restrictions for each node, the ability of a particular link to successfully access the channel is dependent on the silencing restrictions imposed at each node. In other words, certain links will have a harder time accessing the channel and will thus not be able to support as large of a traffic demand as more capable links. Additionally, a link that requires a large number of nodes to be silenced for a transmission to occur requires a larger percentage of network resources because reuse of the channel is more restricted. An additional byproduct of our MAC protocol is that we are able to associate a channel-access cost to each link in the network. Specifically, the number of nodes in a silencing set has a direct correlation to the ability of a node to access the channel. We consider the effect of disabling certain links in the network which in turn allows the remaining links to

have more access to channel resources. In doing this, we are implicitly assuming that some higher-layer protocol is able to govern traffic routes to avoid certain links. Although implementation of a routing protocol is beyond the scope of this thesis, we seek to evaluate the effect that link selection has on network performance. In doing so, we are relating a derived link cost to overall network performance which can in turn be used as a partial metric to a link-cost based routing protocol. Exploitation of this relationship between channel-access restrictions and higher-layer protocols is an area of future research.


In order to establish a heuristic guideline for link utilization, we define the *neighborhood silencing threshold*, θ , as the maximum fraction of neighbors that a particular link is allowed to silence. Specifically, if a node b has P neighbors and the link from node a to node b requires that $|S_{b,a}|$ nodes be silenced, then the link will be used only if $|S_{b,a}|/P < \theta$. As nodes are added to the silencing set via initialization or the MAC backup process, a utilization check occurs to determine if the neighborhood silencing threshold has been exceeded. Once the threshold has been exceeded, the particular link is no longer used for packet transmissions.

We define *link utilization*, denoted U , for a particular network as the percentage of connected links that are used for packet transmissions. Previously, we defined the link connectivity factor, denoted C , as the fraction of links that are within communications range for a given network. We scale the link connectivity factor by the link utilization factor to provide the total number of links used for packet transmissions. As links are removed as a result of the neighborhood silencing threshold, the link utilization decreases. Hence, the total number of active links in the network is $U * C * Q * (Q - 1)$ where


Q is the number of nodes in the network. We now revisit the base throughput measure previously defined for fully-connected networks by introducing the *utilized base throughput* which is equal to $U * C * Q * (Q - 1) * T_u$, where T_u is the minimum throughput for all active links. This measure provides a metric for networks that are not fully connected so that individual link performance and fairness can be evaluated in addition to total network throughput. Part A of Figures 4.20, 4.21, 4.22, and 4.23 provides measures of link utilization as a function of the neighborhood silencing threshold θ for 20 node networks distributed in various topology sizes. We consider operation of the adaptive with reuse protocol employing capacity fairness for these results. For this evaluation, the consecutive MAC failure silencing threshold is set at $\alpha = 3$ so that three consecutive MAC silencing failures triggers an addition to the silencing set for the link in question. It should be noted that a neighborhood silencing threshold of 1.0 corresponds to a link that requires all neighboring nodes to be silenced for a transmission to occur.

Corresponding to the link utilization curve of Figures 4.20, 4.21, 4.22, and 4.23 is a throughput performance curve in part B of each figure. The performance curves give measures of utilized base throughput and total throughput as a function of the neighborhood silencing threshold. Throughput performance corresponding to a particular neighborhood silencing threshold must be considered jointly with the corresponding link utilization measure. As can be seen in the following figures, an increase in throughput performance is obtained via a decrease in link utilization. However, a small decrease in utilization can result in a significant increase in throughput performance. For example, the base throughput in the 6000 by 6000 meter topologies shown in Figure 4.22, can be increased by 102% for a reduction in link utilization from 85% to 81%. This corresponds

to a change in the silencing threshold from $\theta = 1.0$ to $\theta = 0.9$. Reducing the silencing threshold θ to 0.7 increases the base throughput by 203%. However, the link utilization corresponding to this reduction changes from 85% to 69% of the connected links.







(a) Link utilization vs. θ





(b) Throughput vs. θ

Figure 4.20 Performance of 2000 by 2000 network with $Q = 20$ and $\alpha = 3$

(a) Link utilization vs. θ (b) Throughput vs. θ Figure 4.21 Performance of 4000 by 4000 network with $Q = 20$ and $\alpha = 3$

(a) Link utilization vs. θ (b) Throughput vs. θ Figure 4.22 Performance of 6000 by 6000 network with $Q = 20$ and $\alpha = 3$

(a) Link utilization vs. θ (b) Throughput vs. θ Figure 4.23 Performance of 8000 by 8000 network with $Q = 20$ and $\alpha = 3$

A tradeoff exists between the throughput obtainable and the link utilization factor. As expected, removing poor links as declared by the neighborhood silencing requirement threshold results in significant total and utilized base throughput performance gains. However, reducing the number of links available for transmission puts a higher demand on a routing protocol to develop efficient routes between all nodes. This study cannot simply conclude that additional removal of links improves performance because we do not explicitly model the effect that the removed links have on higher-layer protocols. However, this study does establish that the correlation between silencing requirements and throughput performance provides a means for a higher-layer protocol to select attractive links. One byproduct of the MAC silencing process described in this study is that a channel-access cost has been established. Physical-layer statistics in conjunction with feedback behavior from the MAC silencing mechanisms provides a metric that is tractable to a routing protocol. This study has established a method to discriminate links that are likely to use excessive channel resources. The results shown in this section suggest that a protocol which avoids these particular links can significantly improve overall network throughput and the throughput for other links.

CHAPTER 5

CONCLUSIONS

In this investigation we consider the interaction of an adaptive-transmission physical-layer protocol with a selective channel reuse protocol at the MAC layer. Evaluation of protocol design in this study suggests that MAC layer operation should be topology dependent and that spatial dependencies can be exploited to improve network throughput performance. We formulate a cross-layer metric, derived from the adaptive-transmission protocol side-information, which allows link comparisons to be established at a receiver so that significant sources of multiple-access interference can be identified. We also describe a distributed heuristic to formulate silencing sets, which are used to govern transmission restrictions for a particular link.

Consideration is given to specific design issues associated with the cross-layer protocol under investigation. Specifically, link selection strategies which weigh previous link throughputs or on-air transmission times are examined to relate the significance that the link selection has on various throughput measures. Additionally, we consider a link utilization metric that provides a means to strategically identify and deactivate poor links based on silencing set requirements. The resulting link utilization curves and the associated throughput performance gains are provided. This study shows that for a variety of network scenarios, our adaptive-transmission, cross-layer protocol significantly outperforms protocols that employ static transmission parameters or layer-independent operation.

APPENDIX

The purpose of this appendix is to describe the simulation model for the PDSQ statistics and an approximation for estimating E_s/N_0 . We also introduce a new approximation for estimates of E_s/N_0 that is applicable if there are significant levels of multiple-access interference. The original PDSQ model is described in [20]. Additional details on estimates of E_s/N_0 derived from PDSQ approximations are found in [21]. For convenience, the work in [20] and [21] is recreated in Section A.1. Section A.2 describes our new model for estimating E_s/N_0 and is based closely on the approach taken in [20] and [21].

A.1 A Possible Approximation for PDSQ

We assume that each signature sequence and each data sequence are random sequences of independent random variables, each of which is uniformly distributed on the set $\{-1, +1\}$. Signature sequences and data sequences for different signals are independent. The chip rate is fixed, so the spreading factor is varied by changing the duration of the data symbols. Let N denote the number of chips per data symbol for the desired signal at the time that the PDSQ statistic is to be determined.

We assume the sampling interval t_0 is equal to the chip duration T_c , and the chip waveform is the rectangular pulse of duration T_c . For the desired signal, one symbol has N chips, so there are N samples per symbol at the matched filter output and only a single sample in the main lobe. Thus, $M = \{1, 2, \dots, N - 1\}$ is the index set for the off-peak or side-lobe samples.

Consider a sequence of symbols indexed by i , and denote the peak statistic for the i th symbol by $Z_{i,0}$. The off-peak statistics for the i th symbol are denoted by $Z_{i,m}$ for $1 \leq m \leq N-1$. The PDSQ statistic for an individual symbol interval is the ratio of the square of the peak statistic to the sum of the squares of the off-peak statistics in the interval and is given by

$$Q_i = \frac{(Z_{i,0})^2}{\sum_{m=1}^{N-1} (Z_{i,m})^2}. \quad (\text{A.1})$$

If S denotes a set of N_s symbol positions for which a measure of the signal quality is desired define

$$Z_0(S) = \sum_{i \in S} (Z_{i,0})^2 \quad (\text{A.2})$$

and

$$X(S) = \sum_{i \in S} \sum_{m=1}^{N-1} (Z_{i,m})^2. \quad (\text{A.3})$$

A suitable PDSQ statistic is

$$Q(S) = \frac{Z_0(S)}{X(S)}. \quad (\text{A.4})$$

The separate summations over $i \in S$ that represent the numerator and the denominator of (A.4) suggest that expected values can be used to obtain an approximation for $Q(S)$. Because the random variables associated with the i th symbol have the same distribution as those associated with any other symbol, then $E\{(Z_{i,0})^2\}$ is the same for each i . Similarly, as long as $m \in M$ then $E\{(Z_{i,m})^2\}$ does not depend on either i or m . Let the random variable Z_0 have the same distribution as the random

variables $Z_{i,0}$. For each $m \in M$, let Z_m have the same distribution as the random variables $Z_{i,m}$. It follows that

$$\frac{E\{Z_0(S)\}}{E\{X(S)\}} = \frac{N_s E\{(Z_0)^2\}}{N_s E\{\sum_{m=1}^{N-1} (Z_m)^2\}} = \frac{E\{(Z_0)^2\}}{E\{\sum_{m=1}^{N-1} (Z_m)^2\}}. \quad (\text{A.5})$$

Consider a system with a single interfering spread-spectrum signal. The peak sample has three terms: one due to the desired signal, one due to the interference, and one due to the noise. If the amplitude of the desired signal is A_0 , the amplitude of the interference is A_1 , and the spectral density of the noise is $N_0/2$, then

$$E\{(Z_0)^2\} = (A_0 NT_c)^2 + (A_1 T_c)^2 N + \frac{N_0}{2} NT_c \quad (\text{A.6})$$

and

$$E\left\{\sum_{m=1}^{N-1} (Z_m)^2\right\} = (N-1)(A_1 T_c)^2 N + (N-1)\frac{N_0}{2} NT_c. \quad (\text{A.7})$$

Note that (A.7) assumes ideal spreading sequences so that autocorrelation sidelobes are zero. Substitute from (A.6) and (A.7) into the right-hand side of (A.5) and divide the numerator and denominator by NT_c to obtain

$$Q(S) = \frac{A_0^2 NT_c + A_1^2 T_c + \frac{N_0}{2}}{(N-1)\left[A_1^2 T_c + \frac{N_0}{2}\right]}. \quad (\text{A.8})$$

Let E_s denote the energy per symbol for the desired signal. It follows that $E_s = A_0^2 NT_c$. Define $E_1 = A_1^2 NT_c$, and notice that E_1 is the energy per symbol in the interference only if the interference signal has the same spreading factor as the desired signal. In general, if N_1 is the spreading factor for the interference symbol then E_1 is N/N_1 times the energy per

symbol in the interference signal. In terms of E_s and E_1 , the proposed approximation for the PDSQ for a set of symbols is

$$Q(S) = \frac{\frac{A_0^2 NT_c}{N_0} + \frac{1}{N} \left(\frac{A_1^2 NT_c}{N_0} \right) + \frac{1}{2}}{(N-1) \left[\frac{1}{N} \left(\frac{A_1^2 NT_c}{N_0} \right) + \frac{1}{2} \right]} = \frac{\frac{E_s}{N_0} + \frac{1}{N} \left(\frac{E_1}{N_0} \right) + \frac{1}{2}}{(N-1) \left[\frac{1}{N} \left(\frac{E_1}{N_0} \right) + \frac{1}{2} \right]}. \quad (\text{A.9})$$

Let L denote the number of channel symbols per packet and let $\{S_1, S_2, \dots, S_n\}$ be a partition of the set of channel symbols. We define the PDSQ statistics vector, denoted \mathbf{Q} , as $\mathbf{Q} = \{Q(S_1), Q(S_2), \dots, Q(S_n)\}$ where each S_i contains L/n symbols. We assume the power in the desired signal P_0 is constant over the entire packet where $P_0 = A_0^2$. Let the power in the interfering signals for each symbol set S_i be represented by $P_1(S_i) = A_1^2$. In our simulations we assume the interference power to be constant over the duration of each set S_i of symbols. Thus there can be at most n different measures of interference for each packet transmitted. Extending (A.8) the PDSQ can also be denoted

$$Q(S_i) = \frac{P_0 NT_c + P_1(S_i) T_c + \frac{N_0}{2}}{(N-1) \left\{ [P_0 + P_1(S_i)] T_c + \frac{N_0}{2} \right\}}. \quad (\text{A.10})$$

where the power in the desired signal is included in the denominator of the statistic. Under the assumption that there is no multiple-access interference for at least one of the n vector components, a reasonable estimate of E_s/N_0 can be derived from (A.10). Let $q = \max\{Q(S_i): 1 \leq i \leq n\}$, which corresponds to the interval with the least interference power present. Doing so assumes that the interference power in that interval is negligible (i.e., $P_1(S_i) = 0$). Also note the symbol energy E_s is given by $P_0 NT_c$. Thus (A.10) simplifies to

$$q = \frac{E_s + \frac{N_0}{2}}{(N-1) \left\{ \frac{E_s}{N} + \frac{N_0}{2} \right\}}. \quad (\text{A.11})$$

Solving for E_s/N_0 in (A.11) yields

$$\frac{E_s}{N_0} \approx f(q, N) = \frac{N \{q(N-1) - 1\}}{2 \{N - q(N-1)\}}. \quad (\text{A.12})$$

A.2 An Additional Approximation for E_s/N_0

Suppose that the power in the interfering signals for a given q interval is not equal to zero (i.e., suppose $P_1(S_i) = \alpha$). Let $I_\alpha = \alpha T_c$ be the energy in the sum of the interfering signals for this particular interval. Substituting into (A.10) gives

$$q = \frac{E_s + I_\alpha + \frac{N_0}{2}}{(N-1) \left\{ \frac{E_s}{N} + I_\alpha + \frac{N_0}{2} \right\}}. \quad (\text{A.13})$$

Solving for E_s/N_0 in (A.13) yields

$$\frac{E_s}{N_0} = f_I(q, N) = \frac{N \{q(N-1) - 1\} + \frac{I_\alpha}{N_0/2} [N \{q(N-1) - 1\}]}{2 \{N - q(N-1)\}} = f(q, N) \left[1 + \frac{I_\alpha}{N_0/2} \right]. \quad (\text{A.14})$$

Note that equation (A.12) is the special case of (A.14) where $I_\alpha = 0$. For a given value of $I_\alpha > 0$, the resulting estimate given by (A.14) would result in an E_s/N_0 approximation that is higher than that when the interference is assumed to be zero. Thus, omitting the I_α value in the q statistic leads to a lower estimate of E_s/N_0 than the actual value. Using I_α in the estimating function provides a more accurate measure of E_s/N_0

when significant interference is present. However, values for I_α are not readily available to include in an estimating function to approximate E_s/N_0 .

The main problem in the above formulation is that there are two unknowns embedded into the available measurements (P_0 and P_1). However, solving (A.10) for P_0 in terms of unknown P_1 , measured value q , and N_0 gives

$$P_0 = \frac{\{q(N-1)-1\}}{\{N-q(N-1)\}} \left[\frac{N_0/2}{T_c} + P_1 \right]. \quad (\text{A.15})$$

In addition to the aforementioned PDSQ statistics is the automatic gain controller (AGC) which provides a measure of the sum of the received power from the desired signal, interfering signals, and thermal noise during a packet reception.

$$\text{AGC} = P_0 + P_1 + \frac{N_0/2}{T_c} \quad (\text{A.16})$$

Substituting (A.15) into (A.16) and solving for P_1 gives

$$P_1 = \left[\frac{N-q(N-1)}{N-1} \right] \text{AGC} - \frac{N_0/2}{T_c}. \quad (\text{A.17})$$

Note that (A.17) provides an estimate of the interference power in terms of known parameters (N and T_c) and measured values (q , AGC, and N_0). Substituting this estimate of P_1 into equation (A.14) provides a new approximation of E_s/N_0 .

$$\frac{E_s}{N_0} = \frac{NT_c \{q(N-1)-1\}}{(N-1)N_0} \text{AGC} \quad (\text{A.18})$$

REFERENCES

1. A. J. Goldsmith and S. B. Wicker, "Design challenges for energy-constrained ad hoc wireless networks," *IEEE Wireless Communications*, vol. 9, no. 4, pp. 8-27, August 2002.
2. V. Bharghavan, A. Demers, S. Shenkar, and L. Zhang, "MACAW: A media access protocol for wireless LAN's," in *Proceedings of ACM SIGCOMM*, pp. 134-140, September 1990.
3. F.A. Tobagi and L. Kleinrock, "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*, vol. 23, pp. 1417-1433, December 1975.
4. M. B. Pursley, "The role of spread spectrum in packet radio networks," *Proceedings of the IEEE*, vol. 75, no. 1, January 1987.
5. M. B. Pursley and D. J. Taipale, "Error probabilities for spread-spectrum packet radio with convolutional codes and Viterbi decoding" *IEEE Transactions on Communications*, vol. 35, no. 1, Jan 1987.
6. F. J. Block and M. B. Pursley, "A protocol for adaptive transmission in direct-sequence spread-spectrum packet radio networks," *IEEE Transactions on Communications*, vol. 52, no. 8, August 2004.
7. S. W. Boyd, M. B. Pursley, and H. B. Russell, "An adaptive-transmission cross-layer protocol with selective MAC layer spatial reuse capabilities for ad hoc networks" in *Proceedings of the 2006 IEEE Military Communications Conference*, October 2006.
8. V. Kawadia and P. Kumar "Principles and protocols for power control in wireless ad hoc networks" *IEEE Journal on Selected Areas in Communications*, vol. 23, no. 1, January 2005.
9. S. Narayanaswamy, V. Kawadia, R. Sreenivas, P. Kumar "Power control in ad-hoc networks: Theory, architecture, algorithm and implementation of COMPOW protocol" in *European Wireless Conference*, February 2002.
10. T. Elbatt and A. Ephremides, "Joint scheduling and power control for wireless ad hoc networks," *IEEE Transactions on Wireless Communications*, vol. 3, no. 1, January 2004.

11. S. Toumpis, A. J. Goldsmith, "Capacity regions for wireless ad hoc networks," *IEEE Transactions on Wireless Communications*, vol. 2, no. 4, July 2003.
12. A. Muqattash and M. Krunz, "Power controlled dual channel (PCDC) medium access protocol for wireless ad hoc networks" in *Proceedings of the IEEE INFOCOM Conference*, vol. 1, pp. 470-480, April 2003.
13. A. Muqattash and M. Krunz, "POWMAC: A single-channel power-control protocol for throughput enhancement in wireless ad hoc networks," *IEEE Journal on Selected Areas in Communications*, vol. 23, no. 5, pp. 1067-1084, May 2005.
14. A. Swaminathan, D. L. Neneaker, and H. B. Russell, "Protocols for efficient spatial reuse of traffic channels in DS spread-spectrum packet radio networks with directional antennas," in *Proceedings of the 2005 IEEE Military Communications Conference*, vol. 5, pp. 2705-2712, October 2005.
15. M. B. Pursley and C. S. Wilkins, "Adaptive transmission for direct-sequence spread-spectrum communications over multipath channels," *International Journal of Wireless Information Networks*, vol. 7, no. 2, pp. 69-77, April 2000.
16. J. H. Gass, Jr., M. B. Pursley, H. B. Russell, and J. S. Wysocarski, "An adaptive-transmission protocol for frequency-hop wireless communication networks," *Wireless Networks*, vol. 7, no. 5, pp. 487-495, 2001.
17. K. Gajaraj, "An adaptive transmission protocol for direct sequence spread spectrum packet radio networks," MS Thesis, 2004.
18. D. L. Neneaker, A. R. Raghavan and C. W. Baum, "The effect of automatic gain control on serial, matched-filter acquisition in DS packet radio communication," *IEEE Transactions on Vehicular Technology*, vol. 50, pp. 1140-1150, July 2001.
19. A. R. Raghavan and C. W. Baum, "An unslotted multichannel channel-access protocol for distributed direct-sequence networks," *Mobile Networks and Applications*, vol. 5, no. 1, pp. 49-56, 2000.
20. M. B. Pursley, "An approximation for PDSQ," Unpublished notes, March 2004.
21. F. J. Block, "Estimate of E_s/N_0 ," Unpublished notes, April 2004.