REPORT DOCUMENTATION PAGE

Form Approved OMB NO. 0704-0188

1. REPORT DATE (DD-MM-YYYY)	2. REPORT TYPE	3. DATES COVERED (From - To)
17-03-2023	Thesis or Dissertation	-

4. TITLE AND SUBTITLE
A Transmission Scheduling Protocol Using Radios With Multiple Antennas For Ad Hoc Networks

5a. CONTRACT NUMBER
W911NF-17-1-0244

5b. GRANT NUMBER

5c. PROGRAM ELEMENT NUMBER
611103

6. AUTHORS
John, Hardaway

5d. PROJECT NUMBER

5e. TASK NUMBER

5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAMES AND ADDRESSES
Clemson University
300 Brackett Hall
Box 345702
Clemson, SC 29634 -5702

8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS (ES)
U.S. Army Research Office
P.O. Box 12211
Research Triangle Park, NC 27709-2211

10. SPONSOR/MONITOR'S ACRONYM(S)
ARO

11. SPONSOR/MONITOR'S REPORT NUMBER(S)
70086-NC-RIP.4

12. DISTRIBUTION AVAILIBILITY STATEMENT
Approved for public release; distribution is unlimited.

13. SUPPLEMENTARY NOTES
The views, opinions and/or findings contained in this report are those of the author(s) and should not contrued as an official Department of the Army position, policy or decision, unless so designated by other documentation.

14. ABSTRACT

15. SUBJECT TERMS

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	15. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			Daniel Noneaker
UU	UU	UU	UU		19b. TELEPHONE NUMBER 864-656-0100

REPORT DOCUMENTATION PAGE (SF298)
(Continuation Sheet)

Continuation for Block 13

Proposal/Report Number: 70086.4-NC-RIP

Report Title: A Transmission Scheduling Protocol Using Radios With Multiple Antennas For Ad Hoc Networks
Report Type: Ph.D. Dissertation

Publication Type: Thesis or Dissertation
Institution: Clemson University
Date Received: 17-Mar-2023 Completion Date: 5/15/20 8:09PM
Title: A Transmission Scheduling Protocol Using Radios With Multiple Antennas For Ad Hoc Networks
Authors: John, Hardaway
Acknowledged Federal Support: **N**

May 2020

A Transmission Scheduling Protocol Using Radios With Multiple Antennas For Ad Hoc Networks

John Benjamin Hardaway
Clemson University, johnhardawayv@gmail.com

A Transmission Scheduling Protocol Using Radios With Multiple Antennas 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
Computer Engineering

---

by
John Benjamin Hardaway V
May 2020

Accepted by:
Dr. Harlan B. Russell, Committee Chair
Dr. Kuang-Ching Wang
Dr. Linke Guo

Abstract

Distributed medium access control (MAC) protocols are essential due to the flexible and self-organizing nature of ad hoc networks. Scheduling protocols have been popular choices, because they guarantee access to the channel for each transceiver. The disadvantage with these scheduled approaches is that they are inefficient when the network has low traffic loads. Consider a time-division multiple access (TDMA) schedule where nodes are assigned time slots in which they are allowed to transmit. If a particular node is scheduled but has no traffic to forward, then the time slot is wasted. Because the channel has been reserved for use by that particular node, other nodes in the network with traffic to forward are unable to do so. We investigate strategies to improve the performance of TDMA scheduling protocols for ad hoc networks using radios with multiple antennas. Multiple antennas at each radio enables the use of a physical layer technique known as multiple-input multiple-output (MIMO) that leverages the spatial dimension. The antennas allow for both spatial multiplexing and interference cancellation. Spatial multiplexing allows for multiple parallel data streams to be transmitted at the same time. Interference cancellation is used to selectively pick neighbors that do not receive interference from a transmission. Our protocol uses both techniques to allow unscheduled nodes to transmit if the slot is not fully utilized. Using a custom simulation, we show that Lyui's scheduling protocol can be extended to support MIMO and time slot sharing. Our new protocol provides performance improvements with regards to end-to-end completion, throughput, and average delay.

Dedication

I dedicate this thesis to my fiancée, Cassidy, for believing in me, giving me the push I needed, and fixing more commas than she agreed to.

Acknowledgments

I would like to start by thanking my adviser, Dr. Harlan Russell, for our weekly meetings in which we discussed and brainstormed many ideas that worked and some that did not. Your guidance has been invaluable. I would also like to thank Dr. Linke Guo and Dr. Kuang-Ching Wang for serving on my committee. Lastly, I would like to thank all of my professors at Clemson University that have assisted me in getting to this point in my academic career.

Table of Contents

List of Tables

List of Figures

Chapter 1

Introduction

Wireless networks have become increasingly popular due to low cost and powerful transceivers. An ad hoc network is an example of a flexible self-organizing wireless network which does not rely on a preexisting infrastructure. In an ad hoc network, a transceiver serves the role of both the traffic source and sink as well as the intermediary router for traffic not destined for itself. Because ad hoc networks do not require preexisting infrastructure, they have been a popular choice for military applications in hostile environments and during disaster relief efforts when existing infrastructure has been damaged.

1.1 Medium Access Control

Wireless networks operate with all nodes using the same medium. This means that if two nodes decide to transmit at the same time, additional interference is created at the intended receivers. If the interference is too large, then a receiver may be unable to decode the intended message, and the packet may be lost. A network's sensitivity to interference caused by simultaneous transmissions depends on the modulation and coding schemes used at the physical layer. Due to the issues posed by simultaneous transmissions, medium access control (MAC) protocols are used to determine when a node may transmit. Because ad hoc networks cannot use a central controller, it is important that all MAC and network protocols work in a distributed manner.

The two classifications of MAC protocols are contention based and contention free. Contention-based MAC protocols allow for nodes to dynamically decide when to transmit. Classic examples of

contention-based MAC protocols are ALOHA [1], CSMA [9], and MACA [8]. Some MAC protocols, such as ALOHA, are unable to handle the hidden terminal problem [19]. The hidden terminal problem occurs when node B is communicable with nodes A and C, but nodes A and C cannot communicate. Node A must not communicate to B if C is simultaneously communicating with B. The challenge of the hidden terminal problem is that nodes A and C cannot coordinate transmissions to B, because they cannot communicate with each other. The MACA protocol is an example of a contention-based MAC protocol that can address the hidden terminal problem by using a RTS/CTS scheme to reserve the channel. This involves a node sending a request-to-send (RTS) and waiting to transmit until it receives a clear-to-send (CTS) notification. Waiting for a CTS is the key to avoiding the hidden terminal problem. Should a node not receive a CTS after a period of time, the RTS is repeated. Commonly, an exponential back off is used to select a random time in the future to send another RTS. This approach works well during periods of low network traffic. However, during periods of high network traffic, delay increases rapidly due to the exponential back off.

In contention-free MAC protocols, channel access is not determined dynamically but is instead reserved ahead of time. The process of reserving portions of the channel is referred to as scheduling. Schedules can be based on time (TDMA), frequency (FDMA), or code (CDMA). We use a TDMA approach. Traditionally, TDMA allocates one node per time slot. In large networks, this approach is inefficient as many nodes outside the communicable range of the transmitter are capable of transmitting. Consequently, we focus on spatial TDMA (S-TDMA), in which multiple, widely separated, nodes are allowed to transmit in the same time slot. Scheduled MAC protocols provide a guaranteed access to the channel for the node to transmit. Because of the guaranteed access to the channel, traffic experiences less variation in delay. This feature is helpful for delay sensitive traffic such as voice packets. Unfortunately, traditional TDMA approaches are inefficient when the scheduled node has no packets to forward. If a node is scheduled to transmit but has no packets to forward, then the entire time slot is wasted, and other nodes may have been able to transmit instead. Often, the performance of ad hoc networks is limited by a significantly smaller set of bottleneck nodes. If the traffic arrival rate at a bottleneck node is greater than the departure rate, then eventually queue overflow will occur, and packets will be dropped from the network. We propose an approach to TDMA scheduling to provide additional transmission opportunities to nodes, including bottlenecks, that helps improve performance.

1.2 Direct Sequence Spread Spectrum

Direct-sequence spread-spectrum (DSSS) modulation is a physical layer technique that uses a psuedorandom chipping sequence to spread the transmitter's energy over a wider portion of the frequency domain. The receiver must be correlated in time and be aware of the chipping sequence used by the transmitter. With this modulation approach links are more robust to interference which is important for ad hoc networks in which multiple access interference varies widely. Links with high signal to interference-and-noise ratios (SINR) are able to shorten the length of the chipping codes to increase the data rate. This reduces the protection from multiple-access interference but results in higher capacity links. The higher capacity robust links provided by DSSS are of particular use when trying to improve the performance of ad hoc networks [21].

1.3 MIMO

Multiple-input multiple-output (MIMO) is a physical layer technique that requires transmitters and receivers to have multiple antennas, and it leverages the spatial dimension of well-conditioned channels. Significant channel capacity increases have been shown using MIMO. For example, MIMO has been used extensively to provide increased throughput for WiFi. Spatial multiplexing is a technique that uses MIMO to create multiple independent data streams. The increase in link capacity is proportional to the number of antennas the transmitter-receiver pair have. Another application of MIMO is interference cancellation, during which a transmitter is able to null out its own signal at another receiver. This means that the receiver is available to receive a transmission from a different transmitter. Additionally, MIMO is independent of the coding and modulation technique used, so DSSS and MIMO can be combined as shown in our protocol. The NSF PAWR [3] program has recently funded experimental platforms for next generation wireless technologies. One such project is the creation of a MIMO test bed for future massive MIMO applications that could use a similar approach to our protocol outlined here.

1.4 Problem Statement

We propose a new distributed MAC protocol that uses MIMO techniques to allow a slot's unused capacity to be reclaimed by secondary transmitters that are not scheduled to transmit in the

slot. Each node has n antennas which are used for both spatial multiplexing and interference cancellation to increase the effectiveness of the unscheduled transmissions. We show that this protocol improves end-to-end completion in large ad hoc networks with varying densities.

The rest of this document is organized in the following order. Background is presented in Chapter 2. A summary of the key assumptions of our system model are provided in Chapter 3. Our new protocol for shared scheduled access is presented in Chapter 4. Simulation details are covered in Chapter 5, and results from the simulation runs are presented in Chapter 6. Lastly, Chapter 7 contains the conclusions and future work.

Chapter 2

Background

This chapter provides background and prior work for TDMA scheduling protocols and MIMO. Our new protocol adapts an existing TDMA protocol to take advantage of transceivers with multiple antennas. The existence of multiple antennas at the transmitter and receiver allows for the use of MIMO.

2.1 Scheduling

Scheduling protocols based on TDMA can be classified into either link scheduling or broadcast scheduling. Link scheduling guarantees that a transmission on a link from node A to node B will be successful. Broadcast scheduling guarantees that a broadcast from node A will be correctly decoded by each 1-neighbor. A k-neighbor is a node that is k hops from terminal A, where a hop indicates a wireless transmission between two nodes. Therefore, a 1-neighbor is one hop away from node A, and a 2-neighbor is two hops away. In order to avoid the hidden terminal problem, broadcast scheduling requires that when a node transmits none of its 1-neighbors or 2-neighbors transmit at the same time. The collective set of 1-neighbors and 2-neighbors is referred to as a node's neighborhood.

The most basic TDMA protocol for broadcast scheduling is to assign every node a single time slot in which to transmit. With N nodes in a network, a node is scheduled to transmit every N time slots. This is extremely inefficient, however, for large networks in which nodes may be separated widely enough to transmit without disrupting other ongoing transmissions. An S-TDMA protocols

take advantage of frequency reuse to allow multiple transmissions to take place at the same time. However, the optimal slot assignment for S-TDMA has been shown to be NP-complete [5].

The classic coloring problem has been been applied in several ways to implement S-TDMA. The protocols assign a color number to each node in the network such that every node has a unique color number in its neighborhood. Then, color numbers are assigned time slots instead of individual nodes. This approach improves the spatial efficiency of the basic TDMA approach. The RAND protocol [13] is an example of a centralized approach to coloring nodes for S-TDMA. The protocol was later revised as a distributed version called DRAND [15]. Lyui's protocol [10] is a distributed protocol based on the global coloring approach as well, but assigns multiple colors per slot to increase the spatial efficiency of the algorithm. We use Lyui's algorithm as a basis for our work. However, our approach works with any broadcast scheduling algorithm.

2.2 MIMO

A narrowband time-invariant wireless channel with n transmit antennas and n receive antennas is described by H, an $n \times n$ deterministic matrix. The received signal, y, is modeled by

$$y = Hx + w \tag{2.1}$$

where x is the transmitted signal, and w is white Gaussian noise. A combination of pre-processing and post-processing as shown in [20] can be used to find the correct symbols to transmit, x, that produce the desired signal, y. Therefore, when both the transmitter and receiver have multiple antennas and the channel has suitable fading conditions, then the spatial dimension can be exploited to get a degree of freedom (DoF) gain. The DoF gain allows for both spatial multiplexing and interference cancellation. Spatial multiplexing allows for n independent streams of data to be sent from antennas on one transmitter to antennas on one or more receivers. Interference cancellation is when a transmitter nulls out its own signal at another receiver. The other receiver is then able to receive a transmission unimpaired by the original transmitter.

One challenge to using MIMO systems is acquiring the channel state information (CSI) used for finding H. Inaccurate approximations of the CSI lead to significant increases in interference at the receivers. A process known as channel sounding is used to find an approximation for the CSI. One approach to measuring the CSI is to broadcast a known sequence of data to the receivers over

the link. The receiver is then able to estimate the CSI based on the received signal and send the resulting matrix, H, back to the transmitter. The performance of the MIMO system depends not only on the accuracy of the approximation, H, but also the frequency with which the channel is sounded. The frequency at which the channel should be sounded depends on the characteristics of the channel. Experiments conducted by Ma [11] have shown that, for stable channels, the CSI provides a good measure of the channel for up to 100 ms. More dynamic channels require more frequent sounding. Another concern is that as the number of antennas, n, increases, the size of the CSI matrix, H, increases like n^2. This means increased overhead in estimating and distributing the CSI as shown in [6].

Due to the complex matrix operations involved at the physical layer for MIMO, a simplified optimistic network model has been developed [2] [12] that only requires numeric computations to keep track of the number of degrees of freedom. Every node's antenna is associated with a DoF. Spatial multiplexing consumes one DoF at both the transmitter and receiver. Similarly, interference cancellation consumes one DoF at the transmitter. A node can use any combination of degrees of freedom for either spatial multiplexing or interference cancellation, but the total number of degrees of freedom used cannot exceed the number of antennas a node has. The other constraint when considering MIMO is the use of power. Each transmitter has a maximum power that can be used and is divided among the antennas. At a network layer, it is helpful to think of the DoF and power resources using budgets. To check if a MIMO transmission is feasible, it is only necessary to check that the required DoF and power resources are within the respective budgets of the nodes involved.

Chapter 3

System Model

Key assumptions of the system model are described in this chapter. The channel model describes how the signal to interference plus noise ratio (SINR) is calculated and the requirements for the transmission to be correctly decoded at a receiver. Next the scheduling algorithm used as basis for our protocol is described. Lastly, details regarding MIMO are discussed for distributing the channel state information and calculating the SINR.

3.1 Channel Model

Communication links are half duplex, meaning that nodes cannot transmit and receive at the same time. Each node uses direct-sequence spread-spectrum (DSSS) modulation, and all nodes are synchronized in time. A packet is considered correctly decoded only if the $SINR$ at the receiving node is greater than a threshold, β. Specifically, if node i is transmitting a packet to node j, then the packet can be decoded correctly only if the $SINR$ at the receiving node, j, satisfies

$$SINR_{i,j} = \frac{P_r(i,j)N_sT_c}{N_0 + \sum_{\forall k \neq i} P_r(k,j)T_c} > \beta \tag{3.1}$$

where $P_r(i,j)$ is the power received at node j from node i, N_s is the spreading factor, T_c is the chip duration, and N_0 is the noise at the receiver. The multiple-access interference at node j from all other nodes is $\sum_{\forall k \neq i} P_r(k,j)T_c$.

Signal fading is assumed to follow an urban area cellular radio path loss model [14]. If the

distance between the transmitter, i, and receiver, j, is denoted $d_{i,j}$, then the received power, denoted by $P_r(i,j)$, is given by

$$P_r(i,j) = P_t(i)(\frac{\lambda}{4\pi d_{i,j}})^{\alpha} \tag{3.2}$$

where $P_t(i)$ is the transmit power from node i, λ is the wavelength of the transmitted signal, and α is the path loss exponent.

The spreading factor, N_s, can be adapted on a per link basis depending on $SINR_{i,j}$. Using the maximum spreading factor, N_{max}, one packet per time slot can be transmitted on a link. Using $N_s = \frac{1}{2}N_{max}$, two packets can be transmitted per time slot at the expense of halving the effective $SINR_{i,j}$. The same holds true for $N_s = \frac{1}{4}N_{max}$ which is a link rate of four packets per time slot at the expense of dropping the SINR by 75%. These are the only three spreading factors employed for studies reported in this thesis.

3.2 Scheduling

Scheduling is based on Lyui's protocol, originally defined in [10] and further analyzed in [7] and [22]. Lyui's protocol improves the time slotted global coloring algorithm by relaxing the requirement that only nodes with the same color number can transmit in a slot. As in the global coloring approach, each node, i, is assigned a color number, c_i, that is unique among i's neighborhood, where the neighborhood is all nodes within 2 hops of node i. The problem of finding the minimum number of color numbers is analogous to the coloring problem [16], a classic computation problem that is known to be NP-complete. A greedy coloring algorithm is used because it can be implemented in a distributed manner. Each node attempts to claim the smallest available color number in its neighborhood, and ties are broken using the node's unique ID number. This approach provides near optimal coloring when the nodes are randomly distributed in the environment. Figure 3.1 shows a possible coloring using this distributed approach for 100 nodes.

Lyui's algorithm specifies which colors are allowed to transmit in a time slot. A node, i, with color number, c_i, is a candidate to transmit in time slot t if there exists an integer, n, such that

$$t = c_i + nP_G(c_i) \tag{3.3}$$

where $P_G(c_i) = 2^k$, and k is the smallest integer such that $2^k \geq c_i$. Table 3.1 shows when the first

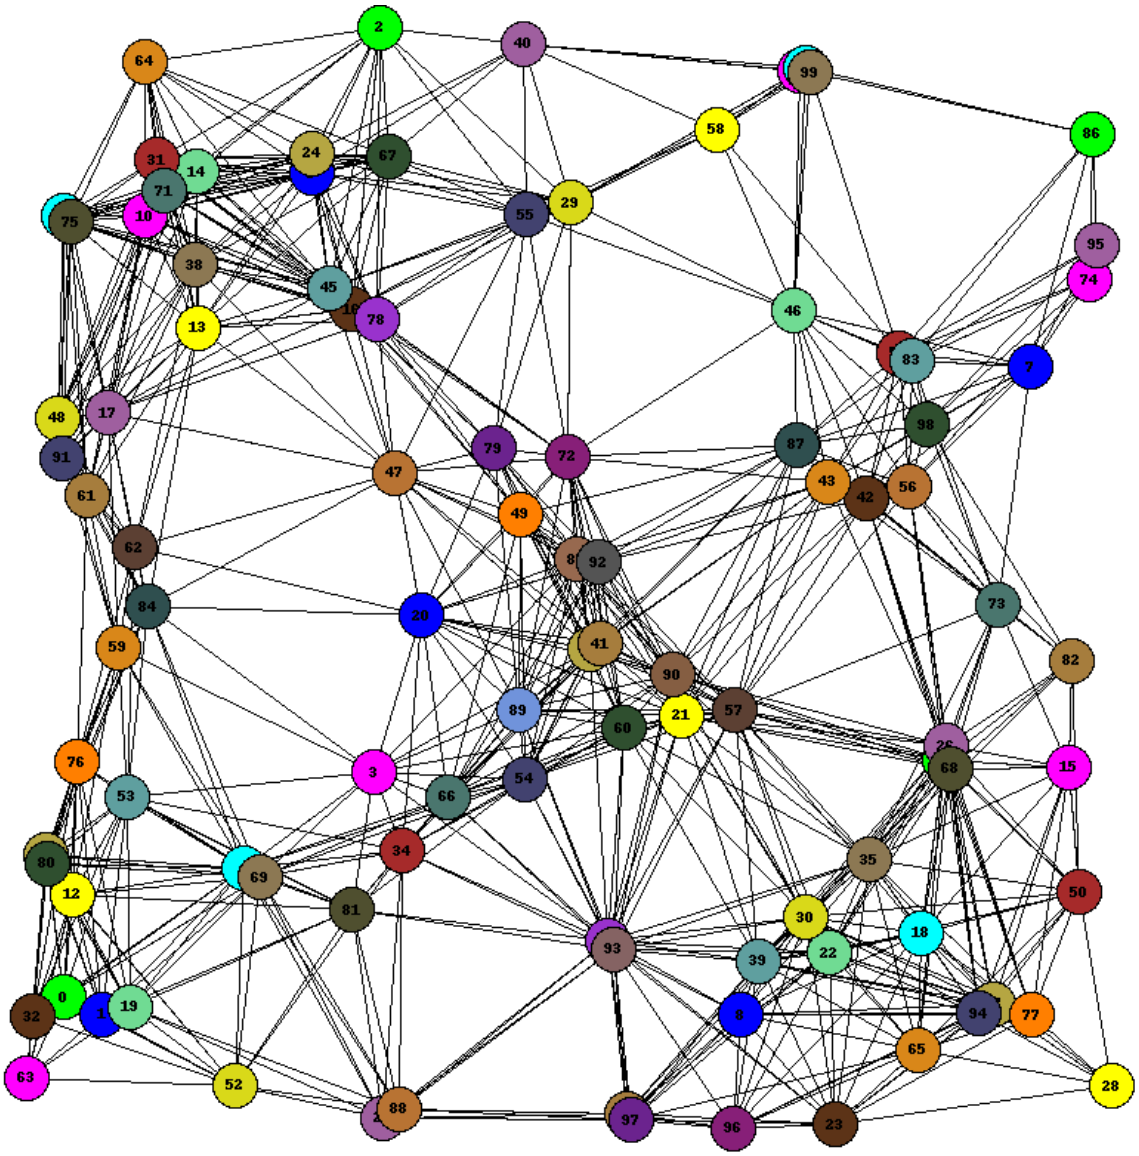

Figure 3.1: Sample network coloring for 100 randomly located nodes

eight color numbers are candidates for transmission according to Equation (3.3).

In each time slot, t, node i forms a candidate set, C, such that $\forall j \in C$ j is either equal to node i or in i's neighborhood, and Equation (3.3) is satisfied. Table 3.1 shows the candidate time slots for the first eight color numbers. Node i will only be assigned to transmit in time slot t if it is a member of the candidate set and has the largest color number among the other candidates. Lyui's algorithm guarantees that each node will be able to transmit at least once per frame. The frame size, f, is dictated by $f = P_G(c_{max})$, where c_{max} is the largest color number in node i's neighborhood.

Color number c_i \ time slot t	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
1	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
2		✓		✓		✓		✓		✓		✓		✓		✓
3			✓				✓				✓				✓	
4				✓				✓				✓				✓
5					✓								✓			
6						✓								✓		
7							✓								✓	
8								✓								✓

Table 3.1: Candidates for transmission, by color number, for Lyui's algorithm

3.3 MIMO

Each node in the network has n antennas, and the channel state information(CSI), H, for each 1-neighbor. The channel state information is distributed using the fact that every node is scheduled at least once per frame. Consider node i with frame size f that is scheduled to transmit in time slot t. Node i transmits a sequence of pilot symbols at the start of time slot t, allowing each 1-neighbor to calculate H. Node i is guaranteed to be scheduled to transmit again in time slot $t + f$. Meanwhile, every 1-neighbor of node i is guaranteed to be scheduled to transmit at least once, allowing each neighbor to respond to node i with the corresponding CSI. Then, beginning in time slot $t + f$, node i is able to begin transmitting using MIMO. The process of learning the CSI is repeated every frame. This approach means that the CSI used will be at least f time slots old. However, as shown in [11], it is possible to use the same CSI for an extended period of time.

Figure 3.2 shows how the beginning of the time slot is divided to broadcast CSI information. At the beginning of the time slot, every scheduled node transmits a sequence of pilot symbols. These pilot symbols allow each 1-neighbor to calculate the CSI matrix, H, for node i. The next portion of the time slot is used by the node to broadcast the CSI matrices calculated for all of its k neighbors over the previous frame.

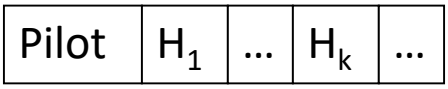

Figure 3.2: Time slot division for CSI

The use of MIMO complicates the SINR calculation in Equation 3.1 slightly. If the transmis-

sion from node i to node j is using MIMO for either spatial multiplexing or interference cancellation, then the multiple-access interference from the other transmissions are included in the MIMO pre-processing associated with the receive antenna at node j. Formally this is expressed as

$$SINR_{i,j} = \frac{P_r(i,j)N_sT_c}{N_0 + \sum_{\forall k \in T_i} P_r(k,j)T_c} > \beta \tag{3.4}$$

where T_i is the set of all transmissions in the current time slot that did not originate at i and were included in the MIMO pre-processing to be canceled at the receive antenna at node j. For example, if node i uses two antennas to transmit to node j, then the interference from one antenna is not included in the multiple-access interference calculation for the other antenna at j. This is because we assume perfect cancellation of signals at the receiver.

Chapter 4

Scheduled Channel Access for MIMO

This chapter describes our new protocol for transmission scheduling using radios with multiple antennas. We leverage interference cancellation to allow multiple nodes to transmit. This protocol is layered on top of a broadcast scheduling protocol. We define a primary transmitter as a node scheduled for the time slot in question. A secondary transmitter is a node not scheduled for the current time slot, but that is allowed to transmit. Our protocol allows one or more secondary transmitters to transmit when a primary transmitter is unable to completely utilize a slot. First, an overview of how we combine spatial multiplexing and interference cancellation with a broadcast scheduling algorithm is described. Second, the details of how to pick secondary transmitters without violating the assumptions of the broadcast scheduling algorithm is described. Next, the specific algorithm for greedily assigning resources for a primary transmitter is covered, and then the algorithm for assigning remaining resources to secondary transmitters is described. Lastly, an example is discussed to clarify how the two algorithms work together.

4.1 Overview

Every node acts as the primary transmitter in time slots when it is scheduled according to Lyui's protocol. Each primary transmitter forms an ordered set of secondary transmitters. Let the

primary transmitter have m secondary transmitters where the size of the ordered set is less than or equal to m_{max}. In a time slot in which the primary transmitter is unable to completely use the slot capacity, secondary transmitters are allowed to transmit. This allows time slots that would otherwise be wasted to be reclaimed. The primary transmitter uses a greedy approach to try to fill the capacity of the slot. The primary node transmits as many packets from its own queue as it can, using both spatial multiplexing and adaptive spreading. The remainder of the slot capacity is given to the set of secondary transmitters which use interference cancellation to send additional packets. The coordination between the primary transmitter and secondary transmitters is done during a short sequence of messages at the start of the slot.

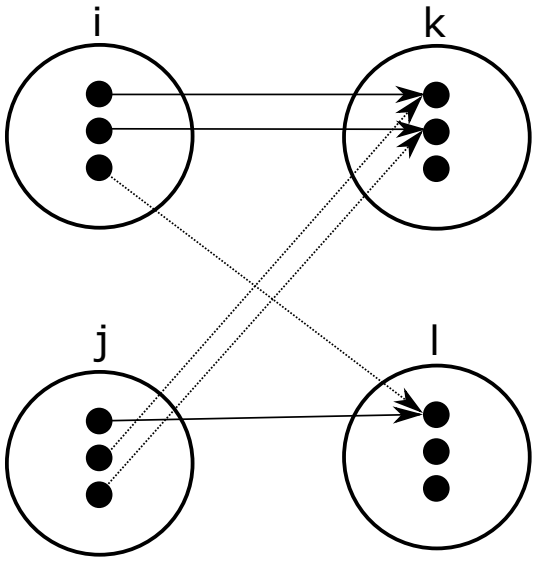

Figure 4.1: Example one slot assignment for MIMO

An example time slot is shown in Figure 4.1. Each node has three antennas allowing for each node to be involved in up to three transmissions. Let i be the primary transmitter in time slot t with secondary transmitter j. In this time slot example, node i has packets in its queue for two transmissions to node k. Node i has not used the entire capacity of the time slot and has no packets remaining in its queue that do not require more resources than are available. The secondary transmitter, j, is then allowed to send a transmission to node l. These three data transmissions are indicated by the solid arrows in Figure 4.1. However, this requires node i canceling out the interference from its two transmissions at node l, and node j canceling out the interference from its transmission at both antennas at node k. The interference cancellation is indicated by the dashed

arrows in Figure 4.1. It is not necessary for each data transmission to only contain a single packet. Adaptive spreading may be used to pack additional packets into each transmission.

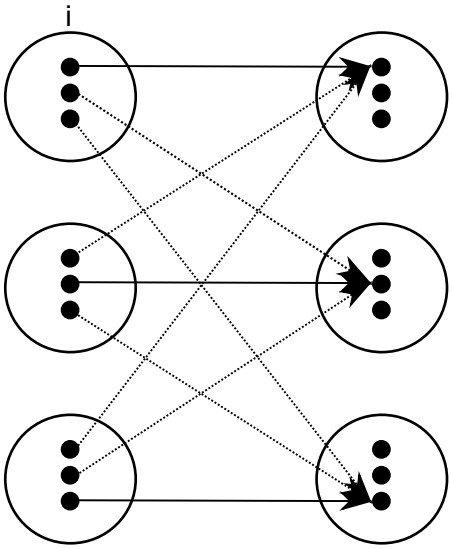

Figure 4.2: Example two slot assignment for MIMO with maximum number of transmitters

Three constraints on two resources control which packets can be transmitted from the queues of the primary and secondary transmitters. The two resources are the degrees of freedom (DoF) and available power for each transmitter. Each node has n antennas, each of which corresponds to a degree of freedom. Additionally, each node has a maximum total transmit power of P_{max}. The final constraint is to limit the total output power of the primary and secondary transmitters to be less than or equal to P_{max}. Limiting the total output power helps prevent drastically changing the multiple-access interference environment. These constraints are monitored using three budgets, B_{DoF}, B_P, and B_{TP}. The degrees of freedom (DoF) budget, B_{DoF}, is monitored on a per node basis and monitors the number of antennas available to use. Thus, the DoF budget is capped at n. The power budget, B_P, is also monitored on a per node basis and monitors the amount of power available for transmitting. For simplicity, the power budget is divided into a finite number of units, and it is assumed that an integer number of units is always used. Figure 4.2 depicts a situation in which the maximum number, n, of unique transmitters occurs. Each node must use its $n-1$ other antennas for interference cancellation. This results in a total of n^2 transmissions. In order to allow the possibility of n parallel transmissions, the power budget must have at least n^2 units. Again for

simplicity, it is assumed that the the power budget is capped at n^2 units. The total power budget, B_{TP}, tracks the total output power of the primary and secondary transmitters. The total power budget is capped at $B_{TP} = n^2$ to prevent creating excessive additional interference.

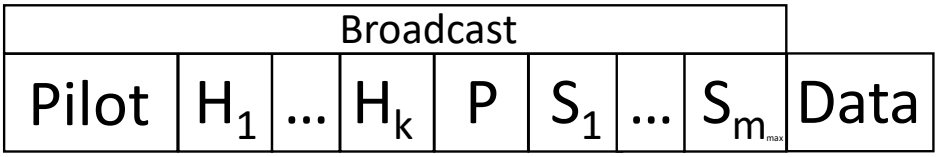

Figure 4.3: Time slot division for our protocol

Figure 4.3 illustrates how the time slot is divided. During the broadcast phase of the time slot, the channel state information is shared with neighbors as described in Chapter 3. This is followed by the coordination of the primary and secondary transmitters. First, the primary announces its intended transmissions. The secondary transmitters then take turns announcing their own intended transmissions with the remaining slot capacity. After the coordination phase is over, the rest of the time slot is spent transmitting data.

4.2 Picking Secondary Transmitters

The method of selecting secondary transmitters is based on a time division approach to sharing slots with secondary transmitters from [17] and [18]. Let node i be the primary transmitter of time slot t, and node j be a 1-neighbor of i. Node j is a possible secondary transmitter of node i if

$$SINR_{i,j} \geq p\beta \tag{4.1}$$

where p is a constant greater than one. The parameter, p, constrains the secondary transmitters to limit creating additional multiple-access interference.

Each 1-neighbor of i is evaluated using Equation (4.1) to produce a candidate set of secondary transmitters. The candidate set is then filtered into the ordered set of secondary transmitters by randomly picking nodes from the candidate set until either the candidate set is exhausted or the

secondary transmitter set has reached its maximum size, m_{max}. As secondary transmitters are randomly selected, the candidate set is further reduced to ensure that each candidate is communicable with every secondary transmitter. Each secondary transmitter must be communicable with every other secondary transmitter so that the initial coordination of primary and secondary transmitters can occur. Picking secondary transmitters randomly from the candidate set helps reduce correlation in the secondary transmitters of neighboring nodes.

4.3 Primary Transmitter Budget Allocation

We use a greedy approach to try to allocate as much of the budget as possible to the primary transmitter because there is no guarantee that secondary transmitters will have packets in their own queues. The goal is to forward as many packets from node i's queue to its next hop as possible without wasting any of the slot capacity. A transmitter can use both adaptive spreading and spatial multiplexing to increase the throughput of high capacity links. Adaptive spreading involves choosing a spreading factor, N_s, that allows one, two, or four packets to be transmitted to the same next hop, j, at the expense of lowering the $SINR_{i,j}$. Spatial multiplexing involves allocating portions of the DoF budget and power budgets to allow up to n independent data flows that do not have to be to the same next hop. As in adaptive spreading, spatial multiplexing also involves dropping the $SINR_{i,j}$ to achieve the performance gain. Each data flow created by spatial multiplexing is also capable of supporting adapting spreading if the SINR on the link is sufficiently high. The challenge is finding the combination of adaptive spreading and spatial multiplexing while working within the given budgets.

At the start of a time slot, each node has its DoF budget, B_{DoF}, initialized to n, and its power budget, B_P, initialized to n^2. During the primary transmitter allocation phase, it is not necessary to consider the total output power budget, because it is equivalent to the individual power budget. The primary transmitter begins searching at the beginning of its queue. Let the next hop for this packet be over link $l_{i,j}$ with associated $SINR_{i,j}$. The primary transmitter then counts the number of packets, k, in its queue with the same next hop, j.

The spreading factor, N_s, is chosen first by finding the minimum value of N_s such that

$$k \geq R(N_s) \tag{4.2}$$

where $R(N_s) = \frac{N_{max}}{N_s}$ is the associated transmission rate. This indicates that no portion of the transmission in the time slot is being wasted. This can be done by first calculating the minimum spreading factor, N_{min}, given $SINR_{i,j}$ and the largest power of two less than or equal to k, $P_L(k)$. The minimum spreading factor that satisfies Equation 5.6 is given by

$$N_s = \frac{N_{max}}{min(N_{min}, P_L(k))} \tag{4.3}$$

For example, consider a link that is capable of transmitting four packets using the spreading factor $\frac{1}{4}N_{max}$. If there are only two packets in i's queue with j as a next hop, then the spreading factor $\frac{1}{2}N_{max}$ is chosen. If instead there are three packets in i's queue with j as the next hop, then the same spreading factor is still be used to transmit two of the packets. This is done to maximize the use of the slot.

The next step is to determine how much of the DoF budget and power budget this transmission will require. Every transmission by the primary transmitter requires one antenna which is equivalent to one degree of freedom from the budget. It is also necessary to determine the power, P, required for the transmission to be correctly decoded by node j given the new effective SINR, $\frac{N_s}{N_{max}}SINR_{i,j}$. The DoF budget is then decremented by one, and the power budget is decremented by P. The first $R(N_s)$ packets in i's queue with next hop, j, are then removed from the queue and prepared for transmission.

If either the DoF budget or the power budget have been exhausted at this point, then node i may transmit as allocated. Otherwise, the process begins again with the next packet in the queue. Initial values of N_s and P are computed in the same way, but now it is necessary to make sure their is room in the respective budgets. The budget check is indicated by the expressions 4.4 and 4.5.

$$B_{DoF} > 0 \tag{4.4}$$

$$B_P - P \geq 0 \tag{4.5}$$

Specifically, this means that the DoF budget must be greater than zero, and the power budget minus the required power, P, be greater than or equal to zero. If these conditions hold, then the first $R(N_s)$ packets in i's queue with next hop j are then removed from the queue and prepared for

transmission. If the conditions do not hold, and the spreading factor is not at its maximum value, then the spreading factor can be increased by a factor of two. The new spreading factor halves the transmission rate and decreases the required power P. The process of iteratively increasing the spreading factor can be repeated until the required resources are within the budget or the spreading factor reaches N_{max}. If the required resources are now within the budget, then the first $R(N_s)$ packets in i's queue with next hop, j, are removed from the queue and prepared for transmission. Otherwise no additional transmissions to j are feasible.

The iterative process of selecting N_s and P is repeated until every packet remaining in i's queue has been examined or either budget is exhausted. At this point, if either the DoF budget or the power budget have been exhausted, then node i may transmit as allocated. Otherwise, there is remaining slot capacity which is offered to the ordered set of secondary transmitters. Algorithm 1 provides a concise description of the primary transmitter budget allocation.

Algorithm 1 Primary Transmitter Budget Allocation

1: Instantiate budgets $B_P = n^2$ and $B_{DoF} = n$.
2: The packet at the front of node i's queue is selected, assume it's next hop is link $l_{i,j}$ with the associated $SINR_{i,j}$.
3: Count the number of packets, k, in the queue with the same next hop destination j.
4: Find the minimum spreading factor, N_s, such that $k \geq R(N_s)$, where $R(N_s)$ is the link rate given $SINR_{i,j}$.
5: Calculate minimum power, P, required for the transmission from i to j given N_s.
6: **if** $B_P - P \geq 0$ and $B_{DoF} > 0$ **then**
7: $\quad B_P \leftarrow B_P - P$
8: $\quad B_{DoF} \leftarrow B_{DoF} - 1$
9: $\quad$ Allocate first $R(N_s)$ packets from i's queue with next hop j for transmission.
10: **else if** $N_s < N_{max}$ **then**
11: $\quad N_s \leftarrow 2N_s$
12: $\quad$ **goto** step 5
13:
14: **if** $B_P = 0$ **then**
15: $\quad$ No more transmissions in time slot can occur.
16: **else if** $B_{DoF} = 0$ **then**
17: $\quad$ No more transmissions in time slot can occur.
18: **else if** i is last packet in queue **then**
19: $\quad$ Continue with secondary transmitter budget allocation.
20: **else**
21: $\quad$ Consider the next packet in i's queue and the new next hop j and **goto** step 3.

4.4 Secondary Transmitter Budget Allocation

If neither budget is filled after the primary transmitter allocation phase, the remaining slot capacity is offered to the ordered set of m secondary transmitters. The first secondary transmitter forwards as many packets from its own queue as possible, using both adaptive spreading and spatial multiplexing. The next secondary transmitter is then offered any remaining slot capacity. This is continued until the slot capacity is exhausted, or all m nodes have been offered the remaining capacity.

The algorithm for allocating a secondary transmitter's budget has two main differences from allocating a primary transmitter's budget. The first difference is the introduction of a new constraint. The total output power of the primary and secondary transmitters is limited to P_{max}, the maximum output power of a single transmitter. This constraint helps to prevent detrimental changes to the interference environment that would break the underlying assumptions of the scheduling algorithm. This constraint is monitored using an additional power budget that measures the total output power of the participating nodes thus far. The second difference is that each secondary transmission must use interference cancellation to every participating receiver. Let T be the set of all previously allocated transmit antennas, and R be the set of all previously allocated receive antennas. In order for a secondary transmitter, i, to forward packets to node j, i must also transmit an interference cancellation signal to every antenna, a_R, in R unless the antenna belongs to j and the transmission to a_R is also from i. In other words, if the antenna a_R is being used for spatial multiplexing from the same source antenna, then i does not need to cancel interference at antenna a_R. Additionally, every transmitter other than i with an antenna in T must transmit an interference cancellation signal to j.

As with the primary transmitter, every secondary transmitter is initialized with a full DoF budget and power budget at the beginning of the time slot. The total power budget is initialized by the first secondary transmitter to the power used by the primary transmitter. The total power budget is then shared by each subsequent secondary transmitter. The sets of participating transmit antennas, T, and participating receive antennas, R, are also inherited from the primary transmitter and shared with subsequent secondary transmitters. Each secondary transmitter begins searching at the beginning of its queue. Let the next hop for this packet be link $l_{i,j}$ with associated $SINR_{i,j}$. Links are half duplex, so node i must not have any antennas in the set R. If so, then the current

secondary transmitter cannot transmit in the time slot so consider the next secondary transmitter. Additionally, because links are half duplex node j must not already have an antenna in the set, T. If so, then consider the next packet in i's queue. Count the number of packets, k, in the secondary transmitter's queue with the same next hop, j.

The spreading factor, N_s, is chosen first by finding the minimum value of N_s such that equation 4.2 holds. Given N_s and $SINR_{i,j}$ the required power, P, for the secondary data transmission can be calculated. These calculations are the same as in the primary transmitter allocation phase. Consider the set T_i, a subset of T, that is only missing allocated antennas from i, and the set $R_{i,j}$, which is a subset of R and only missing receive antennas used for the link $l_{i,j}$. An additional degree of freedom for every antenna in the set $R_{i,j}$ is also needed, so the degree of freedom budget check is given by Equation 4.6. Additionally, each interference cancellation from i to r requires power, P_r, for each receiver r in the set $R_{i,j}$. The power budget check is given by Equation 4.7. Additionally, each transmitter, t, in T_i must satisfy Equations 4.4 and 4.5 for the interference cancellation from t to j. The total power budget check is given by Equation 4.8.

$$B_{DoF} - |R_{i,j}| > 0 \tag{4.6}$$

$$B_P - P - \sum_{\forall r \in R_{i,j}} P_r \geq 0 \tag{4.7}$$

$$B_{TP} - P - \sum_{\forall r \in R_{i,j}} P_r - \sum_{\forall t \in T_i} P_t \geq 0 \tag{4.8}$$

If the conditions do not hold and the spreading factor is not at its maximum value, then the spreading factor can be increased by a factor of two. The new spreading factor, N_s, halves the transmission rate and the required power P. The process of iteratively increasing the spreading factor can be repeated until the required resources are within the budgets or the spreading factor reaches N_{max}. If the required resources are within the budget, then the first $R(N_s)$ packets in i's queue with next hop, j, are then removed from the queue and prepared for transmission. Otherwise, no additional transmissions to j are feasible.

The iterative process of selecting N_s and P is repeated until every packet remaining in i's queue has been examined or either budget is exhausted. At this point, if the DoF budget or

either power budget have been exhausted, then node i may transmit as allocated. Otherwise, there is remaining slot capacity which is offered to the remaining secondary transmitters. A concise description of this process is given in algorithm 2.

Algorithm 2 Secondary Transmitter Budget Allocation

1: Instantiate budgets $B_P = n^2$, $B_{DoF} = n$.
2: Inherit budget $B_{TP} = B_P(primary)$, the set of transmit antennas, T, and the set of receive antennas, R, from the primary transmitter.
3: The packet at the front of node i's queue is selected, assume it's next hop is link $l_{i,j}$ with the associated $SINR_{i,j}$.
4: Calculate subsets T_i and $R_{i,j}$.
5: **if** $i \in R$ **then**
6: Allow the next secondary transmitter to use the remaining slot capacity.
7: **if** $j \in T$ **then**
8: Consider the next packet in i's queue and the new next hop j and **goto** step 4.
9: Calculate subsets T_i and $R_{i,j}$.
10: Count the number of packets, k, in the queue with the same next hop destination j.
11: Find the minimum spreading factor, N_s, such that $k \geq R_{i,j}$, where $R_{i,j}$ is the link rate given $SINR_{i,j}$.
12: Calculate minimum power necessary, P, for the transmission from i to j given N_s.
13: **if** $B_{DoF} - |R_{i,j}| > 0$ and $B_P - P - \sum_{\forall r \in R_{i,j}} P_r \geq 0$ and $B_{TP} - P - \sum_{\forall r \in R_{i,j}} P_r - \sum_{\forall t \in T_i} P_t \geq 0$ **then**
14: $B_P \leftarrow B_P - P - \sum_{\forall r \in R_{i,j}} P_r$
15: $B_{TP} \leftarrow B_{TP} - P - \sum_{\forall r \in R_{i,j}} P_r - \sum_{\forall t \in T_i} P_t$
16: $B_{DoF} \leftarrow B_{DoF} - |R_{i,j}| - 1$
17: Allocate first $R_{i,j}$ packets from i's queue with next hop j for transmission.
18:
19: **if** $B_P = 0$ **then**
20: No more transmissions in time slot can occur.
21: **else if** $B_{DoF} = 0$ **then**
22: No more transmissions in time slot can occur.
23: **else if** $B_{TP} = 0$ **then**
24: No more transmissions in time slot can occur.
25: **else if** i is last packet in queue **then**
26: Allow the next secondary transmitter to use the remaining slot capacity.
27: **else**
28: Consider the next packet in i's queue and the new next hop j and **goto** step 4.

4.5 Example

This section provides an example time slot for both the primary and secondary transmitters. Consider Figure 4.4 in which node i is the primary transmitter of the current time slot and node k is a secondary transmitter. In this example, node m is elsewhere in the network and not communicable

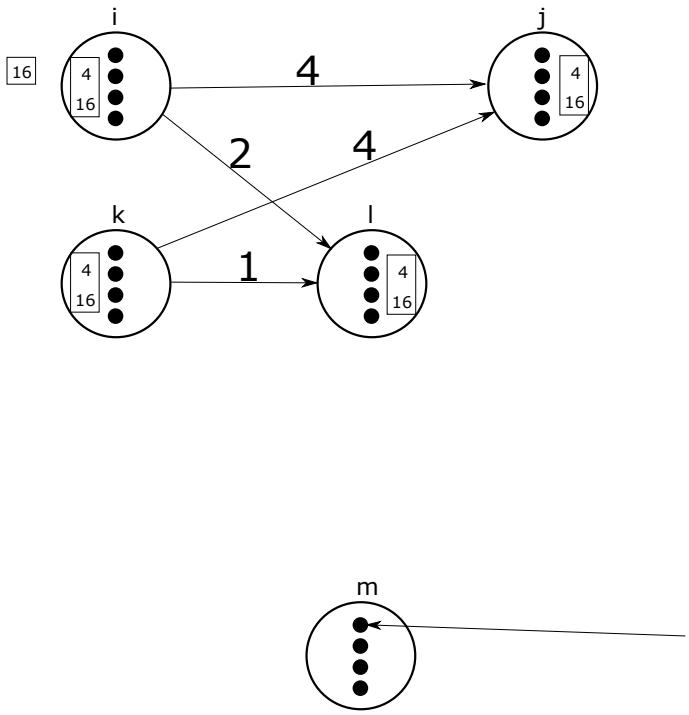

Figure 4.4: Example three slot assignment for MIMO setup

with i but happens to be receiving a packet in the same time slot from another node. The numbers on the link indicate the minimum units of power required for a transmission to be correctly decoded at the receiver. Each node has four antennas indicating that the DoF budget is limited to four, and the power budget is limited to 16. The DoF and power budget for each node at the start of the time slot is shown for each node. The total output power budget is also included in Figure 4.4. A simple example is to assume that node i has four packets in its queue with next hop j. Using the spreading factor $\frac{1}{4}N_{max}$ all four packets could be packed onto the link consuming all 16 units of power and one degree of freedom. Consider instead a slightly more complicated scenario in which node i has one packet in its queue and the next hop of the packet is j, and let node k have three packets in its queue, all of which have the next hop l.

Node i is the primary transmitter so it will attempt to transmit as many packets from its queue as possible. However, in this scenario there is only one packet in the queue. The transmission to j requires four units of power and one degree of freedom. Figure 4.5 illustrates the state of the budgets after the resources for the first transmission have been allocated.

Node i has no remaining packets in its queue, so the first secondary transmitter, k, is offered

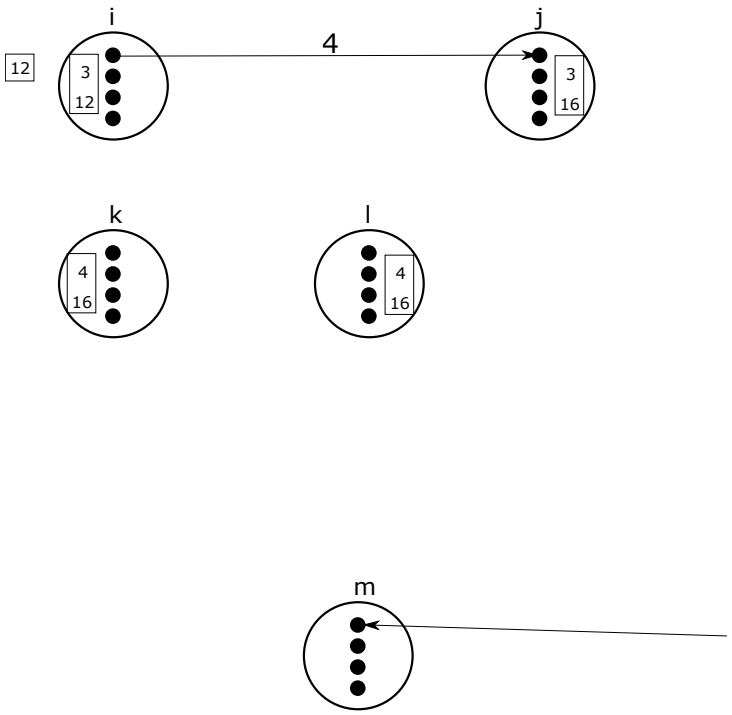

Figure 4.5: Example three slot assignment for MIMO part one

the remaining slot capacity. All three packets in k's queue have next hop l. The SINR on the link from k to l is also good enough to support the spreading factor $\frac{1}{4}N_{max}$, which transmits four packets. However, with only three packets with next hop l the spreading factor $\frac{1}{2}N_{max}$ is instead chosen to transmit two packets. Halving the spreading rate means the required power is increased by a factor of two to two units. The transmission from k to l requires both node i and k to transmit interference cancellation signals to the intended receive antennas. Figure 4.6 illustrates the resource requirements for allocating the second transmission. Both node i and k only require two degrees of freedom and six power units for a grand total of 12 power units. The power and DoF budgets have not been consumed, each node has two DoF remaining and four power units for the total power budget.

The secondary transmitter, k, still has one packet in its queue for l that may be able to be transmitted. If the maximum spreading factor is used for this transmission from k to l, then only one unit of power is required Figure 4.7. Node i must use interference cancellation at the new receive antenna. However, because node k is using spatial multiplexing, it does not need to transmit any additional interference cancellation. At this point, both nodes have consumed three degrees of freedom and a total of 15 units of power. The capacity of the slot has not been completely consumed,

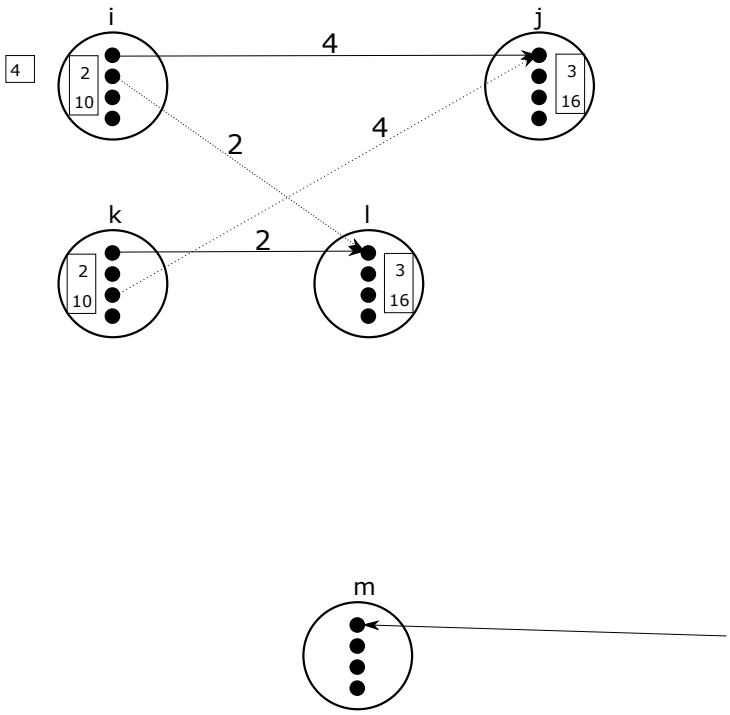

Figure 4.6: Example three slot assignment for MIMO part two

but there are no more secondary transmitters to offer the extra capacity, and the remaining unit of power cannot be further subdivided. It is important to note that all three transmitting antennas are creating interference at node m. The broadcast scheduling algorithm determined it was safe for i, not k, to transmit when m can receive. The extra multiple-access interference, caused by node k transmitting, could be enough to prevent m from correctly decoding its packet. Our protocol includes two constraints to limit the effect of multiple access interference in the event of this situation occurring. First, we constrain the secondary transmitters to have high SINR links with the primary transmitter in hopes that the produced multiple-access interference is similar to the interference produced by the primary node. Second, we limit the power available to the secondary transmitter by limiting the total output power of all the transmitters.

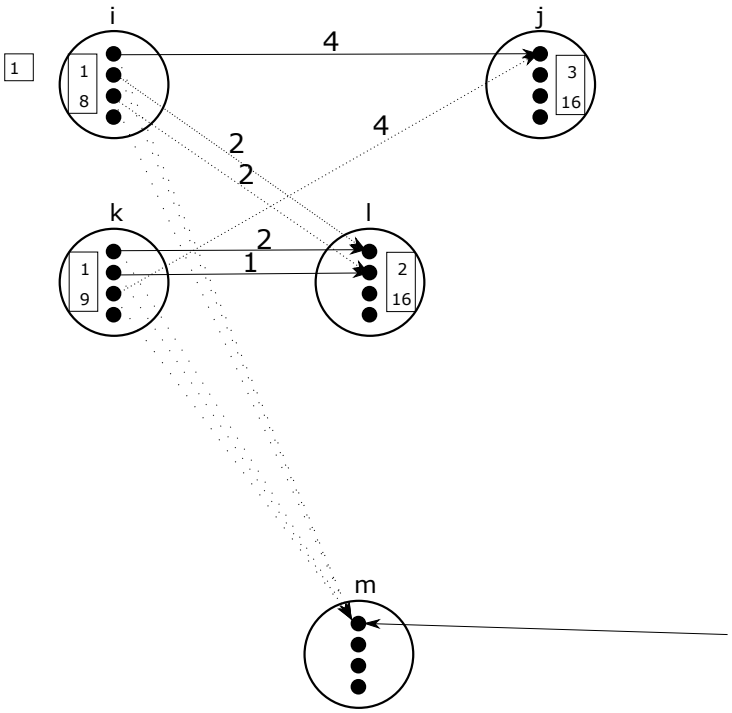

Figure 4.7: Example three slot assignment for MIMO part three

Chapter 5

Simulation

A custom simulation was developed in C++ to test our new protocol. The simulation operates in a time slotted manner and uses the physical layer and link layer techniques discussed previously. The passing of messages for sounding the channel and coordinating the secondary transmitters are not simulated. This chapter discusses the details relevant to the simulation, including routing and traffic generation.

5.1 Channel

The parameters used for the channel model are given in Table 5.1, where R denotes the the maximum distance between a transmitter and receiver such that Equation (3.1) holds, assuming the use of the largest spreading factor, N_{max}, the maximum transmit power, P_{max}, and no multiple-access interference. Given R, we define

$$P_{max} = (\frac{4\pi R}{\lambda})^{\alpha} \frac{\beta N_0}{T_c N_{max}} \tag{5.1}$$

where P_{max} is the maximum transmit power of each node.

When node i receives a packet from node j, the instantaneous SINR, S_{est}, is measured. It is assumed that the link is approximately symmetric, so the measured incoming SINR from node j is used as an estimate for the outgoing SINR to node j. However, the received SINR depends on the spreading factor, N_s, and transmit power, P_t, so the estimated SINR must be normalized before

λ	0.125 m
β	8
N_{max}	96
T_c	$2.9e-7$ s/Chip
α	3.5
N_0	$4.0e-21$ J/Hz
R	200 m

Table 5.1: Channel model parameters

being recorded. Equation 5.2 is used to normalize the estimated SINR, S_{est} to S_{norm}.

$$S_{norm} = \frac{N_{max}P_{max}}{N_s P_t} S_{est} \tag{5.2}$$

In order to calculate lower bounds on the expected outgoing SINR on the link an exponential weighted moving average (EWMA) is used to track the average and variance of the normalized SINR samples. The formulas for calculating the new average, S_n, and variance, V_n, given a new sample, S_{norm}, are shown below. The smoothing factor used is $\alpha = 0.15$, and the variance, V_0, is initialized to zero. For the purposes of the simulation, the average S_0 is initialized to the SINR estimate assuming maximum spreading factors and transmit power as well as no multiple-access interference. The standard deviation of the link is given by $\sigma_n = \sqrt{V_n}$. An estimated lower bound of the outgoing SINR on the link, $l_{i,j}$, is given by $S_n - 3\sigma_n$.

$$\delta_n = S_{norm} - S_{n-1} \tag{5.3}$$

$$S_n = S_{n-1} + \alpha\delta_n \tag{5.4}$$

$$V_n = (1-\alpha) \times (V_{n-1} + \alpha\delta_n^2) \tag{5.5}$$

Adapting the spreading factor, N_s, allows the data rate of a link to be increased at the expense of decreasing the SINR of the received signal. The exact multiple access interference cannot be anticipated, so the link SINR estimate is used to pick the value of N_s. To try to guarantee that the message will be correctly decoded at the receiver, the following condition is used $\frac{N_s}{N_{max}}(S-3\sigma) \geq f\beta$ where S and σ are the current estimates of the SINR and standard deviation, respectively, and f acts

as a buffer in case the multiple-access interference is more than was estimated. For the simulation, a buffer of $f = 1.5$ is used. Solving for S yields Equation 5.6 which is used to pick the spreading factor, N_s, as shown in Table 5.2.

$$S \geq \frac{N_{max}}{N_s} f\beta + 3\sigma \tag{5.6}$$

Link SINR Estimate	Link Spreading Factor	Link Rate
$\beta < S < 3\beta + 3\sigma$	N_{max}	1 packet per slot
$3\beta + 3\sigma \leq S < 6\beta + 3\sigma$	$\frac{N_{max}}{2}$	2 packets per slot
$6\beta + 3\sigma < S$	$\frac{N_{max}}{4}$	4 packets per slot

Table 5.2: Spreading factors

5.2 MIMO

The power budget is distributed into n^2 units with the total transmit power P_{max}. Let S be the estimated SINR and σ be the estimated standard deviation if the full transmit power (n^2 units) is used to transmit. If only a units of power are used, the effective SINR is $\frac{a}{n^2}S$. then, the minimum units of power to ensure that the SINR is acceptable is

$$a = \lceil \frac{n^2 f\beta}{S - 3\sigma} \rceil \tag{5.7}$$

where as explained before $S - 3\sigma$ is an expected lower bound of the SINR, and f provides a buffer to ensure the SINR is not reduced directly to β. For the purposes of the simulation, the value $f = 1.5$ is used. The same calculation for required power is used for data transmissions and spatial multiplexing.

In the simulation, by default, all nodes are given four antennas, the candidate secondary transmitter set constraint is ten, and the maximum number of secondary transmitters is five. These default parameters are given in Table 5.3.

n	4
p	10
m_{max}	5

Table 5.3: MIMO parameters

5.3 Routing

Packets are routed from the source nodes to the destination nodes using forwarding tables generated from Dijkstra's algorithm [4]. Routes are calculated every 1000 time slots. The metric used to calculate link weights from node i to node j is based on [21] and given by

$$w_{i,j} = \frac{\phi(S_{i,j})(1+U_j)}{ETR(i)R(i,j)}. \tag{5.8}$$

The term $\phi(S_{i,j})$ is a scaling function used to de-emphasize routing over low SINR links, and $S_{i,j}$ is the current SINR estimate over the link from i to j. This term is defined as:

$$\phi(x) = \begin{cases} \infty & x \leq \beta \\ 1 - \ln(\frac{x-\beta}{\beta}) & \beta < x \leq 2\beta \\ 1 & x > 2\beta \end{cases} \tag{5.9}$$

The term U_j approximates the utilization of node j by using an exponential weighted moving average (EWMA), and the fraction of slots assigned to j, in which j transmits a packet. Every time j is a scheduled as the primary transmitter for the slot t, j updates its utilization estimate as follows:

$$U_j' = (0.95)U_j + (0.05)T(j) \tag{5.10}$$

where the function $T(j) = 1$ if j transmits a packet in the current time slot, and $T(j) = 0$ otherwise. This term attempts to spread traffic through the network to prevent nodes from becoming bottlenecks.

The term $ETR(i)$ is the effective transmit rate of node i, which is calculated as the number of slots in each frame that node i is scheduled as the primary transmitter divided by the frame length.

The final term, $R(i,j)$, is an approximation of the link rate from node i to node j. A simple metric is calculated to approximate the link rate. The calculation begins by finding the minimum power, a, from Equation (5.7) for the transmission to be correctly decoded at the receiver. Given the new effective SINR $\frac{a}{n^2}S_{i,j}$, the minimum spreading factor, N_s is chosen. The rate supported on a link using the minimum power is given by $\frac{N_{max}}{N_s}$. The next calculation is to see how many of the

minimum power transmissions can be supported via spatial multiplexing. At most, there can be n parallel transmissions. Equation (5.11) gives the number of parallel links, k. Note that k can take values 0 through n.

$$k_{i,j} = \begin{cases} n & a < n \\ \lfloor \frac{n^2}{a} \rfloor & a \geq n \end{cases} \tag{5.11}$$

The resulting link rate is the product of the number of parallel links, k, and the data rate achievable on each of these links as shown in Equation 5.12.

$$R(i,j) = k_{i,j} \frac{N_{max}}{N_s} \tag{5.12}$$

5.4 Traffic Generation

A global traffic generation rate is specified as G for a network of N nodes. Each node then generates traffic at an average rate of $\frac{G}{N}$ using a Bernoulli process, letting $p = \frac{G}{N}$. When a packet is generated, the destination is uniformly distributed across the remaining $N - 1$ nodes.

Acknowledgments are not sent for packets that reach the final destination. Packets dropped by the network for any reason are not retransmitted. There are three possible reasons for a packet being forwarded from node i to node j to be dropped. First, a packet may be dropped due to queue overflow at j, if j is not the final destination, and the queue size of j is equal to Q_{max}. Second, if the $SINR_{i,j} \leq \beta$, the packet fails to be decoded by receiver j. Third, if node j has no entry in its routing table for the next hop, then the packet is dropped.

5.5 Simulation Parameters

Simulations consist of two parts, a warm-up phase and steady-state phase. The warm-up phase allows the network to reach equilibrium. Only during the steady-state phase are statistics recorded. Each node has a maximum queue size of Q_{max}. Table 5.4 details the simulation parameters used for the analysis.

The N nodes participating are distributed uniformly in a square whose dimensions are given by $s = \sqrt{\frac{N}{\rho}}$ where ρ is the specified network density in units of nodes per m^2. Table 5.5 details the densities investigated. The density statistics were measured using 2000 unique, random networks

N	500
Q_{max}	40
Warm up Length	1000
Steady state Length	20000

Table 5.4: Simulation parameters

for each density. In order to calculate the network density statistics, simulations were run using the basic Lyui's algorithm and min hop routing where each link weight is one. The network diameter is the length of the longest shortest path between any two nodes. The average number of hops is the average number of times a node is forwarded until it reaches the final destination. The average number of 1-neighbors is the average number of nodes within one hop. However, our simulation does not use min hop routing. The routing metric given by Equation 5.12 emphasizes taking short, fast hops over long, slow hops. Consequently, the maximum and average path lengths are much higher in simulations as shown in Table 5.6. The simulation runs used to determine these statistics were using the default simulation parameters listed in this chapter.

	ρ	Diameter	Average # Hops	Average # 1-neighbors
Low	$\frac{1}{100^2}$	20.0	7.8	11.6
Med	$\frac{1}{75^2}$	13.9	5.4	20.1
High	$\frac{1}{50^2}$	8.9	3.5	42.8

Table 5.5: Network densities

	ρ	Maximum # Hops	Average # Hops
Low	$\frac{1}{100^2}$	48.4	15.7
Med	$\frac{1}{75^2}$	40.5	13.0
High	$\frac{1}{50^2}$	32.2	9.9

Table 5.6: Network densities routing path lengths

5.6 Statistics

During each simulation run, three measurements (end-to-end completion rate, throughput, and average delay) are recorded to compare network performance. End-to-end completion is the ratio of the number of packets that reach the final destination to the total number of packets that are generated. Throughput is the average number of packets delivered to the final destination in a given time slot. Average delay is the average number of time slots from when a packet is generated

until it is delivered to the final destination.

Chapter 6

Results

Results from simulation investigations are discussed in this chapter. First, the benefits of spatial multiplexing and secondary transmissions are examined. Then, the parameters for determining the set of secondary transmitters (p, m_{max}) are examined to justify the choices used elsewhere in this chapter. Then the benefits of using additional antennas (n) is discussed. Lastly, the importance of the routing metric is demonstrated. For each of the following experiments, the global generation rate, G, is swept from 0.1 to 4.0. For each generation rate, G, 200 trials were performed, each with a unique random network. The average across all 200 trials is depicted in the plots below with 95% confidence error bars.

6.1 Performance of Secondary Transmitters

The main inefficiency with standard time slotted scheduling MAC protocols is that when a node has no packets in its queue, the time slot is wasted. In these wasted time slots, the scheduled transmitter may be blocking nearby nodes that have packets which would otherwise be forwarded. The main contribution of this MAC protocol is to apply MIMO techniques that allow the time slot to be shared with neighboring nodes when the primary transmitter is unable to completely utilize the time slot. The benefit of sharing the time slots is illustrated in Figures 6.1, 6.4, and 6.5 which detail the end-to-end completion rate for networks with varying densities. When comparing performance, it is useful to observe when each curve crosses the 90% completion rate. The 90% thresholds for the three examined densities are listed in Table 6.1.

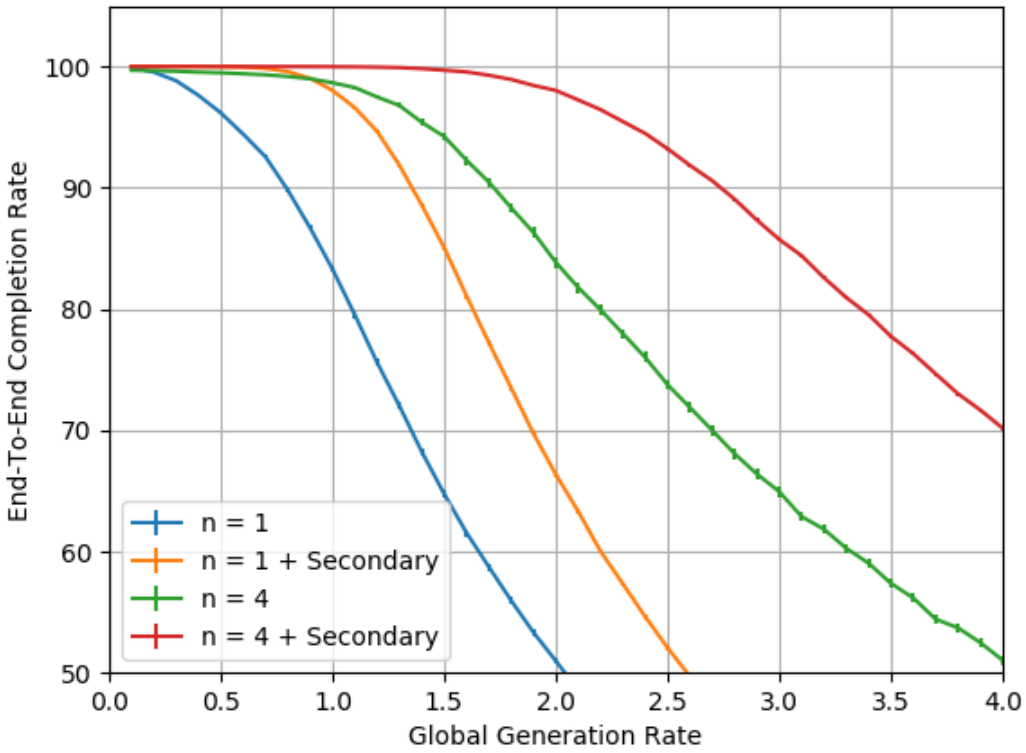

Figure 6.1: Completion rate versus secondary transmissions and spatial multiplexing in high density networks ($p = 10$, $m_{max} = 5$)

The performance of the set of high density networks is illustrated in Figure 6.1. The value of n indicates the number of antennas each node has, and the presence of "secondary" indicates that secondary transmissions have been enabled. The baseline to compare performance is given by $n = 1$, which constrains each node to only have a single antenna and disables secondary transmissions. Under these conditions, nodes are unable to use either spatial multiplexing or interference cancellation, because each node has only one degree of freedom. However, even with only one antenna, if secondary transmissions are enabled, and the primary transmitter's queue is empty, a secondary transmitter is still able to transmit. Even without the benefits of MIMO techniques, the sharing of the time slot increases the 90% threshold by a factor of 1.72. The spatial multiplexing gain of four antennas is given by $n = 4$, where the 90% threshold increases by a factor of 2.18 compared to the base case of $n = 1$. This large performance gain is because higher capacity links are taken advantage of. With just a single antenna, links have a maximum link rate of four packets.

However, with four antennas, links have a maximum link rate of 16 packets. Enabling secondary transmissions when $n = 4$ raises the 90% threshold by a factor of 3.47 compared to the base case of $n = 1$. The additional improvement is due to the interference cancellation that permits multiple transmitters if the primary is unable to utilize the slot capacity.

The performance of the set of high density networks is also illustrated in Figures 6.2 and 6.3 which show the average delay and throughput, respectively. In both cases, the network performance improves as the number of antennas, n, increases or secondary transmissions are allowed. For very low generation rates, the delay is minimized for a single antenna. This is because routing emphasizes fewer hops when high capacity links are unavailable. However, using delay to compare results can be misleading, because it is only measured for packets that reach their final destination, and at high generation rates, many packets never reach the final destination. Throughput and end-to-end completion are indirectly related, so all the observations already made have analogous observations that can be drawn from Figure 6.3. To avoid cluttering the results section, focus will be on evaluating network performance based on completion rate for the remainder of this chapter. Plots for all remaining experiments pertaining to average delay and throughput are included in Appendix B.

The performance of the set of networks with medium density are shown in Figure 6.4. The baseline performance with $n = 1$ and no secondary transmissions is better than for the high density network. The decreased density means smaller frames, so nodes get to transmit more often than in high density networks. By allowing secondary transmissions, the 90% threshold increases by a factor of 1.58 compared to $n = 1$. The spatial multiplexing gain for four antennas is 1.75 compared to the base case of $n = 1$. Lastly, the performance gain from spatial multiplexing and secondary transmissions is 2.46 compared to $n = 1$.

The performance of the set of networks with low density is shown in Figure 6.5. Similar trends are observed for low density networks as for previous scenarios. The performance gain from allowing secondary transmissions while restricted to one antenna is 1.36 compared to the base case of $n = 1$. However, the performance in this scenario exhibits a higher variability from point to point than experiments using other densities. In low density networks, there is a higher probability of a link with low SINR being needed to ensure network connectivity. These links act as bottleneck links, and our protocol is of little help in those cases. The increased variability is an indication of the sensitivity to the network topology. The spatial multiplexing gain of using four antennas is 1.42

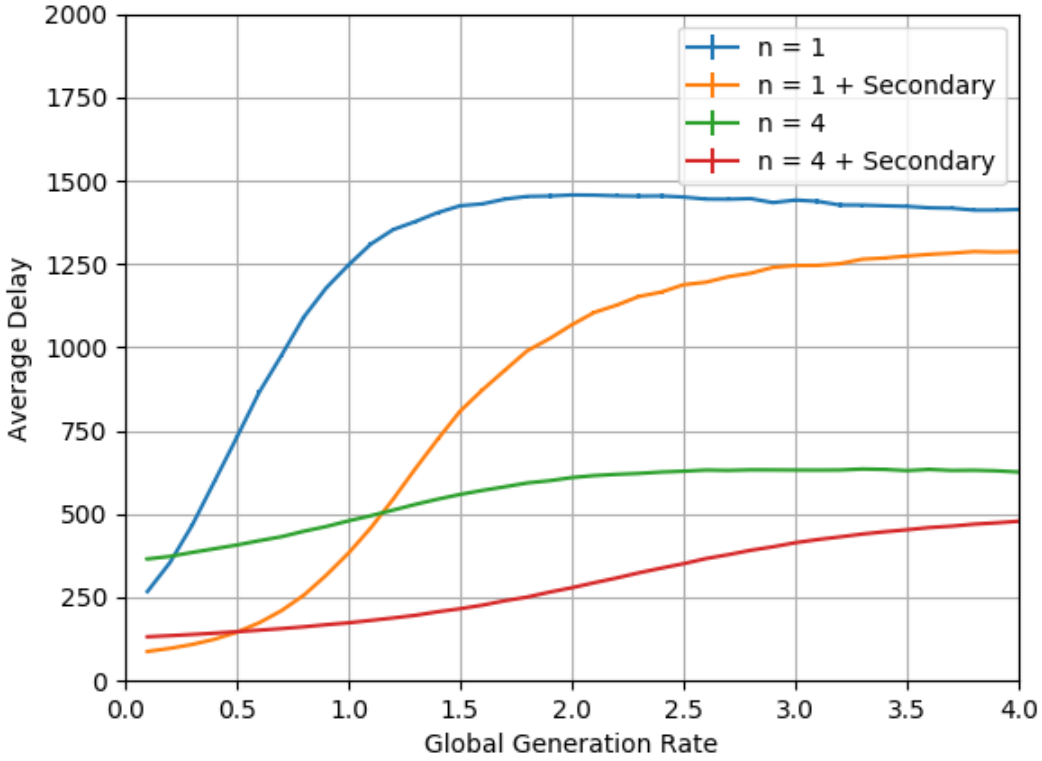

Figure 6.2: Average delay versus secondary transmissions and spatial multiplexing in high density networks ($p = 10$, $m_{max} = 5$)

compared to $n = 1$. Lastly, the performance gain from spatial multiplexing with four antennas and enabling secondary transmissions is 1.70 compared to $n = 1$.

In practice, the network density is not a parameter that can be controlled, instead it is a property of the network. We show that our protocol works well in a variety of scenarios. However, our protocol with four antennas and secondary transmissions provides the largest performance gain, compared to the base case of $n = 1$ for high density networks. This can be attributed to three reasons. First, the base performance with one antenna and no secondary transmissions is worse for high density networks because of larger frame sizes that result in fewer transmission opportunities. Second, the high density network has more high capacity links that are taken advantage of by spatial multiplexing. Third, increased network density means a larger secondary transmitter candidate set. Larger candidate sets generally means larger sets of secondary transmitters, which increases the chances of some node having traffic for the remaining slot capacity. The combination of higher

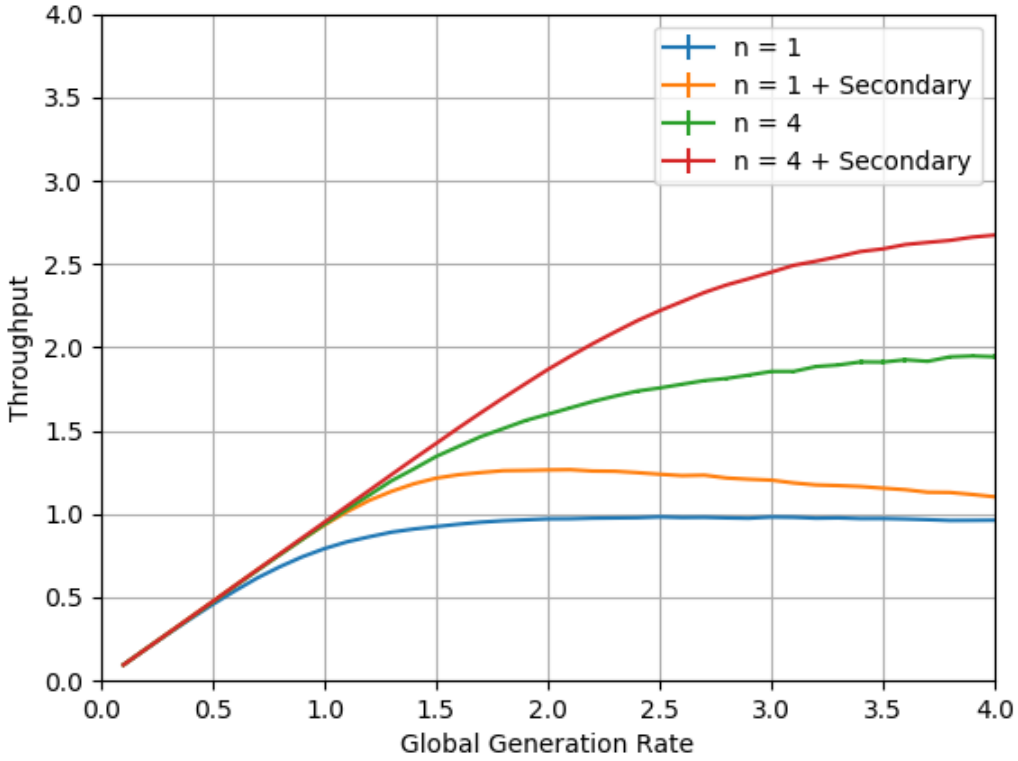

Figure 6.3: Throughput versus secondary transmissions and spatial multiplexing in high density networks ($p = 10$, $m_{max} = 5$)

average throughput on a per link basis and increased transmission opportunities as a secondary transmitter provide the performance gains described earlier.

	Low	Medium	High
$n = 1$	0.98	1.00	0.79
$n = 1$ & secondary	1.33	1.58	1.36
$n = 4$	1.39	1.75	1.72
$n = 4$ & secondary	1.67	2.46	2.74

Table 6.1: 90% thresholds of spatial multiplexing and secondary transmissions

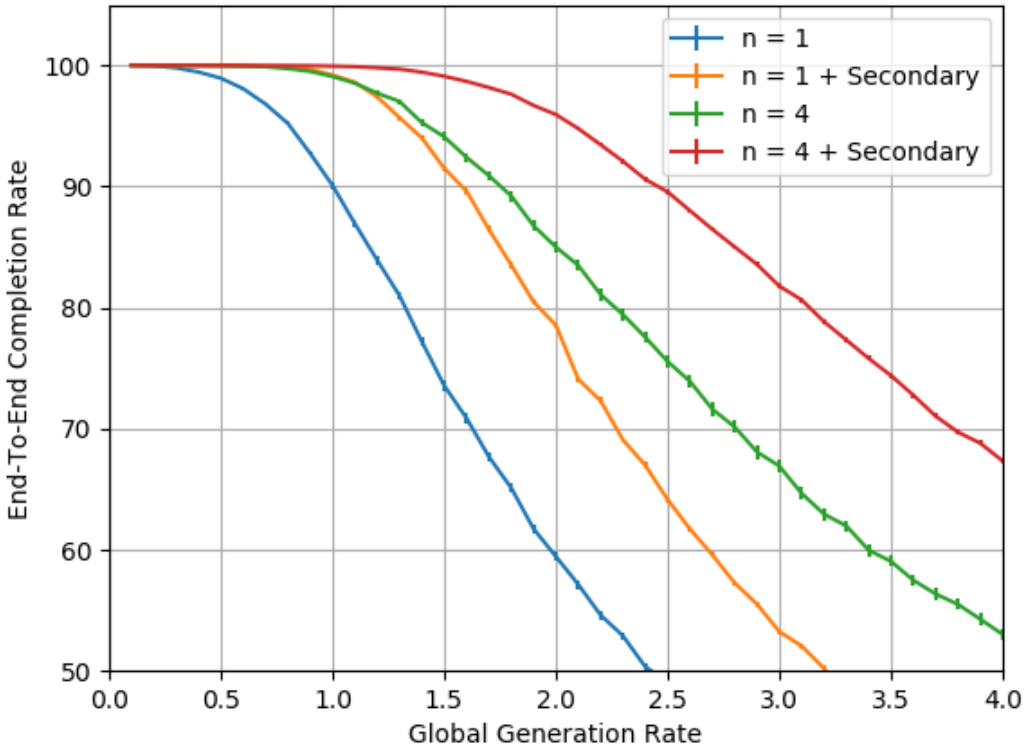

Figure 6.4: Completion rate versus secondary transmissions and spatial multiplexing in medium density networks ($p = 10$, $m_{max} = 5$)

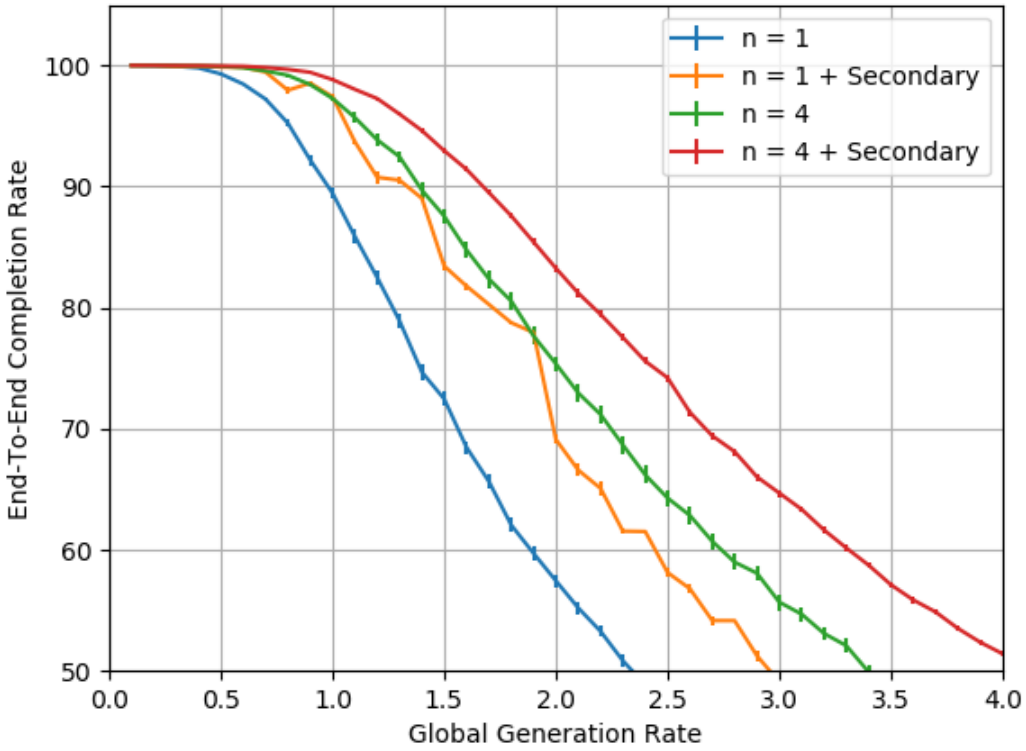

Figure 6.5: Completion rate versus secondary transmissions and spatial multiplexing in low density networks ($p = 10$, $m_{max} = 5$)

6.2 Secondary Transmitter Candidates

The parameter, p, constrains the set of secondary transmitters with the goal of limiting changes to the interference environment. For instance, $p = 1$ allows every 1-neighbor to be a secondary transmitter candidate and gives the largest possible candidate set. Larger sets of secondary transmitters translates to more chances for secondary transmitters to have packets to reclaim a slot. Allowing $p = 1$ also allows for nodes to potentially create different multiple-access interference by transmitting in slots not assigned to them. For this reason, it is advisable to use $p > 1$ to limit creating additional interference. However, if p is over constrained, the set of candidates for secondary transmitters may be reduced to the point that the candidate set is empty, and secondary transmissions will not occur. Selecting an appropriate value of p needs to balance the concerns of added interference and the number of secondary transmitters. Figures 6.6, 6.8, and 6.9 detail the end-to-end completion rate for networks with varying densities. The 90% thresholds for the three examined densities are listed in Table 6.2.

p	Low	Medium	High
1	1.86	2.43	2.50
5	1.74	2.50	2.68
10	1.67	2.46	2.74
20	1.58	2.33	2.75
40	1.49	2.13	2.68
100	1.42	1.93	2.38
200	1.42	1.86	2.22

Table 6.2: 90% thresholds for candidate secondary transmitter set (p)

For the high density network in Figure 6.6, as p increases, performance improves until $p = 20$, and then falls to be worse than the initial performance with $p = 1$. This illustrates the balance between additional interference and limiting secondary transmitters. Initially, allowing every 1-neighbor to be a secondary transmitter creates too much additional interference, which causes packets to be dropped because the SINR falls below β. Figure 6.7 illustrates the rapid drop off of the number of link errors as p is increased. Then as $p \gg 1$ the candidate set of secondary transmitters shrinks, resulting in fewer nodes that can take advantage of the remaining slot capacity. The correct balance of the two concerns with high density networks appears to be between $p = 10$ and $p = 20$.

For the medium density network in Figure 6.8, as p is increased, performance remains fairly

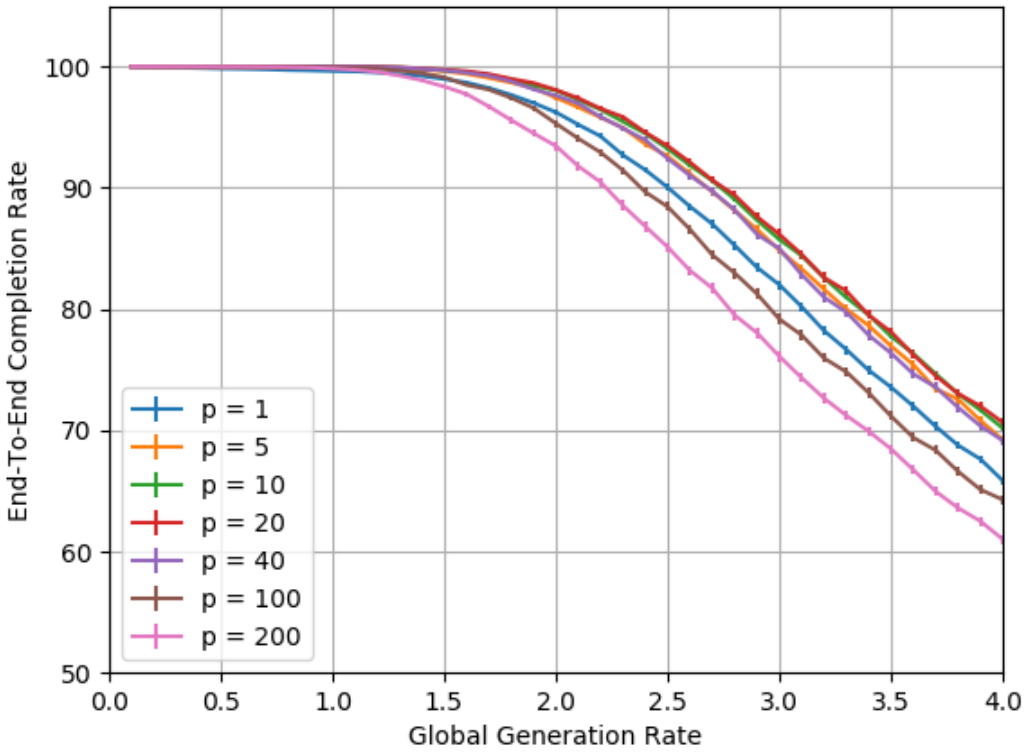

Figure 6.6: Completion rate versus secondary transmitter candidate set size in high density networks ($n = 4$, $m_{max} = 5$)

stationary until $p = 20$ where performance begins to drop. This set of networks illustrates a similar balance compared to the high density network, because performance drops after $p = 10$. However, unlike the high density networks, performance does not improve initially as p is increased. This is likely due to the decrease in network density reducing the secondary transmitter candidate set.

The low density network detailed in Figure 6.9 performs best for $p = 1$, and performance drops as the candidate set is further restricted. This set of networks does not show the trade off seen with high density networks, because the size of the secondary transmitter candidate set is already so restricted for $p = 1$. As with the previous experiment, the lower density network scenarios have a higher variability from point to point. This is because low density networks are more sensitive to the random topology of the network, where it is more likely for the network bottlenecks to be single links with low SINR that can not be taken advantage of.

The value $p = 10$ is utilized for the remaining investigations presented in this thesis. It

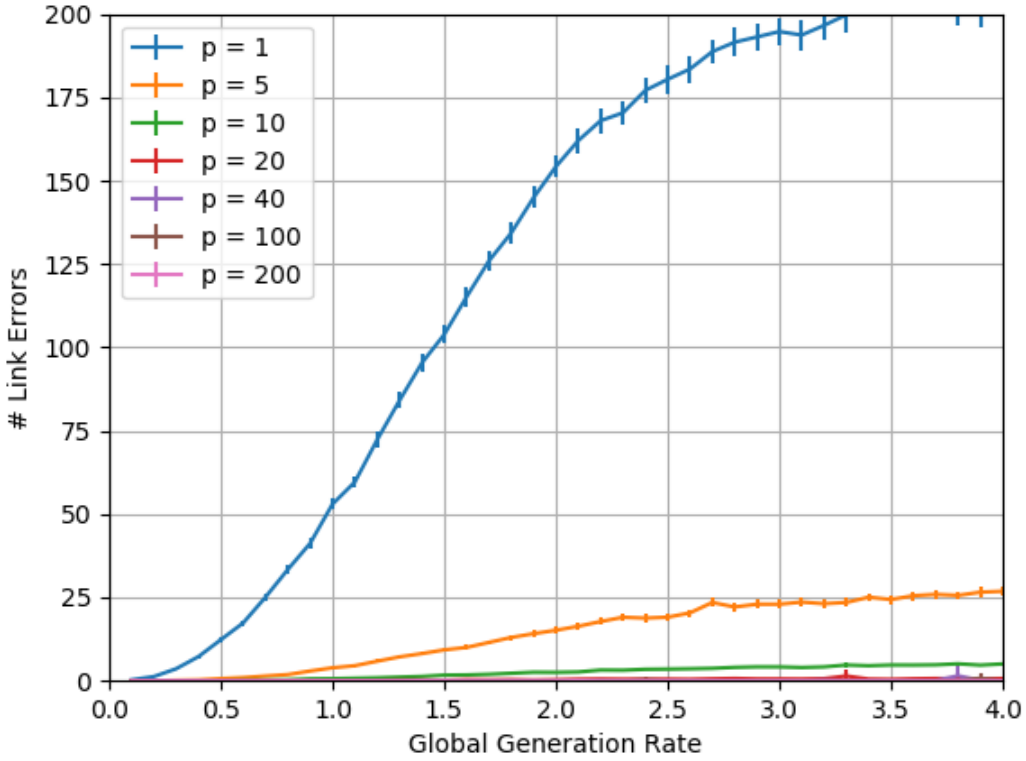

Figure 6.7: Link errors versus secondary transmitter candidate set size in high density networks ($n = 4$, $m_{max} = 5$)

provided the best performance for high and medium density network scenarios investigated in this work, without being far below the best value of p for low density networks.

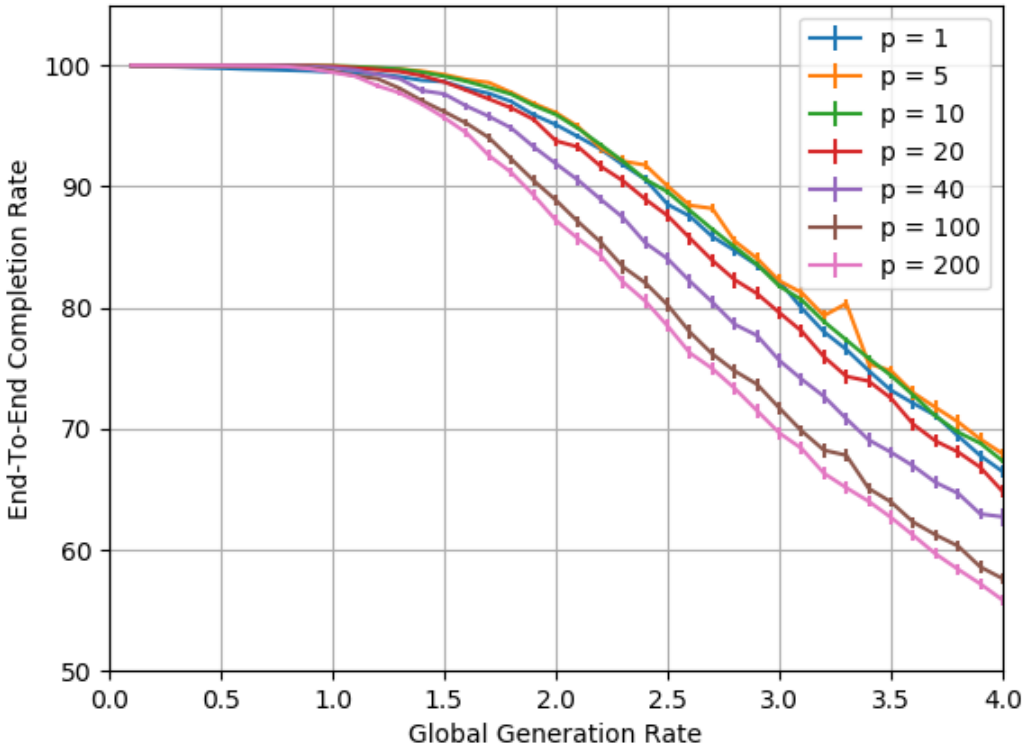

Figure 6.8: Completion rate versus secondary transmitter candidate set size in medium density networks ($n = 4$, $m_{max} = 5$)

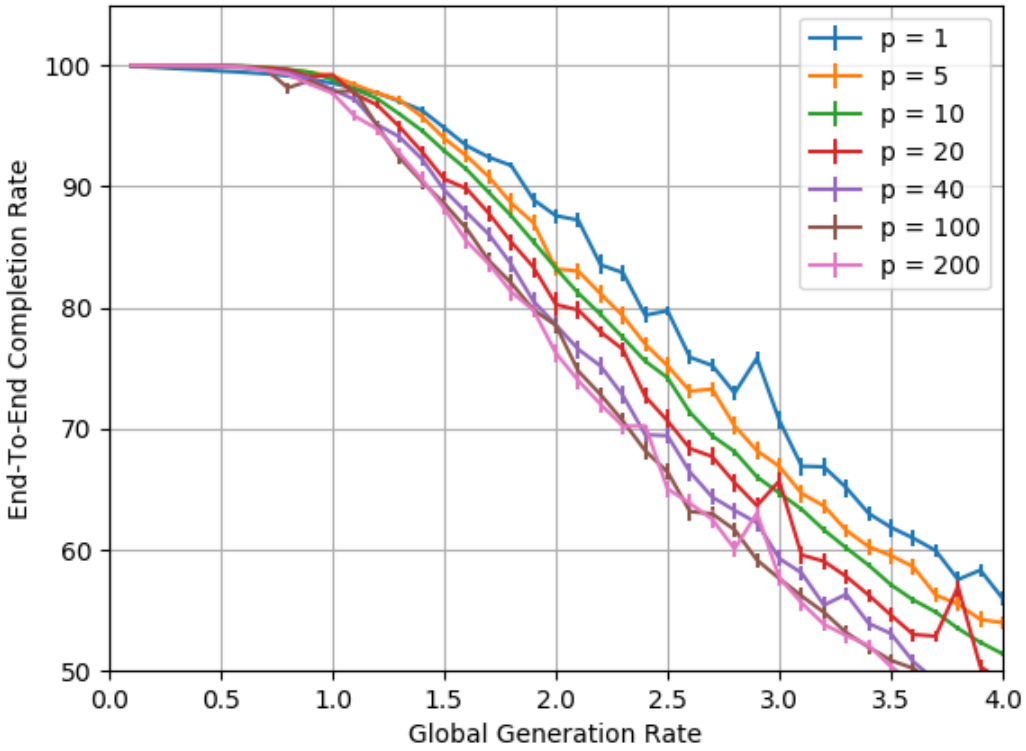

Figure 6.9: Completion rate versus secondary transmitter candidate set size in low density networks ($n = 4$, $m_{max} = 5$)

6.3 Maximum Number of Secondary Transmitters

The parameter, m_{max}, determines the maximum number of secondary transmitters that can be selected from the candidate set. While at most n, where n is the number of antennas, secondary transmitters may transmit in a given time slot, it is not necessary to limit the number of secondary transmitters to n. This is because secondary transmitters are only of use if they have traffic that they can forward. By allowing the set of secondary transmitters to be larger, the probability of a secondary transmitter having a packet to forward is increased. The effect of m_{max} is demonstrated in Figures 6.10, 6.11, and 6.12. The 90% thresholds are given in Table 6.3. There are 500 nodes in the simulation, so setting $m_{max} = 500$ provides the maximum secondary transmitter set size. The only reason for a candidate transmitter to not be included is if it is not communicable with every other secondary transmitter.

m_{max}	Low	Medium	High
1	1.50	1.92	2.05
2	1.61	2.12	2.29
3	1.63	2.28	2.46
4	1.67	2.38	2.61
5	1.67	2.46	2.74
6	1.69	2.53	2.83
7	1.71	2.54	2.90
8	1.68	2.54	2.95
9	1.67	2.54	3.01
10	1.65	2.53	3.02
500	1.63	2.58	3.12

Table 6.3: 90% thresholds for maximum number of secondary transmitters (m_{max})

Figure 6.10 shows how as the value of m_{max} increases, so does the network performance. The best performance is given by $m_{max} = 500$ which allows for the largest possible secondary transmitter sets. The 90% threshold for this case is a factor of 1.52 better than only allowing one secondary transmitter. Approximately 52% of the improvement occurs by the time $m_{max} = n = 4$. As m_{max} increases past 4, we experience diminishing returns, and the performance approaches the upper bound given by $m_{max} = 500$. However, increasing m_{max} also increases the overhead. The overhead is illustrated by Figure 4.3. As the size of the secondary transmitter set increases, a larger fraction of the time slot is consumed by coordinating the primary transmitter and m secondary transmitters. The trade off is not present in these works, because the coordination phase of the time slot is not simulated. Similar behavior is observed for medium density networks in Figure 6.11 where

approximately 85% of the limit is reached when $m_{max} = n = 4$. For low density networks depicted in Figure 6.12, performance does not improve for $m_{max} > 2$. The decreasing impact of m_{max} as the network density decreases is because the average size of the candidate set is shrinks, and m_{max} does not have an impact on the size of the secondary transmitter set.

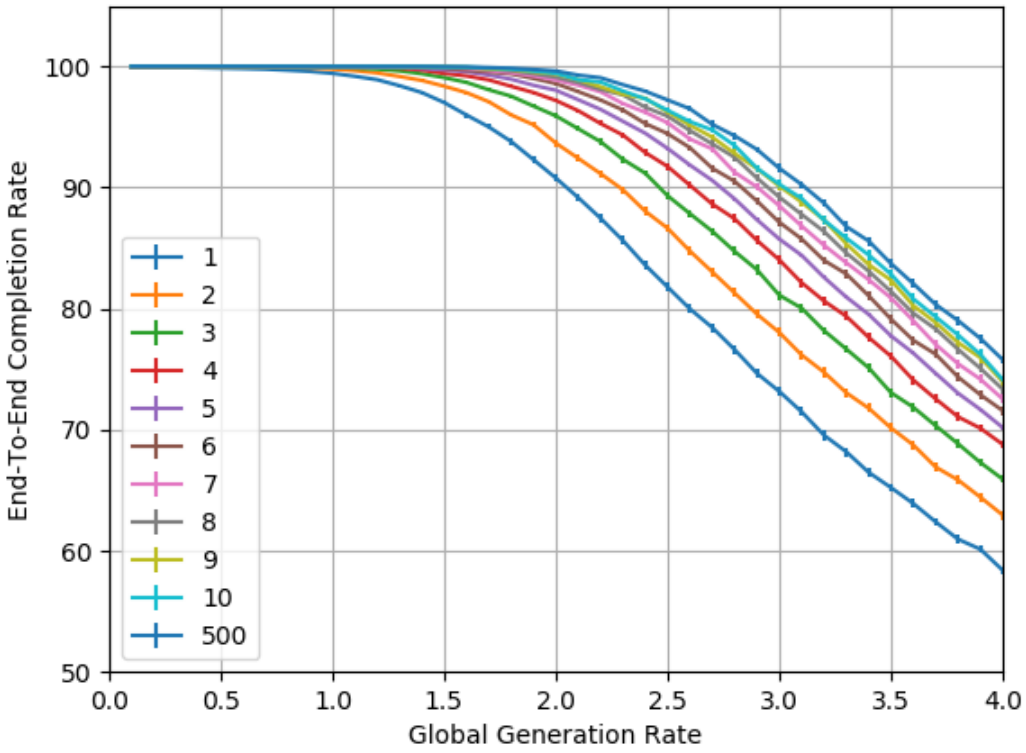

Figure 6.10: Completion rate versus maximum number of secondary transmitters in high density networks ($n = 4$, $p = 10$)

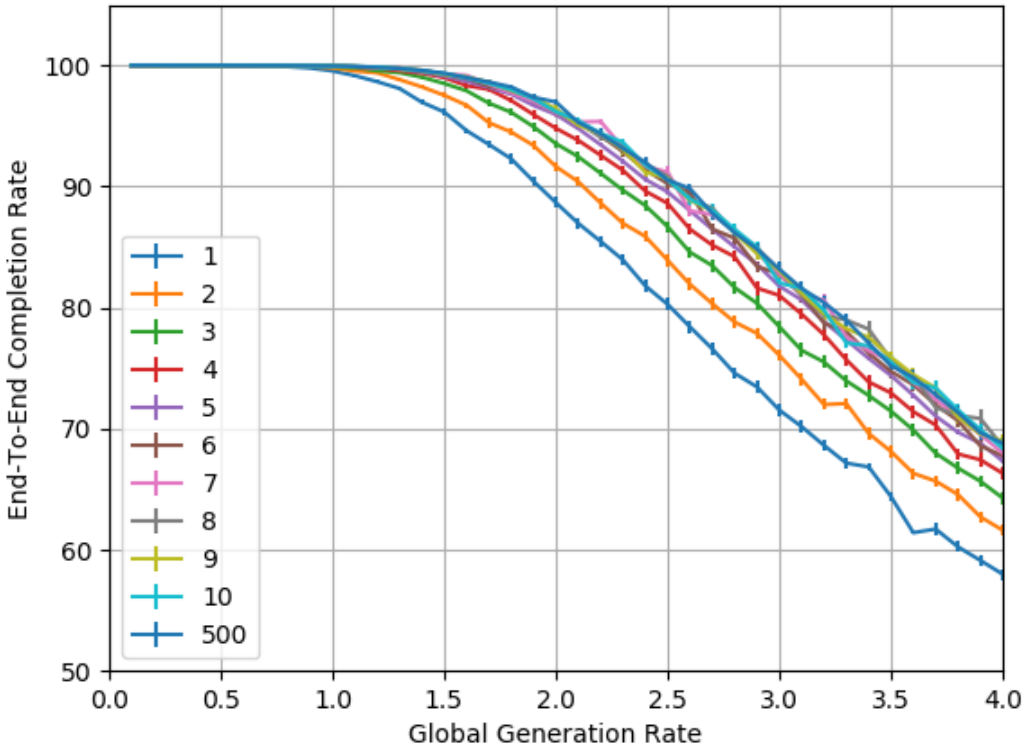

Figure 6.11: Completion rate versus maximum number of secondary transmitters in medium density networks ($n = 4$, $p = 10$)

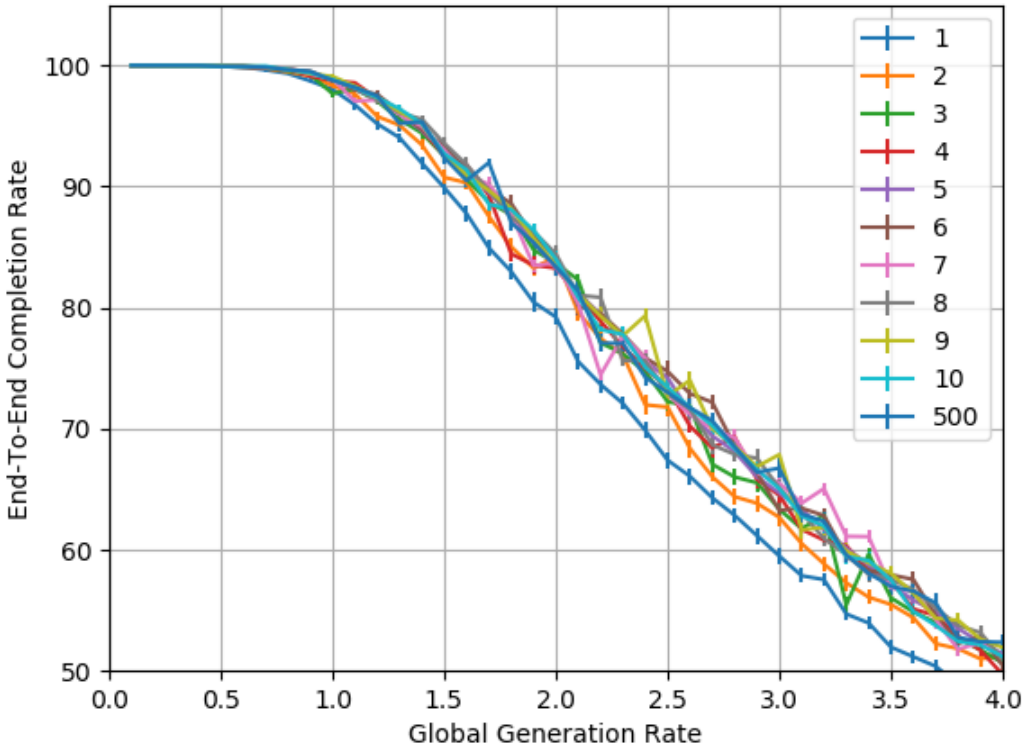

Figure 6.12: Completion rate versus maximum number of secondary transmitters in low density networks ($n = 4$, $p = 10$)

6.4 Number of Antennas

The number of antennas, n, determines the number of degrees of freedom each node has, as well as the possible link rates. Figures 6.13, 6.14, and 6.15 detail the end-to-end completion rate for networks with varying densities. The 90% thresholds are listed in Table 6.4.

n	Low	Medium	High
1	1.33	1.58	1.36
2	1.52	1.91	1.87
3	1.60	2.23	2.36
4	1.67	2.46	2.74
5	1.72	2.63	3.04
6	1.73	2.76	3.22
7	1.76	2.87	3.42
8	1.76	2.98	3.60
9	1.80	3.11	3.76
10	1.79	3.14	3.88

Table 6.4: 90% thresholds for number of antennas (n)

As shown in Figures 6.15 and 6.14 for low and medium density networks, as more antennas are added, performance improves until a limit is reached. It appears that the limit for low density networks is approximately 1.80, and for medium density networks, the limit is approximately 3.14 each with 10 antennas. Similar behavior is observed of high density networks shown in Figure 6.13. However, the limit is approximately 4.76 with $n = 20$. It is important to note that while performance does improve as n increases, the channel state matrix, H, increases like n^2, so the overhead of distributing the CSI increases as well. The overhead of distributing the CSI is not simulated, so no trade offs are seen in the results.

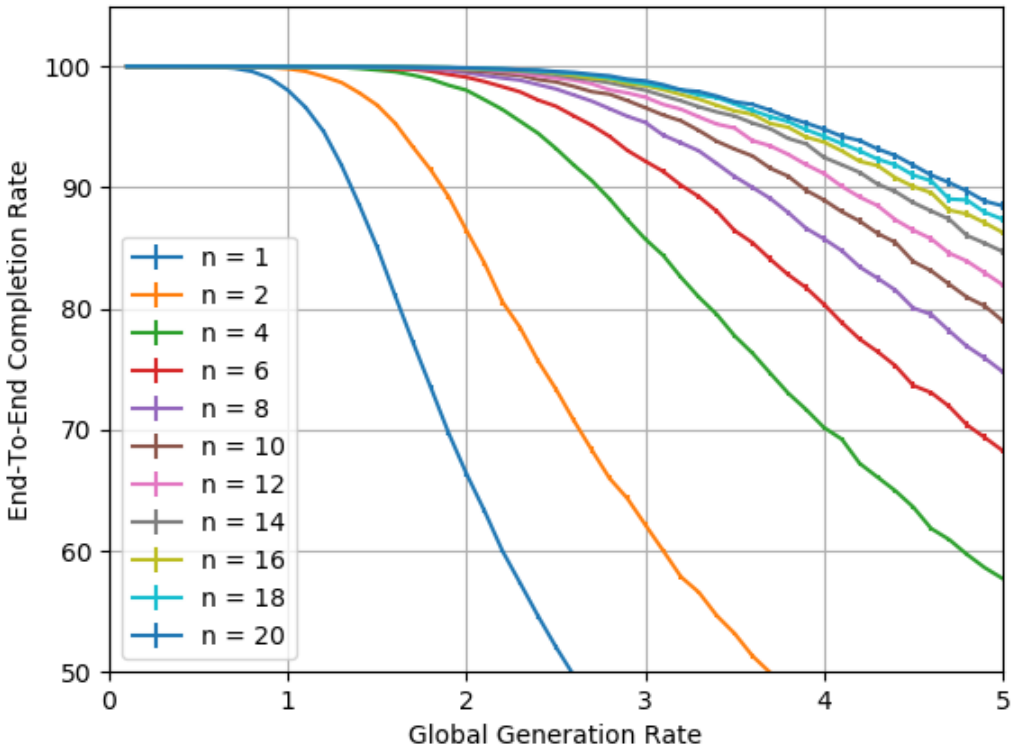

Figure 6.13: Completion rate versus number of antennas in high density networks ($p = 10$, $m_{max} = 5$)

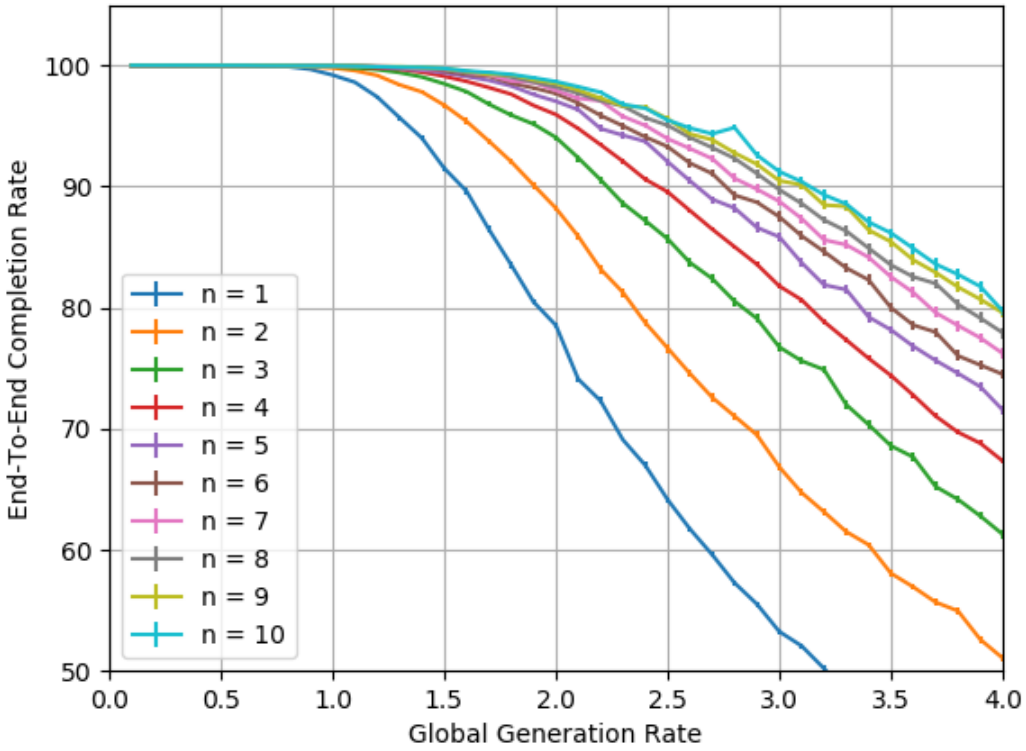

Figure 6.14: Completion rate versus number of antennas in medium density networks ($p = 10$, $m_{max} = 5$)

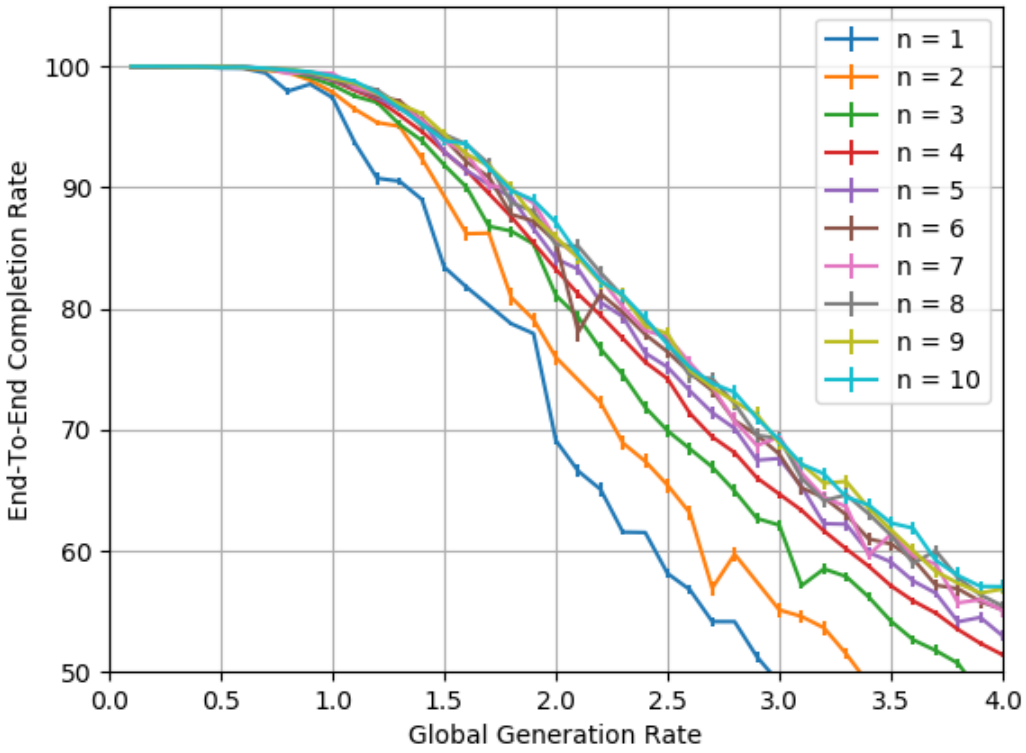

Figure 6.15: Completion rate versus number of antennas in low density networks ($p = 10$, $m_{max} = 5$)

6.5 Routing Metric

Traditional views of networks prefer to treat separate layers as independent entities. However, wireless ad hoc networks have a wide distribution of links, and their performance is often limited by a few bottleneck nodes. Cross-layer routing metrics enable networks to take advantage of high capacity links as well as achieve load balancing at network bottlenecks. The importance of the routing metric or link weight is give in Figures 6.16, 6.17, and 6.18. The routing metric "MinHop_Basic" assigns every link weight to be one. The routing metric "NodeUtil_ETR_LinkRate" is the link weight described by Equation 5.8. The 90% thresholds are given in Table 6.5. When min hop is used to generate link weights, performance does not seem to depend on the network density. However, high density networks, using the "NodeUtil_ETR_LinkRate" link weights, perform a factor of 4.81 better than min hop.

Link Weight	Low	Medium	High
MinHop_Basic	0.56	0.51	0.57
NodeUtil_ETR_LinkRate	1.67	2.46	2.74

Table 6.5: 90% thresholds for routing metric

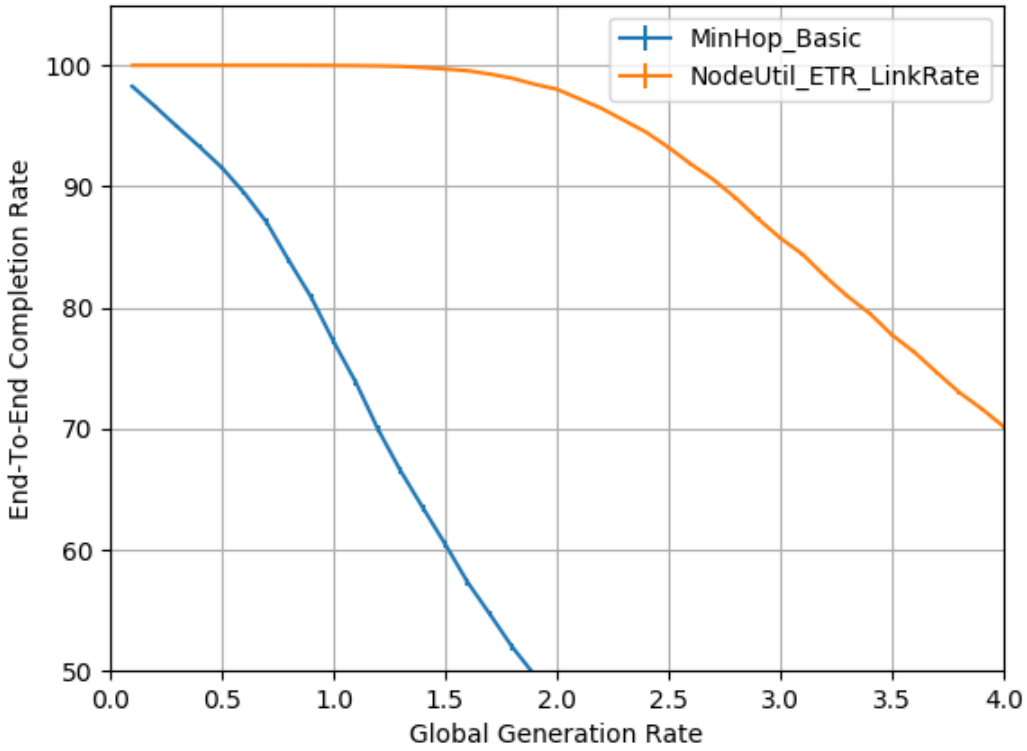

Figure 6.16: Completion rate versus routing metric in high density networks ($p = 10$, $m_{max} = 5$, $n = 4$)

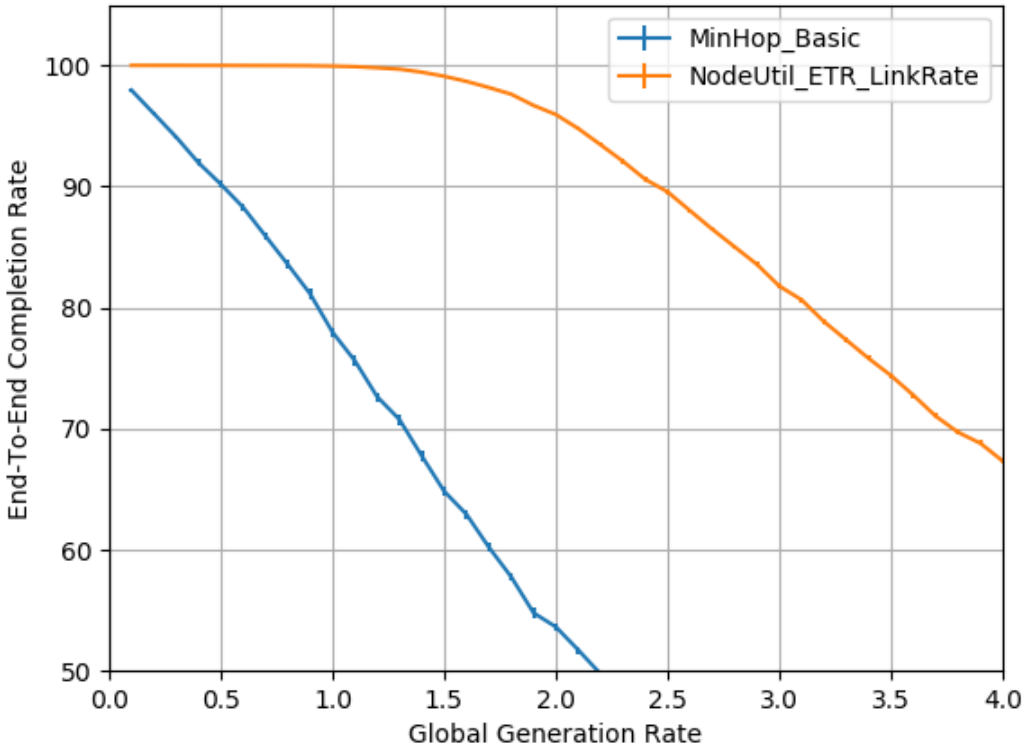

Figure 6.17: Completion rate versus routing metric in high density networks ($p = 10$, $m_{max} = 5$, $n = 4$)

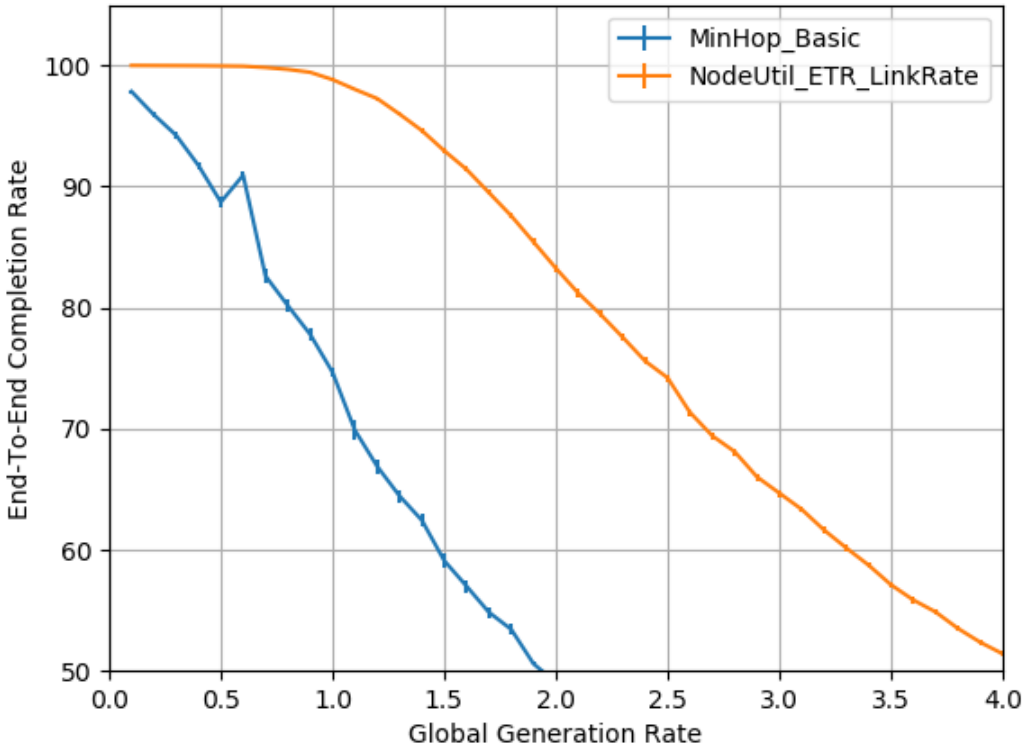

Figure 6.18: Completion rate versus routing metric in high density networks ($p = 10$, $m_{max} = 5$, $n = 4$)

Chapter 7

Conclusions

We have designed and investigated a new protocol for ad hoc networks that use MIMO techniques to capture unused capacity of time slots in scheduled TDMA protocols. In many scheduling protocols, if a node is scheduled to transmit in a time slot but does not have any packets in its queue, the slot is unused. This leads to poor network performance, because scheduled nodes prevent other nodes with packets to forward from transmitting. We propose a protocol that uses interference cancellation to allow nodes not scheduled in a particular time slot to transmit. This approach requires overhead in distributing the CSI and coordinating between primary and secondary transmitters. Each node must transmit pilot symbols to determine the CSI and then distribute the information to neighboring nodes over the course of a frame. The other source of overhead occurs when primary and secondary transmitters coordinate at the beginning of the slot, so that the data can be encoded correctly. We also use adaptive spreading and spatial multiplexing to create high capacity links that are capable of being used by either primary or secondary transmitters.

We use Lyui's scheduling algorithm to demonstrate our approach. However, our protocol is independent of the scheduling algorithm as long as slot assignments ensure broadcast capabilities. Our protocol preserves the idea of fairness, because each node is still guaranteed to be a primary transmitter at least once per frame. Two ideas work together to ensure that the underlying scheduling algorithm, like Lyui's algorithm, continues to work despite unscheduled nodes transmitting. First, we require high SINR between primary and secondary transmitters to hopefully provide similar multiple-access interference. Second, we limit the total output power of the primary and secondary transmitters to be the same as if the primary is the only transmitter.

We use simulations to show that our protocol improves end-to-end completion, throughput, and delay in large random ad hoc networks with varying densities. Additional investigations determined appropriate constraints for the secondary transmitter candidate set and the maximum size of the secondary transmitter set. Investigations also show that as more antennas are used, gains are limited. We also confirm the importance of routing in allowing the network to take advantage of our new protocol.

Future work could provide further improvements through several avenues. Nonrandom, intelligent ways to filter the candidate set of secondary transmitters may lead to higher slot utilization. Routing may provide additional performance gains if the concept of secondary transmissions is somehow incorporated. Another way that routing could lead to gains is to modify the forwarding tables to reflect forward progress. Then, during the budget allocation phases, additional packets could be packed onto links via adaptive spreading even if the link is not the designated next hop.

Appendices

Appendix A Interference Analysis of Secondary Transmitters

The parameter p in Equation (4.1) is used to constrain the candidates for secondary transmitters. The smaller the value of p, the more secondary transmitter candidates. Unfortunately, smaller values of p can cause the interference environment to change more than larger values of p. An analysis of the impact of p on the interference environment based on the underlying channel model follows.

Consider the scenario depicted in Figure 1, in which nodes i, j, and k are co-linear with the $SNR_{i,j} = SNR_{i,k} = f\beta$ where $f \leq 1$. Nodes j and k are thus not communicable with i. However, the interference caused by node i is received by nodes j and k. Let node i have all m secondary transmitters located at point A, where $SINR_{i,A} = p\beta$ and $p \geq 1$. Nodes transmitting at A in node i's time slot will create different interference environments than Lyui's algorithm considers.

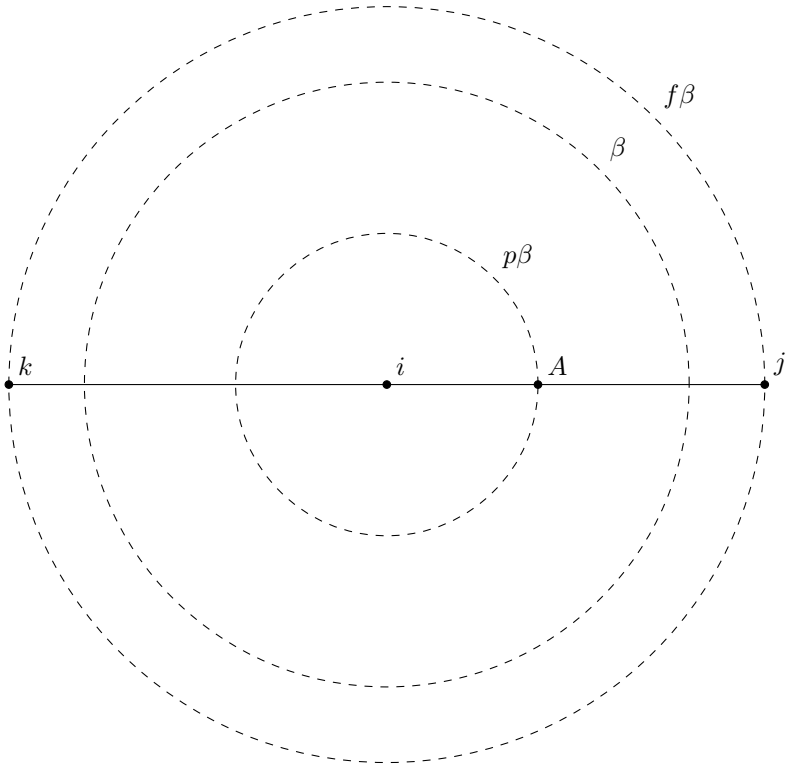

Figure 1: Worst case and best case interference environments.

In this scenario, let node i not have any packets in its queue, so the secondary transmitters

at A use the entire power budget P_t. The measurement $\Phi(x)$ measures the interference gain at node x caused by the secondary transmitters compared to the interference that node i would have created. We define

$$\Phi(x) = \frac{SINR_{A,x}}{SINR_{i,x}} \tag{1}$$

when there is no multiple-access interference. In the scenario outlined in Figure 1, it is clear that $\Phi(x)$ is maximized for node j, denote this $\Phi_{worst} = \Phi(j)$. Conversely, $\Phi(x)$ is minimized for node k, denoted $\Phi_{best} = \Phi(k)$. Using the Equations 3.1 and 3.2 that model the channel, the following can be derived:

$$\Phi_{worst} = [1 - (\frac{f}{p})^{\frac{1}{\alpha}}]^{-\alpha} \tag{2}$$

$$\Phi_{best} = [1 + (\frac{f}{p})^{\frac{1}{\alpha}}]^{-\alpha} \tag{3}$$

Figure 3 shows the worst case interference gain, Φ_{worst}, caused by different values of p depending on the distance from the primary transmitter when the communicable range, R, is $200m$. Figure 2 is a similar plot showing the best case interference gain Φ_{best}. Because we have established upper and lower bounds for the interference gain, we know that the observed interference gain will fall somewhere between the bounds for any other network configuration. For instance, consider the value $p = 10$ and a node located $2R = 400m$ from the primary node. The interference gain $\Phi(x)$ is bounded $0.45 \leq \Phi(x) \leq 3$ independent of the network topology.

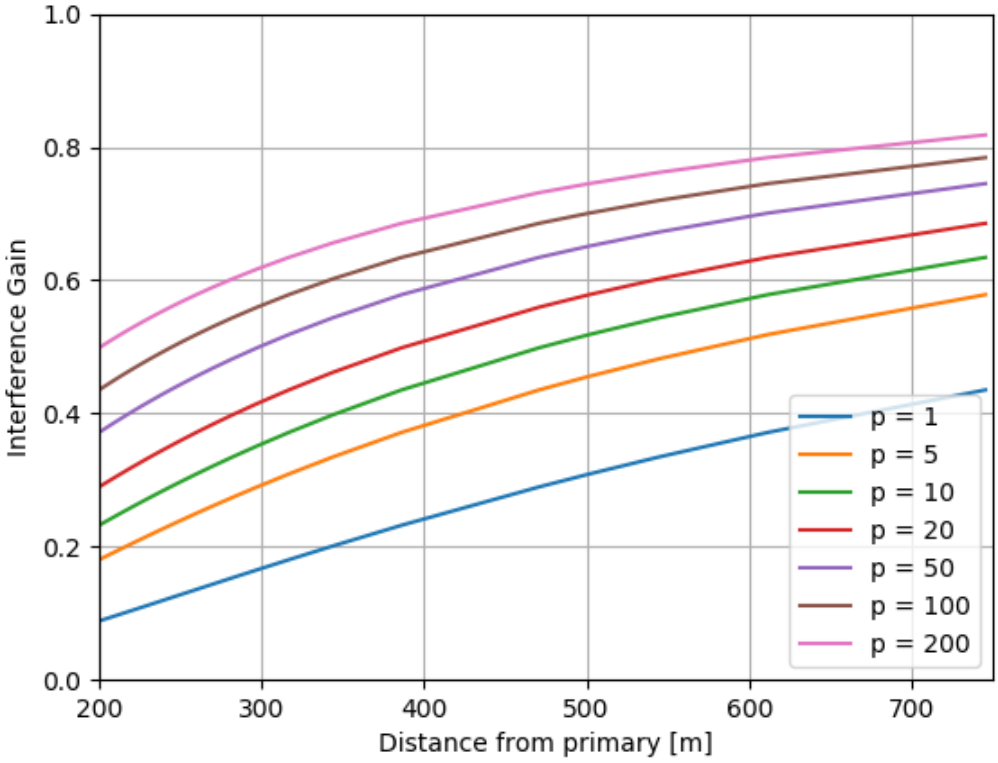

Figure 2: Best case interference gain when selecting secondary transmitters.

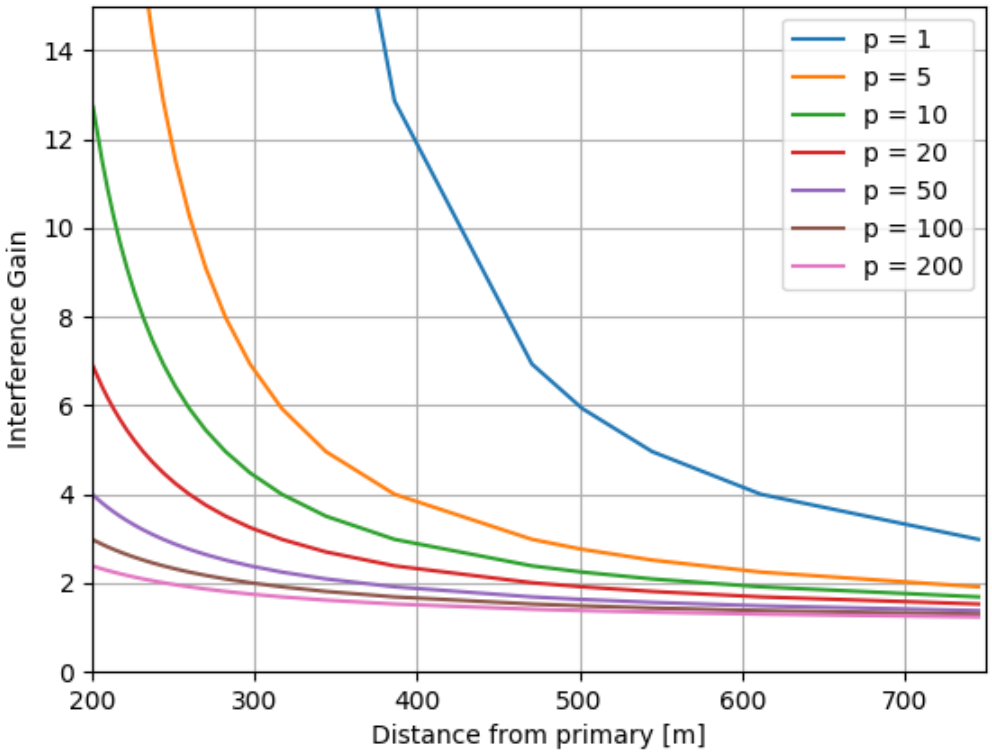

Figure 3: Worst case interference gain when selecting secondary transmitters.

Appendix B Additional Results

B.1 Importance of Secondary Transmitters

Figures 4 and 5 show the impact that allowing secondary transmissions has on the average delay and throughput, respectively, for all three tested densities. In general, the combination of using four antennas and secondary transmissions provides the best results regardless of density for both statistics.

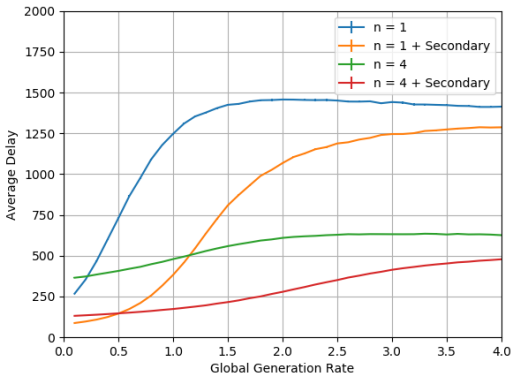

(a) High Density

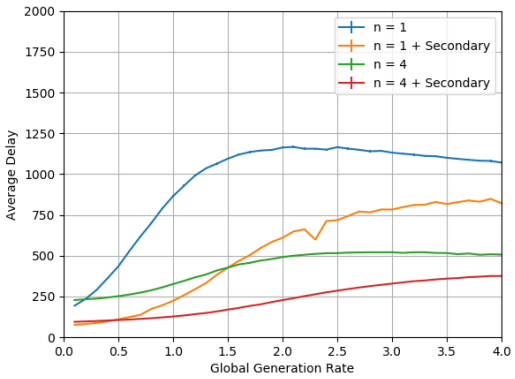

(b) Medium Density

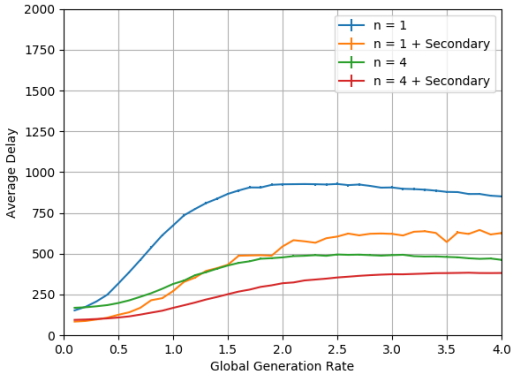

(c) Low Density

Figure 4: Average delay versus secondary transmissions and spatial multiplexing ($p = 10$, $m_{max} = 5$)

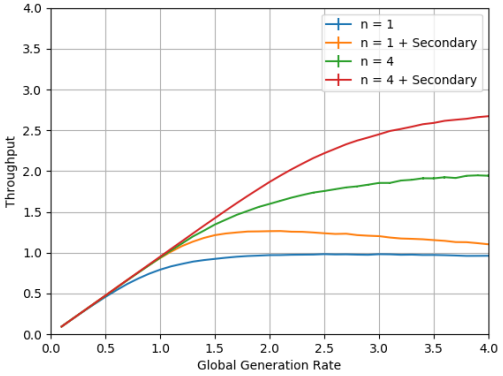

(a) High Density

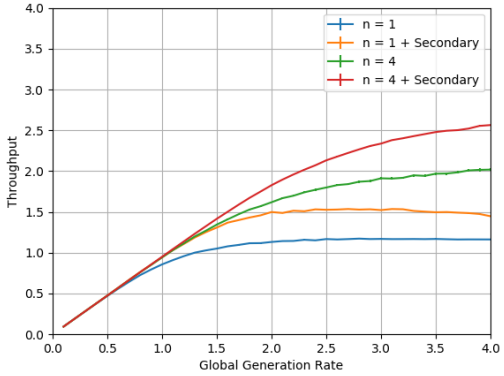

(b) Medium Density

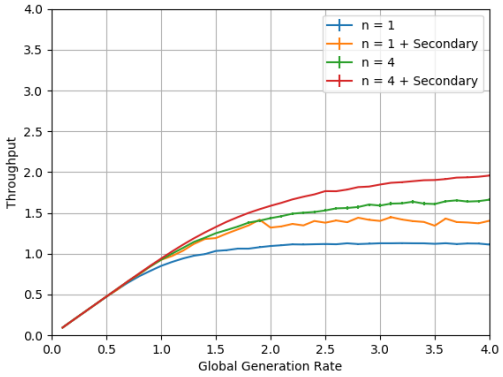

(c) Low Density

Figure 5: Throughput versus secondary transmissions and spatial multiplexing ($p = 10$, $m_{max} = 5$)

B.2 Secondary Transmitter Candidates

Figures 6 and 7 show the impact that the secondary transmitter candidate set has on the average delay and throughput, respectively, for all three tested densities. The value of p hardly makes a difference in either statistic, regardless of the density. This is likely because both statistics only measure packets that successfully reach the final destination. Figure 8 shows the number of link errors as p is changed. In general, as long as $p > 5$ there does not seem to be a significant difference in the number of link errors, regardless of network density.

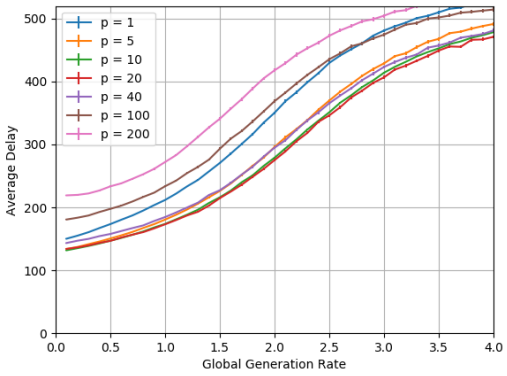

(a) High Density

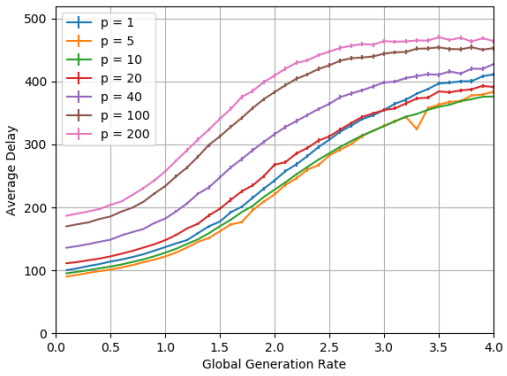

(b) Medium Density

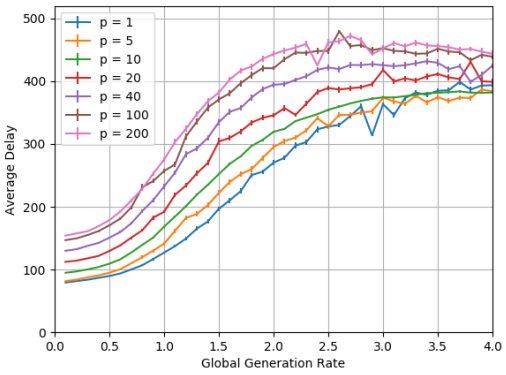

(c) Low Density

Figure 6: Average delay versus secondary transmitter candidate set size ($n = 4$, $m_{max} = 5$)

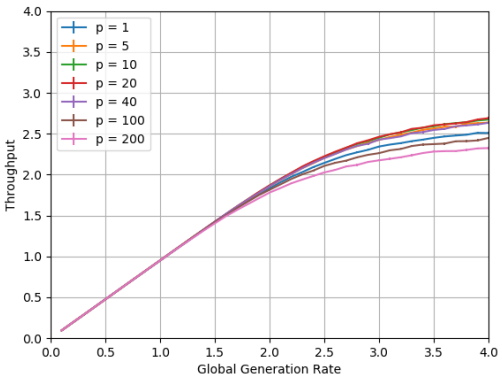

(a) High Density

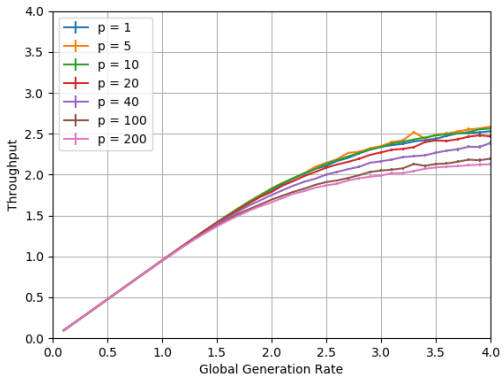

(b) Medium Density

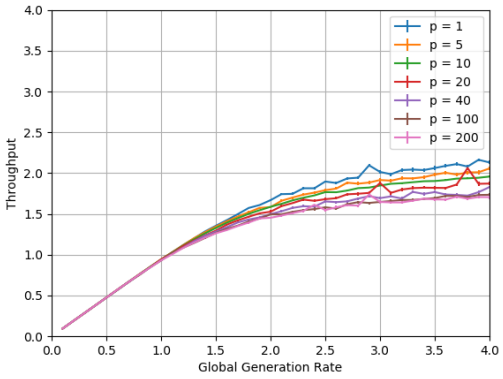

(c) Low Density

Figure 7: Throughput versus secondary transmitter candidate set size ($n = 4$, $m_{max} = 5$)

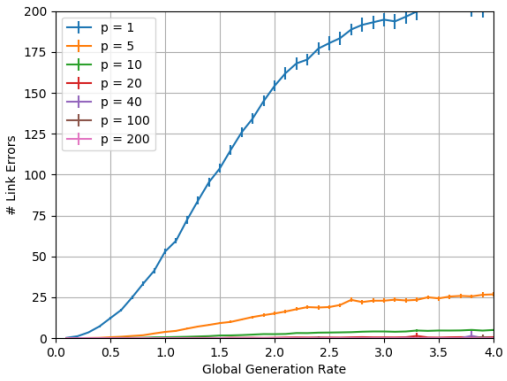

(a) High Density

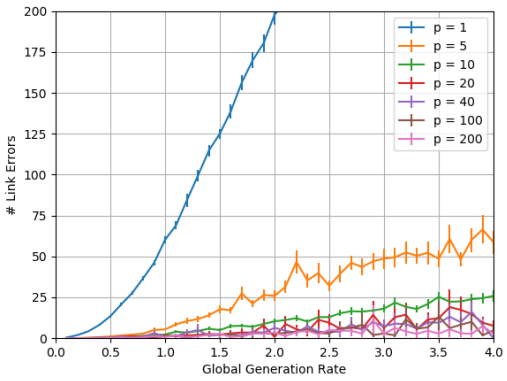

(b) Medium Density

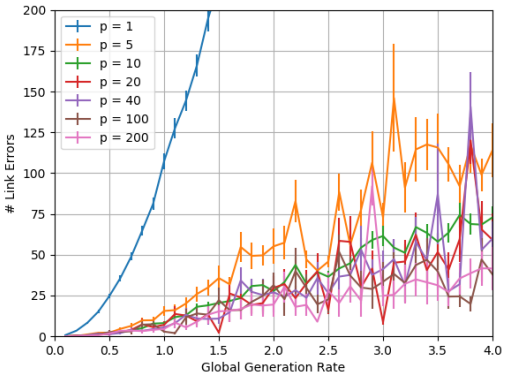

(c) Low Density

Figure 8: Link errors versus secondary transmitter candidate set size ($n = 4$, $m_{max} = 5$)

B.3 Max Number of Secondary Transmitters

Figures 9 and 10 show the impact that the maximum size of the secondary transmitter set has on the average delay and throughput, respectively, for all three tested densities. For high density networks, the delay is slightly improved as the maximum set size is increased. This is because nodes are more likely to be offered a chance to be a secondary transmitter, allowing them to transmit more often. However, for lower density networks the improvement seen by increasing is reduced. This is because the smaller density networks means a smaller candidate set, and the parameter m_{max} does not come into play at all.

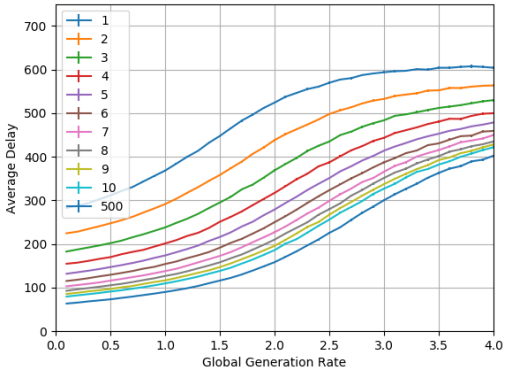

(a) High Density

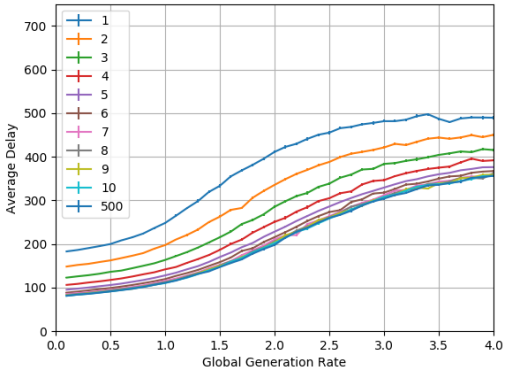

(b) Medium Density

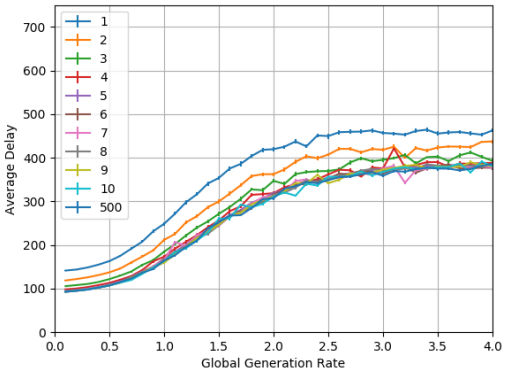

(c) Low Density

Figure 9: Average delay versus maximum number of secondary transmitters ($n = 4$, $p = 10$)

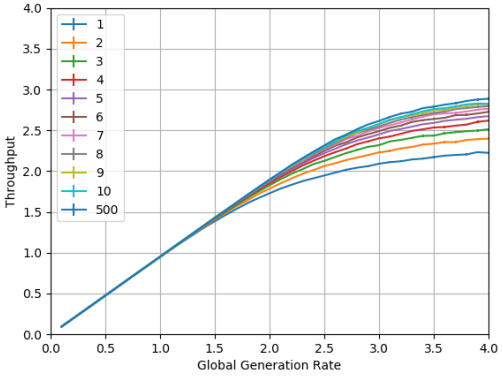

(a) High Density

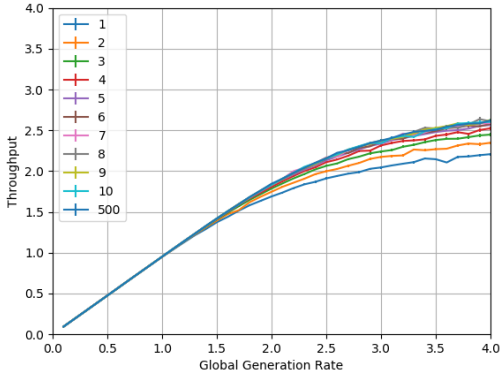

(b) Medium Density

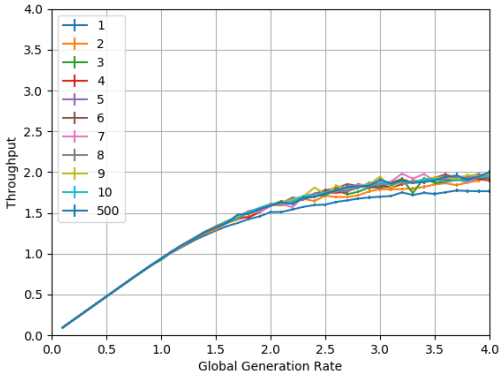

(c) Low Density

Figure 10: Throughput versus maximum number of secondary transmitters ($n = 4$, $p = 10$)

B.4 Number of Antennas

Figures 11 and 12 show the impact that the number of antennas has on the average delay and throughput, respectively, for all three tested densities. In general, the delay decreases as more antennas are added to a node regardless of network density. However, the lower density networks approach a limit faster than the high density network. The same pattern is observed in Figure 12 which shows that the throughput increases as more antennas are added until an upper bound is reached. The reason the high density networks have better performance bounds than lower density networks is because of the higher probability of high capacity links, which can take advantage of more antennas.

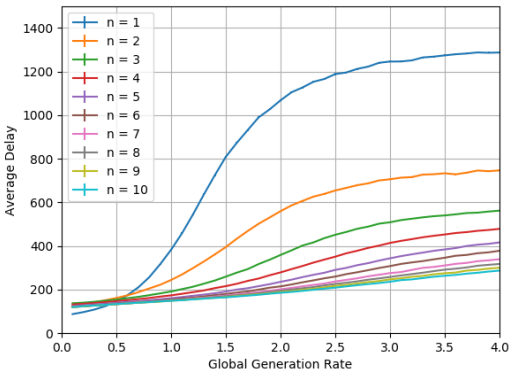

(a) High Density

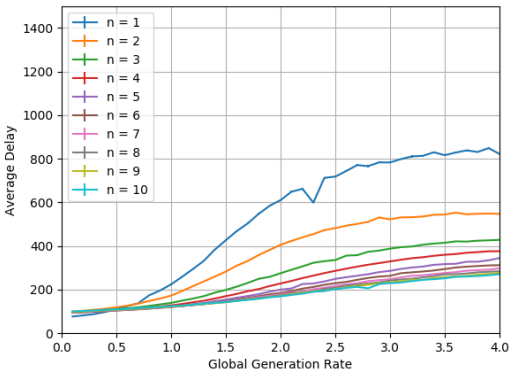

(b) Medium Density

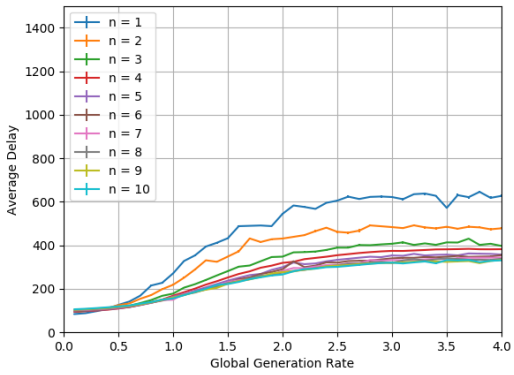

(c) Low Density

Figure 11: Average delay versus number of antennas ($m_{max} = 5$, $p = 10$)

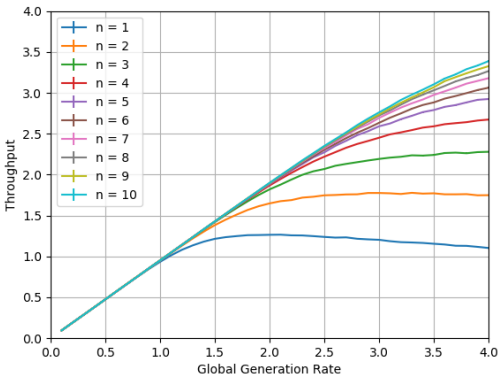

(a) High Density

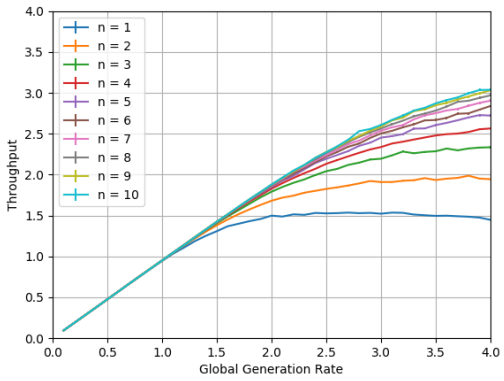

(b) Medium Density

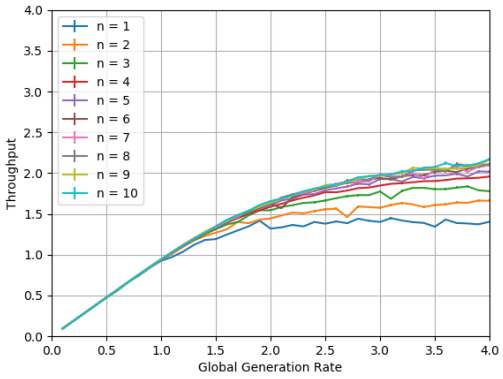

(c) Low Density

Figure 12: Throughput versus number of antennas ($m_{max} = 5$, $p = 10$)

B.5 Routing Metric

Figures 13 and 14 show the impact that the routing link weights have on the average delay and throughput, respectively, for all three tested densities. The "MinHop_Basic" term uses a link weight of one for every link where as the "NodeUtil_ETR_LinkRate" term is the one given by Equation 5.8. The min hop link weight actually provides better performance from the delay point of view, because on average traffic is relayed through fewer hops. However, delay is only measured for packets that successfully reach the final destination, so the measurement is biased. Figure 14 illustrates the throughput improvement from the more advanced link weight.

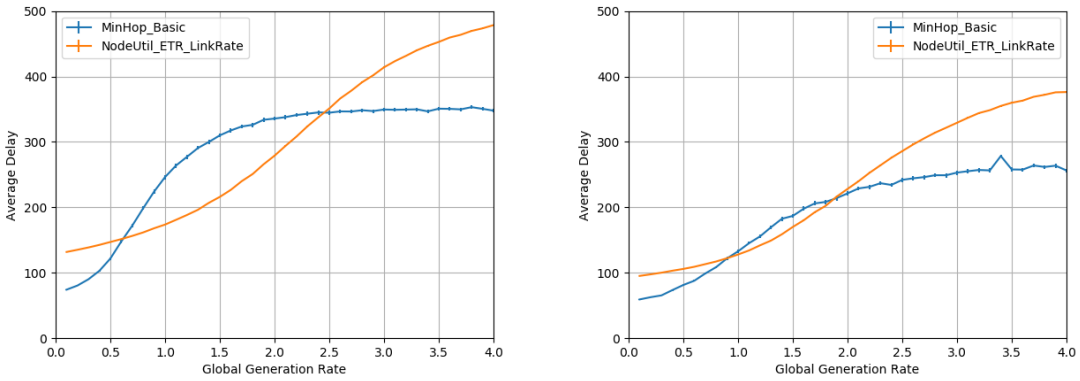

(a) High Density

(b) Medium Density

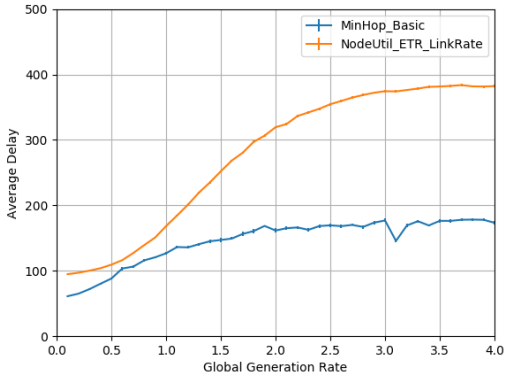

(c) Low Density

Figure 13: Average delay versus routing metric ($n = 4$, $m_{max} = 5$, $p = 10$)

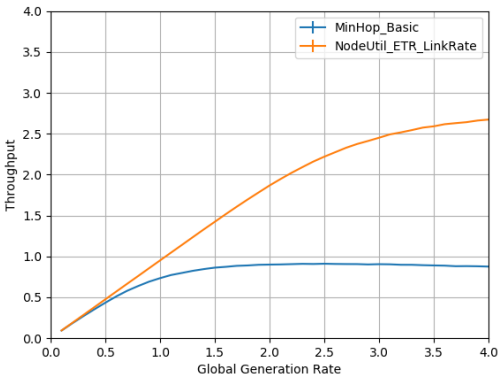

(a) High Density

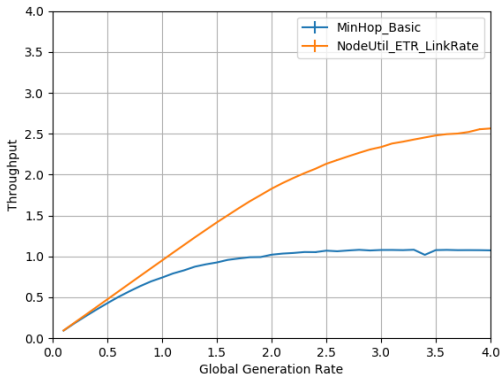

(b) Medium Density

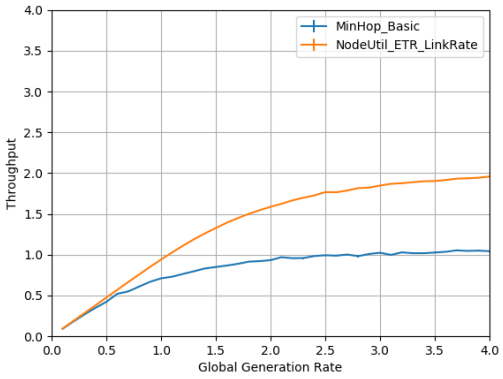

(c) Low Density

Figure 14: Throughput versus routing metric ($n = 4$, $m_{max} = 5$, $p = 10$)

Bibliography

[1] N Abramson. The Aloha system - Another approach for computer communications. *Proceedings of Fall Joint Computer Conference, AFIPS*, 37:281–285, 1970.

[2] Randeep Bhatia and Li Li. Throughput optimization of wireless mesh networks with MIMO links. *Proceedings - IEEE INFOCOM*, pages 2326–2330, 2007.

[3] Jack Brassil. Platforms for Advanced Wireless Research (PAWR): Establishing the PAWR Project Office (PPO) (PAWR/PPO).

[4] E. W. Dijkstra. A note on two problems in connexion with graphs. *Numerische Mathematik*, 1(1):269–271, 1959.

[5] Anthony Ephremides and Thuan V. Truong. Scheduling Broadcasts in Multihop Radio Networks. *IEEE Transactions on Communications*, 38(4):456–460, 1990.

[6] Ezzeldin Hamed, Hariharan Rahul, Mohammed A. Abdelghany, and Dina Katabi. Real-time distributed MIMO systems. *SIGCOMM 2016 - Proceedings of the 2016 ACM Conference on Special Interest Group on Data Communication*, pages 412–425, 2016.

[7] Joseph L. Hammond and Harlan B. Russell. Properties of a transmission assignment algorithm for multiple-hop packet radio networks. *IEEE Transactions on Wireless Communications*, 3(4):1048–1052, 2004.

[8] Phil Karn. MACA.pdf. *ARRI/CRRI Amateur Radio 9th Computer Networking Conference*, pages 134–140, 1990.

[9] Leonard Kleinrock and Fouad A. Tobagi. Packet Switching in Radio Channels: Part I Carrier Sense MultipleAccess Modes and Their Throughput Delay Characteristics. *IEEE Transactions on Communications*, COM-23(12), 1975.

[10] W. P. Lyui. *Design of a new operational structure for mobile radio networks*. Ph.d. dissertation, Clemson University, 1991.

[11] Xiaofu Ma, Qinghai Gao, Ji Wang, Vuk Marojevic, and Jeffrey H. Reed. Dynamic sounding for multi-user MIMO in wireless LANs. *IEEE Transactions on Consumer Electronics*, 63(2):135–144, 2017.

[12] Xiaoqi Qin, Xu Yuan, Yi Shi, Y Thomas Hou, Wenjing Lou, and Scott F Midkiff. Joint Flow Routing and DoF Allocation in Multihop MIMO Networks. *IEEE Transactions on Wireless Communications*, 15(3):1907–1922, 2016.

[13] S. Ramanathan. Unified framework and algorithm for channel assignment in wireless networks. *Wireless Networks*, 5(2):81–94, 1999.

[14] T. S. Rappaport. *Wireless Communications: Principles And Practive: Vol 2*. Prentice Hall, New Jersey, 1996.

[15] Injong Rhee, Ajit Warrier, Jeongki Min, and Lisong Xu. DRAND: Distributed randomized TDMA scheduling for wireless ad hoc networks. *IEEE Transactions on Mobile Computing*, 8(10):1384–1396, 2009.

[16] Gerhard Ringel. *Map Color Theorem*. Springer-Verlag, New York, 1st edition, 1974.

[17] Vikas Bollapragada Subrahmanya. *A Protocol to Recover the Unused Time Slot Assignments in Transmission Scheduling Protocols for Channel Access in Ad Hoc Networks (MSc Thesis)*. Ms thesis, Clemson University, 2016.

[18] Vikas Bollapragada Subrahmanya and Harlan B. Russell. Recovering the reserved transmission time slots in an ad hoc network with scheduled channel access. *Proceedings - IEEE Military Communications Conference MILCOM*, pages 67–72, 2016.

[19] Fouad A. Tobagi and Leonard 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*, 23(12):1417–1433, 1975.

[20] David Tse and Viswanath Pramod. *Fundamentals of wireless communication*, volume 1. 2005.

[21] Brian Wolf. *Cross-Layer Scheduling Protocols for Mobile Ad Hoc Networks Using Adaptive Direct-Sequence Spread-Spectrum Modulation*. Ph.d. dissertation, Clemson University, 2010.

[22] Brian J. Wolf, Joseph L. Hammond, Daniel L. Noneaker, and Harlan B. Russell. A protocol for construction of broadcast transmission schedules in mobile ad hoc networks. *IEEE Transactions on Wireless Communications*, 6(1):74–78, 2007.