

THE EFFICACY OF SOURCE RATE CONTROL IN ACHIEVING FAIRNESS IN WIRELESS MESH NETWORKS

by

Lily (Lei) Li

A thesis
presented to the University of Waterloo
in fulfilment of the
thesis requirement for the degree of
Master of Applied Science
in
Electrical and Computer Engineering

Waterloo, Ontario, Canada, 2007

©Lily (Lei) Li 2007

AUTHOR'S DECLARATION FOR ELECTRONIC SUBMISSION OF A DISSERTATION

I hereby declare that I am the sole author of this thesis. This is a true copy of the thesis, including any required final revisions, as accepted by my examiners.

I understand that my thesis may be made electronically available to the public.

ABSTRACT

The use of 802.11-based wireless mesh networks (WMNs) as an alternative network-backbone technology is growing rapidly. The primary advantages of this approach are ease of deployment and lower cost. However, such networks typically exhibit poor fairness properties, often starving nodes if they are too many hops distant from the gateway. Researchers have shown a growing interest in this problem in recent years. Many solutions proposed amount to some level of source rate control, either by policing directly at the source, or *via* TCP congestion control reacting to a gateway-enforced rate limit. However, there has been limited study on the effectiveness of source rate control.

In this thesis we first demonstrate that source rate control can only partially solve the fairness issue in 802.11-based WMNs, with some routers experiencing an undesirable degree of unfairness, which we call *structural unfairness*. We then identify the four necessary factors that cause structural unfairness. If we can eliminate or reduce any one of these conditions, we can eliminate or ameliorate the unfairness problem. We first investigate two techniques to improve 802.11 MAC scheduling: fixing the contention window and packet spacing at every router node, both means achievable with commodity 802.11 hardware. We show that the combination of these mechanisms provides a significant gain in fairness. We also perform case studies using another three techniques, channel re-assignment, routing changes, and careful router placement, to remove or reduce other necessary conditions. We demonstrate that these techniques, whenever applicable, can eliminate the unfairness problem entirely at times, or at least improve the situation.

ACKNOWLEDGEMENTS

First, I would like to acknowledge my supervisor, Professor Paul Ward, for his ongoing encouragement and guidance. This thesis will not have been made possible if not for him.

Second, I would like to acknowledge Szymon Chachulski and Alex T.K. Lau, who worked with me on the initial work on fairness in Wireless Mesh Network. This work has led me to my thesis work. Szymon's assistance on this work has been invaluable.

Third, I would like to thank the wireless research group led by Professor Paul Ward, in particular, Kamran Jamshaid and Evan Jones. I have been constantly inspired by the valuable discussions and collaborations with them.

Next, I would like to thank my co-workers, Peter Hansen, Jennifer Lee, Leo Tolisano and Jeff Haase for their understanding and support. Even their nagging has been extremely helpful. :)

Last, but not the least, I thank my parents, for putting up with all my frustration, and for providing me with a comfortable home.

CONTENTS

1	INTRODUCTION	1
1.1	FAIRNESS IN WMNS	4
1.2	CONTRIBUTIONS	6
1.3	THESIS ORGANIZATION	7
2	BACKGROUND AND RELATED WORK	9
2.1	WIRELESS TRANSMISSION BASICS	9
2.2	THE 802.11 MAC	10
2.3	MULTI-HOP WIRELESS NETWORKS	14
2.4	MODELING WMNS	15
2.4.1	MATH MODEL	17
2.5	FAIRNESS	19
2.5.1	FAIRNESS MODELS	19
2.5.2	FAIRNESS MECHANISMS	21
2.5.3	MAC-LAYER COORDINATION PROBLEM	24
2.5.4	SUMMARY	27
3	SOURCE RATE CONTROL	29
3.1	EXPERIMENT DESIGN	29
3.1.1	NS2	30
3.1.2	TOPOLOGY GENERATION	30
3.1.3	FLOW GENERATION	33
3.1.4	MEASUREMENTS	33

3.2	EXPERIMENTAL RESULTS	35
3.2.1	ACCURACY OF THE CLIQUE MODEL	39
3.2.2	EFFECTIVENESS OF SOURCE RATE CONTROL	42
3.2.3	SEVERITY OF UNFAIRNESS WITH SOURCE RATE CONTROL	43
3.3	SUMMARY	46
4	CAUSES OF STRUCTURAL UNFAIRNESS	49
4.1	TOPOLOGY REQUIREMENT	51
4.2	LINK UTILIZATION	54
4.2.1	EFFECT OF MULTI-HOP FLOWS	58
4.3	HOP DISTANCE	61
4.3.1	EFFECT OF ACK	65
4.3.2	EFFECT OF DOMINATING NODES	69
4.4	SUMMARY	73
5	IMPROVING MAC SCHEDULING	77
5.1	FIXING THE CONTENTION WINDOW	77
5.2	PER-NODE RATE LIMITING	80
5.3	EXPERIMENTAL EVALUATION	81
5.4	SUMMARY	85
6	ELIMINATING STRUCTURAL UNFAIRNESS CONDITIONS	87
6.1	IDENTIFY THE PROBLEMATIC LINKS	87
6.2	CASE STUDY 1: CHANNEL RE-ASSIGNMENT	88
6.3	CASE STUDY 2: RE-ROUTING	89
6.4	CASE STUDY 3: CAREFUL NODE PLACEMENT	93
6.5	SUMMARY	93

7	CONCLUSION AND FUTURE WORK	95
7.1	FUTURE WORK	95
	REFERENCES	97

FIGURES

1.1	DSL AND CABLE TO HOME AND OFFICE	1
1.2	CELLULAR SYSTEM	2
1.3	ETHERNET AND WLAN	3
1.4	WMN: COMMUNITY NETWORK	4
1.5	EXTEND WLAN WITH WMN ROUTER	5
2.1	WIRELESS TRANSMISSION	10
2.2	802.11 BASIC ACCESS MECHANISAM (CSMA/CA)	13
2.3	MULTI-HOP WIRELESS NETWORK	15
2.4	5-LINK WMN AND CORRESPONDING LINK-CONTENTION GRAPH	18
3.1	CHAIN TOPOLOGY	31
3.2	3X3 GRID	32
3.3	7-HOP CHAIN	36
3.4	THROUGHPUT VS. INPUT RATE FOR 7-HOP CHAIN, WITHOUT RTS/CTS	36
3.5	TRANSMISSION LINKS IN SAMPLE RANDOM TOPOLOGY	37
3.6	WITHIN INTERFERENCE RANGE FOR SAMPLE RANDOM TOPOLOGY . .	38
3.7	THROUGHPUT VS. INPUT RATE WITHOUT RTS/CTS	40
3.8	THROUGHPUT VS. INPUT RATE WITHOUT RTS/CTS	40
4.1	TOPOLOGY REQUIREMENTS	51
4.2	SINGLE-CHANNEL UNEVEN-DISTANCE CHAIN	55
4.3	ASYMMETRIC UNFAIRNESS IN SINGLE-CHANNEL CHAIN	55
4.4	MULTI-CHANNEL UNEVEN-DISTANCE CHAIN	57

4.5	SEVERE UNFAIRNESS IN MULTI-CHANNEL CHAIN	57
4.6	MULTI-HOP FLOW EFFECT IN SINGLE-CHANNEL CHAIN	60
4.7	MULTI-HOP FLOW EFFECT IN MULTI-CHANNEL CHAIN	60
4.8	SINGLE-CHANNEL EVEN-DISTANCE CHAIN	63
4.9	MULTI-CHANNEL EVEN-DISTANCE CHAIN	63
4.10	FAIRNESS IN SINGLE-CHANNEL CHAIN	63
4.11	FAIRNESS IN MULTI-CHANNEL CHAIN	67
4.12	MULTI-CHANNEL CHAIN WITH ACK INTERFERENCE	67
4.13	EFFECT OF ACK INTERFERENCE WITH POWER CAPTURE	68
4.14	EFFECT OF ACK INTERFERENCE WITHOUT POWER CAPTURE	68
4.15	7X7 GRID	71
4.16	SCENARIOS WITH/WITHOUT DOMINATING NODES	71
4.17	EFFECT OF DOMINATING NODES	72
4.18	WITHOUT DOMINATING NODES	72
4.19	DEMONSTRATION OF DOMINATING NODES	74
4.20	EFFECT OF REMOVING DOMINATING NODES	75
4.21	EFFECT OF ADDING DOMINATING NODES	75
5.1	IMPLEMENTATION OF PER-NODE RATE LIMITING IN NS2	82
5.2	FIXED $CW_{max} = 31$	83
5.3	PER-NODE RATE LIMITING	83
5.4	COMBINED EFFECT	86
5.5	FAIRNESS VS. CW_{max}	86
6.1	CHANNEL RE-ASSIGNMENT	90
6.2	REMOVED SU TOPOLOGY AFTER CHANNEL RE-ASSIGNMENT	90
6.3	ORIGINAL ROUTING WITH SU TOPOLOGIES	91

6.4	AFTER RE-ROUTING WITHOUT SU TOPOLOGIES	91
6.5	STRUCTURAL UNFAIRNESS WITH ORIGINAL ROUTING	92
6.6	FAIRNESS WITH ROUTING CHANGES	92

TABLES

3.1	30-NODE RANDOM TOPOLOGY ACCURACY RESULTS	42
3.2	FAIRNESS AND EFFICIENCY IN 30-NODE TOPOLOGY	43
3.3	SEVERITY OF SU IN 30-NODE RANDOM TOPOLOGY	46
3.4	SEVERITY OF SU IN 5X5 GRID TOPOLOGY	46
4.1	COLLISION RESULTS ON EFFECT OF ACK PACKETS	69
4.2	COLLISION RESULTS ON EFFECT OF DOMINATING NODES	74
5.1	IMPROVEMENT OF FIXED CW AND PER-NODE RATE LIMITING	84

1 INTRODUCTION

The demand for broadband data applications, such as Video On Demand (VOD) and image sharing, has been growing. As a result, people require broadband Internet access from everywhere. Broadband Internet access networks today fall into two main categories: last-mile broadband access and Local Area Network (LAN). Many technologies have been made available for Internet access networks. In the last-mile broadband access arena, DSL (Digital Subscriber Line) and cable have been widely used to deliver broadband Internet to homes and business offices, as shown in Figure 1.1. 3G has been developed to allow cellular users to have broadband service (Figure 1.2). In the LAN arena, Ethernet is used in buildings and Wireless LAN (WLAN) has been used both in buildings and Hot Spots to provide Internet access to mobile users (Figure 1.3).

However, these technologies require expensive up-front set-up cost, which hinders their deployment. DSL and cable modems require expensive rights-of-way, acting as a significant barrier-to-entry of competition. In particular, the cost has limited broadband-deployment in rural areas. WLAN typically uses Ethernet as its backbone network for all




Figure 1.1: DSL and Cable to Home and Office

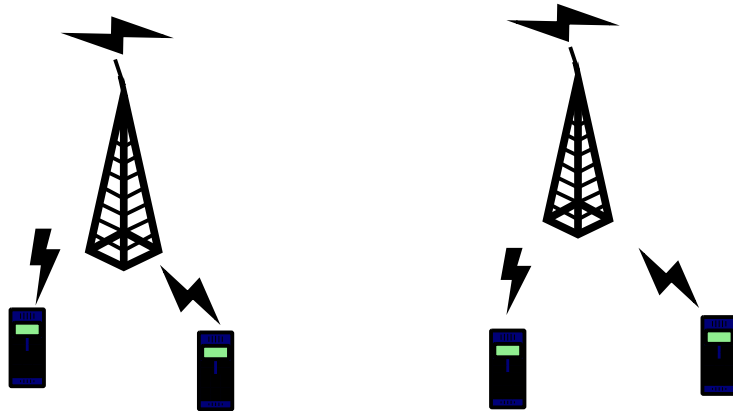


Figure 1.2: Cellular System

the Access Points (APs). Therefore, both Ethernet and WLAN require wire deployment, which is costly in many circumstances. Even though 3G is wireless, it requires licensed spectrum and antennas mounted on high towers, both of which are expensive.

Wireless mesh networks (WMNs; Figure 1.4), a type of multi-hop wireless network, have been proposed as an alternative technology for next-generation Internet access networks. Such networks consist of minimal-mobility mesh routers, together with their clients. The mesh routers communicate with each other over multi-hop wireless links, while the clients typically connect to their preferred router, either *via* wire or over a (possibly orthogonal) wireless channel.

As an example of why these networks are useful, consider an old hotel that is going to hold a conference. There is no existing network infrastructure in the building. However, most attendees have laptops and would like to have Internet access. A WMN can provide a temporary network with little cost. All that is required is to deploy wireless mesh routers in the building. The routers can be taken away after the conference if they are no longer required.

There are many scenarios in which a WMN can provide cheap solutions for broadband connections. In rural areas, where DSL and cable are too expensive to deploy, a




Figure 1.3: Ethernet and WLAN

community network can be set up using a WMN. Alston Cybermoor in Alston, UK [51] is an example of this model. In a wireless community network (Figure 1.4), one or few data connections (*e.g.*, a leased line) are set up as gateways to access the Internet. Then, mesh routers can be laid out, likely on the rooftop, to provide a multi-hop wireless network over the community.

WMNs can also be used to extend existing WLANs. In a WLAN, wireless clients access the network by sending packets to and receiving packets from the access points (APs). The APs are connected typically *via* Ethernet cables to the Internet. Therefore, to save on wiring cost, instead of laying more cables each time a new AP is added in a WLAN, one can simply place one or several WMN routers (Figure 1.5) and have these routers forward traffic to an existing AP over a wireless hop. This solution not only avoids the wiring cost, but also enables incremental deployment. In particular, it is extremely convenient if a temporary extension is needed. The router can be removed after use.




Figure 1.4: WMN: Community Network

1.1 FAIRNESS IN WMNS

There are a significant number of WMN testbeds and deployments today, including the MIT RoofNet [5], the Rice TAPs project [26], the Microsoft Mesh Connectivity Layer [31], various deployments of Nortel Mesh equipment [33], *etc.* A common feature of the majority of these testbeds is that they use commodity 802.11 hardware. This usage is not for lack of money, but rather a key design requirement for mesh networks to be successful, *viz.* they must be based on cheap commodity hardware if they are to be a plausible alternative network technology. Cost is the major barrier for the competition (*i.e.*, DSL, cable modems, 3G, and wireless LAN). If WMNs are as expensive, their case is not compelling. Given the ubiquity of 802.11 and Bluetooth hardware, this cost constraint leaves these as the only extant wireless options, with Bluetooth failing on account of limited range and bandwidth. This will remain the case for the foreseeable future. We therefore focus on 802.11-based WMNs, with all the constraints that implies.

As with other networks, one of the fundamental problems in WMNs is performance. A network is useful to a client only if the client gets a reasonable throughput and latency.




Figure 1.5: Extend WLAN with WMN router

However, without suitable mechanisms, WMNs can exhibit extreme network-layer unfairness, to the point of starving some mesh routers and their associated clients. In particular, it has been demonstrated that nodes close to the gateway receive substantially more throughput than those that are more hops away [8, 17, 24, 25]. Source rate control has been suggested to alleviate this fairness problem, either by direct policing at the source [12, 38, 45] or *via* TCP congestion control reacting to a gateway-enforced rate limit [21]. However, there has been no comprehensive study on the effectiveness of source rate control over many topologies and varied flows, and little work on the interaction between MAC-layer scheduling and network-layer fairness.

This thesis focuses on studying the efficacy of source rate control in achieving fairness in WMNs. Since prior research has shown that fairness only becomes an issue when the network is congested and there are unsatisfied user demands [25, 36], our work always assumes congested network. We use the well-known clique model [11, 12, 32] to estimate the carrying capacity of the network. We first show that in general the clique model is fairly accurate in predicting the fair-share rate. We then demonstrate that while source rate control is effective in many topologies, it fails to completely address the fairness problem. In particular, in more than half of the topologies examined, some nodes

experience more than 15% less throughput than others. We study and identify the factors that cause the problem and present various techniques to completely or substantially eliminate the problem. Keeping the key design requirement of WMNs in mind, our solutions meet the following goals: first, they do not require hardware changes; second, applying the techniques improves the fairness as well as maintaining or improving the aggregate throughput of the network vs. source rate limiting alone.

1.2 CONTRIBUTIONS

The main contributions of this work are as follows:

1. We show that in single-channel WMNs, network capacity calculated by the clique model is in general accurate. Most flows in WMNs can achieve the throughput calculated using the clique model.
2. We demonstrate that source rate control cannot completely achieve fairness. We quantify the severity of the problem, and identify the causes of the problem.
3. We propose two techniques to improve 802.11 MAC scheduling: fixed contention window and packet spacing at each router. We show that the combination of these mechanisms provides a significant gain in fairness without sacrificing aggregate throughput.
4. We perform case studies using three techniques, channel re-assignment, routing changes, and careful router placement, to remove the necessary factors identified. We show that these techniques completely eliminate the problem.

1.3 THESIS ORGANIZATION

The remainder of this thesis is organized as follows. In Chapter 2 we describe the 802.11 protocol and WMNs, and review the previous approaches to solving the fairness problems and their limitations. In Chapter 3 we demonstrate that source rate control, while it effectively alleviates the fairness problem in many scenarios, fails to solve the problem in a number of common cases. We describe this problem as structural unfairness. We study in depth the causes of the structural-unfairness problem in Chapter 4. We propose in Chapter 5 two techniques that improve 802.11 MAC scheduling, thereby substantially ameliorating the problem. In Chapter 6 we show three additional techniques with case studies that remove the required conditions for the problem to occur and demonstrate that they completely eliminate the problem. Finally, our conclusions and future work are presented in Chapter 7.

2 BACKGROUND AND RELATED WORK

In this chapter we first discuss the basics of wireless transmission and the essentials of the 802.11 MAC. We then define our formal model of WMNs. Next, we review the fairness models and mechanisms in prior research. Researchers have demonstrated that the fairness mechanisms designed for wireline networks, such as weighted fair queuing (WFQ), are not suitable for wireless networks. We observe that the fairness mechanisms proposed for WMNs are mostly doing rate control in one way or the other. Finally, we review the related work on MAC-related issues.

2.1 WIRELESS TRANSMISSION BASICS

Wireless networks are a type of network where interconnections between nodes are implemented without the use of wires or fibers. We model wireless transmission using two ranges: transmission range and interference range. When the receiver is within transmission range of the sender, the receiver can receive the signal from the sender with sufficient fidelity to ensure correct decoding of the message. When the receiver is out of transmission range of the sender but within interference range of it, the receiver can sense but cannot receive the signal. Moreover, the signal from the sender could potentially interfere with the reception of a message by the receiver from another sender. If the receiver is out of interference range of the sender, the receiver cannot detect the signal from the sender and the sender cannot interfere with the receiver, either [44].

This leads to the *hidden-terminal* and *exposed-terminal* problems [40, 44]. Consider the scenario depicted in Figure 2.1 (a). Nodes A and C are both within transmission range of B , but are out of interference range of each other. Therefore, node C will sense




Figure 2.1: Wireless Transmission

the medium idle when node A transmits to B , and may try to transmit to B at the same time, interfering with the packet from A to B , and causing a collision. The problem is that a collision in wireless transmission is perceived by the receiver, but not by the sender. Thus, a sender could fail to detect other senders that can cause potential collisions in the receiver. This is called the hidden-terminal problem.

The exposed-terminal problem is the opposite of the hidden-terminal problem. Consider the scenario in Figure 2.1 (b). Since the two senders B and C are within interference range of each other, they will attempt to avoid collision by not transmitting when the other is transmitting. Therefore, when B is transmitting to A , C will not transmit to D , and vice versa. However, since B 's receiver, A , is out of range of sender C , and C 's receiver, D , is out of range of sender B , these two transmissions can in fact succeed at the same time. This problem is called the exposed-terminal problem.

Thus the hidden-terminal problem causes collisions, while the exposed-terminal problem causes unnecessary delay, both leading to inefficient usage of spectrum.

2.2 THE 802.11 MAC

There are many Medium Access Control (MAC) mechanisms designed to deliver data reliably over the noisy and unreliable wireless media (*e.g.*, Aloha [44], 802.11 [18],

CDMA [15]). Based on how the resources are shared among multiple nodes, these and the other protocols can be divided into FDMA (Frequency Division Multiple Access), TDMA (Time Division Multiple Access), CDMA (Code Division Multiple Access) [40] and CSMA (Carrier Sense Multiple Access), among others. The CSMA protocol is designed for random access networks, where central control is not applicable. In CSMA protocols, nodes listen to the carrier and only transmit when the medium is free. There are two versions of CSMA protocols: p-persistent and nonpersistent [40, 44].

In a p-persistent CSMA protocol, if the medium is idle, the node will transmit with a probability p and defer with a probability $q = 1 - p$. In case it defers, it waits until the next time slot. If the medium is still idle in the next slot, the node will again transmit with a probability p . The process will be repeated until either the node transmits or another node transmits. In the latter case, the non-transmitting node considers it a collision, waiting for a random time period and then starting the process all over again. A 1-persistent CSMA protocol is a special case of p-persistent CSMA, where p is 1. In this case, if a node becomes ready when another node is transmitting it will transmit immediately after the other node has completed its transmission. Therefore, if two nodes have to transmit, they will both wait until the medium becomes free and transmit simultaneously, which leads to a collision. If p is smaller than 1, there will be fewer collision in this case. However, it may result in longer delay.

In a nonpersistent CSMA protocol, if the node senses the medium free, it starts transmitting. However, if the medium is busy, the node does not wait for the medium to become idle. Instead, it waits for a random delay and then senses the medium again, transmitting if the medium is idle. This protocol leads to fewer collisions than the 1-persistent protocol and lower latency than the p-persistent protocol.

CSMA with collision avoidance (CSMA/CA) is used as one of the access schemes in the IEEE 802.11 standard. The 802.11 MAC provides two operation modes: Distributed

Coordination Function (DCF) and Point Coordination Function (PCF). The PCF mode is optional and designed to be used in infrastructure mode to provide contention-free medium access *via* central control, while DCF mode is mandatory and does not require central control. In DCF mode, 802.11 uses nonpersistent CSMA/CA [10, 40, 44].

The 802.11 MAC defines three important parameters to control the waiting time before medium access: DIFS (DCF Inter-Frame Spacing), PIFS (PCF Inter-Frame Spacing) and SIFS (Short Inter-Frame Spacing). DIFS is the longest waiting time and SIFS is the shortest. Therefore, nodes that use DIFS as their waiting time have the lowest channel-access priority and those that use SIFS have the highest channel access priority.

We now illustrate the CSMA/CA mechanism in DCF operation in detail. Figure 2.2 shows how a sender gets access to the channel. A node wishing to transmit listens to the wireless channel first to check whether there are any other transmissions on the same channel. If the medium is idle for DIFS, the node can transmit immediately. Otherwise, the node enters into a contention stage by choosing a uniformly distributed random number that represents the backoff time that the sender has to wait before attempting to transmit again. The random number is chosen within a range called the Contention Window (CW). The node counts down the backoff counter by one each time the medium is sensed idle for one slot time (*e.g.*, $20\mu\text{s}$ in 802.11b). If the medium is sensed busy during a slot, the counter is not decremented. When the random backoff counter reaches zero, the node obtains the channel and sends its data immediately.

For each unicast packet, 802.11 requires the receiver to send back an acknowledgement (ACK) after successfully receiving that packet. The ACK is sent at SIFS time after the receiver gets the packet regardless if the channel is busy or idle at the time. If the sender does not receive an ACK from the receiver within a timeout period, the sender presumes a collision has occurred. It will double the size of the CW (subject to a maximum), increment the retry counter, and then it will attempt to retransmit the packet after




Figure 2.2: 802.11 Basic Access Mechanism (CSMA/CA)

going through the contention stage again, as described above. After a maximum number of retries it gives up the transmission attempt and reports the failure to the network layer. By default, the minimum CW size is 32 and the maximum is 1024. This algorithm is called Binary Exponential Backoff (BEB). It is an attempt to avoid collisions in the presence of an unknown number of senders.

The 802.11 MAC defines an optional mechanism to deal with the hidden-terminal problem. A sender sends a request-to-send (RTS) control packet first, which identifies the receiver and the time needed for the whole transmission, including the data transmission and the ACK. Nodes that receive a RTS packet will adjust the earliest time when they should try to access the medium again. The receiver, upon receiving the RTS packet, will send back a clear-to-send (CTS) packet after SIFS, if it also senses the medium idle. The CTS packet also contains a duration field. Nodes which receive the CTS packet will adjust their time-to-transmit accordingly. If the sender does not receive the CTS packet within a timeout period, it perceives a collision and invokes BEB. This mechanism is designed to reduce the likelihood of the hidden-terminal problem in single-hop WLAN networks.

2.3 MULTI-HOP WIRELESS NETWORKS

Multi-hop wireless networks are an extension of single-hop wireless networks. In such networks, two nodes that are not within transmission range of each other communicate *via* intermediate nodes. Consider the wireless network in Figure 2.3. If node *A* needs to communicate with node *D*, it has to go through three hops *via* nodes *B* and *C*. This type of network is called a multi-hop wireless network because the traffic is typically relayed over multiple wireless hops.

This type of network introduces three classes of contention that do not exist in single-




Figure 2.3: Multi-hop Wireless Network

hop networks.

1. Since a network flow traverses multiple hops, the maximum throughput of a flow is not determined by how much bandwidth it can get on a single link. Instead, the throughput is limited by the *bottleneck* link on its route that allocates the least bandwidth to it.
2. A network flow contends with another network flow if their routes intersect. This means, in particular, that traffic originating at a node competes with traffic forwarded by the node.
3. A flow contends with itself over multiple hops because each hop may use the same spectrum.

2.4 MODELING WMNS

We now present our model of WMNs. A WMN is composed of minimal-mobility mesh routers, together with their clients. The mesh routers communicate with each other over multi-hop wireless links, while the clients typically connect to their preferred router.

While some WMNs allow for mesh clients to participate in the routing and forwarding of messages, for the purpose of this thesis we will restrict ourselves to clients that merely send or receive their own messages and do no forwarding. We assume that clients

send messages to, or receive messages from, mesh routers either *via* a wired connection, or over an orthogonal wireless channel. Thus, in this model a client treats the router to which it is connected as either its network gateway (wired) or access point (wireless). In this manner clients can operate entirely unaltered from their normal operation using a standard LAN or WLAN. Clients may or may not be mobile and/or run on battery power.

In contrast to mesh clients, we consider mesh routers to be powered from the electricity grid. This implies that they are neither power constrained nor mobile. This in turn implies that the topology of a WMN is mostly static. Topological changes are caused not by router mobility but by incremental deployment and node failure. Such topological changes should be rare relative to changes in network traffic. The router nodes communicate over a single channel or over multiple channels, forwarding messages as necessary. Not all router nodes are within range of each other.

We assume that routing is effectively static, based on the fact that in WMNs nodes are stationary, and likely quite reliable. What we mean by effectively static is that changes in routing will be significantly less frequent than changes in flow activity.

In general clients do not communicate with each other, but with servers, either local or remote. In either instance, messages will be sent to/from a specific router node or nodes, which we will designate to be the gateway node(s). For simplification purposes, we will presume there is only one such gateway node per mesh network, with all traffic moving to/from that gateway. This is almost certainly not true, but, given static routing, for each node there will be a single gateway. We thus partition a multi-gateway WMN into disjoint WMNs, each of which has a single gateway. While there may be interference between the resulting set of WMNs, this is a problem that must already be dealt with insofar as there may be interference from any number of other sources.

We thus have a system model of WMNs that is essentially akin to a WLAN in which the distribution network is a multi-hop wireless network, either single or multi-channel.

This is consistent with systems such as those developed by Nortel [33], and a generalization of the Transit-Access-Points (TAP) model [12], where the mesh router in our model is akin to the TAP in their model. As with the TAP approach which focuses on TAP-aggregate fairness, the focus of this thesis will be solely on the mesh routers, and ensuring their fair access to the network. In the remainder of the thesis, we use “mesh router” and “node” interchangeably. We use “flow” to refer to “network-layer flow” unless explicitly specified otherwise. If a flow originates at the gateway, and terminates at a mesh router, we term it a *downstream flow*. Conversely, if the flow originates in the WMN, and terminates in the gateway, we term it an *upstream flow*. It then follows that if there are N mesh routers there can be a maximum of $2N$ flows, N upstream flows and N downstream flows.

2.4.1 MATH MODEL

Given these assumptions, a WMN can be modeled as a connectivity graph $G = (V, E)$, where V is the set of vertices that represent the mesh routers, and E is the set of edges that represent the links between mesh routers. An edge exists between two vertices if their corresponding routers are within transmission range and share the same wireless channel.

We use the clique model to determine the capacity of such a network [11, 19, 29, 32]. In this model, we capture contention using a link-contention graph $G_C = (V_C, E_C)$, where V_C is the set of all links in the connectivity graph, and $(u, v) \in E_C$ iff links u and v contend. We define two links as contending if either node from one link is within interference range of either node of the other link. Links u and v cannot be active at the same time if they have an edge between them. In the link-contention graph, this implies that at any given time only one link may be active in any *clique* of the contention graph. A maximal clique can then be used to compute an upper bound on the maximum capacity




Figure 2.4: 5-link WMN and corresponding link-contention graph

of a network [19], as follows.

The channel resource is viewed as being divided among the cliques in this graph. Let $C(l)$ be the throughput that link l carries. Define $B(u)$ to be the available bandwidth in each distinct region u (*i.e.*, in each clique). Since all links in a clique contend with each other, only one link in the clique can be active at any instant. We can thus define the channel-resource constraints of the clique model as:

$$\sum_{i:i \text{ in clique } u} C(l_i) \leq B(u) \quad (2.1)$$

Note that if each wireless router transmits at the same rate, per our system definition, the value of $B(u)$ can be reasonably approximated as the throughput that can be achieved at the MAC layer in a one-hop network [23]. If routers transmit at different rates, a weighted contention graph would be needed, but that is not the focus of this work.

Figure 2.4 (a) shows a simple chain topology with nodes 200m apart. Assuming 802.11b, which has a transmission range of 250m and interference range of 550m, then in this topology the nodes that are two hops apart are within interference range. For

example, node 2 is within interference range of node 4. Therefore, link l_1 contends with link l_4 because node 2 is an end point of link l_1 and node 4 is an end point of link l_4 . Figure 2.4 (b) shows the corresponding link-contention graph. We can see that there are two maximal cliques in this graph, u_a and u_b , each containing four links. The channel-resource constraints for these two cliques can then be written as:

$$C(l_1) + C(l_2) + C(l_3) + C(l_4) \leq B(U_a)$$

$$C(l_2) + C(l_3) + C(l_4) + C(l_5) \leq B(U_b)$$

2.5 FAIRNESS

There has been extensive research on fairness in networks, both wired and wireless, since it is an important network property that must be addressed before a network can be used effectively. There are two aspects to achieving fairness: one is the definition of fairness (*i.e.*, the fairness model); the second is how the fairness is achieved (*i.e.*, the fairness mechanisms). We first discuss various fairness models, then we describe existing and proposed fairness mechanisms. In the end, we review the work related to MAC-layer coordination problems.

2.5.1 FAIRNESS MODELS

A fairness model defines a formal objective to be used as a fairness and performance target in a network [12]. It aims at allocating shared network resources *fairly*.

There are two aspects to a fairness model: resource shared and fairness policy. A fairness model could use either bandwidth or time as the basic network resource to be

fairly shared. If bandwidth is considered as the shared resource, throughput is used as the fairness measure. Therefore, the flows that use links that have better channel conditions and therefore have higher bit rates will be penalized by the flows with low-bit-rate links. In comparison, when channel time is to be fairly shared, if a flow is more cost-effective, it can get higher throughput [12].

The most popular fairness policies are *max-min* fairness and *proportional* fairness. Max-min fairness aims at fair throughputs among the nodes. Let x_i denote the rate allocated to node i . An allocation of rates is max-min fair iff $\forall_{i,j}$ no rate x_i can increase without a lesser rate, $x_j < x_i$, being reduced [6]. When there is only one bottleneck link in the network, max-min fairness becomes absolute fairness, in which case, $\forall_{i,j} x_i = x_j$. Proportional fairness is developed to take into consideration the usage of the network resources, and aims at maximizing an objective function representing the overall utility of the flows in progress [6, 30]. Formally, an allocation of rates x is proportionally fair iff for any other feasible allocation x^* , the aggregate of the proportional change is 0 or negative:

$$\sum_i \frac{x_i^* - x_i}{x_i} \leq 0 \quad (2.2)$$

The fairness granularity can be categorized into per-packet, per-node and per-flow. 802.11 aims at providing fair access to the wireless medium on a per-packet basis. In the ideal case, all nodes get the same chance of channel access to transmit one packet. Therefore, nodes that transmit longer packets tend to get more bandwidth than those that transmit shorter packets. With per-node fairness, all the flows of a single node share the fair share that is allocated to that node. This is the approach taken by the TAP system [12]. With per-flow fairness, each node could have multiple flows and each flow is treated individually and gets its own fair share. This is the approach taken by TCP [7].

Fairness can be achieved in the MAC layer or the network layer. MAC-layer fairness mechanisms aim at providing fair access in a single-hop network. Our work in WMNs aims at per-node network fairness over multiple hops. However, the fairness model is not the main focus of this thesis, but a means to study various fairness mechanisms. For this reason, we deliberately restrict our analysis to max-min and absolute TAP-aggregated network-flow fairness using throughput as our metric.

2.5.2 FAIRNESS MECHANISMS

In wired networks, variations of fair-queuing schemes are often used to achieve fairness [2, 14, 43, 53]. However, these mechanisms do not extend easily to shared wireless networks [38]. Jun and Sichitiu [24] explore various fair-queuing schemes employed at each hop in multi-hop wireless networks. Assigning a queue for each network flow can provide fairer access for each network flow over every hop. However, this scheme assumes perfect MAC-layer scheduling, since if the MAC-layer scheduling cannot guarantee the access each hop requires to traverse the network flows, the node that does not get enough channel access will build up its buffer and finally drop packets when the buffer reaches its limit, in which case the flows through this node achieve lower throughput. Gevros *et al.* [36] point out that fairness only becomes an issue when the network is congested and there are unsatisfied user demands. Jun and Sichitiu [25] also demonstrate that when the demands exceed network capacity, congestion causes serious unfairness among multi-hop flows. However, fair queuing cannot slow down the traffic and cannot deal with congestion gracefully.

There is also much work in achieving MAC-layer fairness in wireless networks [28, 29, 32, 35, 37, 42, 46]. While important, it has been shown that fairness over single-hop flows does not lead to multi-hop network-layer fairness [22, 32]. Therefore, with the use of multi-hop wireless networks becoming more prevalent, more attention has been given

to fairness in such networks.

Since fairness only becomes an issue when there is congestion in the network, most of the work concentrates on providing fairness in a congested network by using congestion-control techniques. The mechanisms proposed by prior works mainly include: hop-by-hop flow control [17, 47, 52, 54] and source rate control [8, 9, 38]. Hop-by-hop flow control adjusts the data transmission rate at each hop (including both the traffic originated by the node and the forwarded traffic for other nodes) when congestion is detected. Source rate control regulates the source of each flow so as not to inject new packets faster than the enforced rate. Sometimes both mechanisms are combined.

Woo and Culler [52] perform hop-by-hop flow control by adapting the p-persistence MAC scheme proposed by Nandagopal *et al.* [32]. The scheme is tailored for the multi-hop setting by separating the originating flow from the relayed flows into two outbound flows managed independently by the p-persistence scheme. It examines an AIMD rate-adjustment strategy in which the additive increase is proportional to the number of descendants of a node, and multiplicative decrease is performed whenever a node detects that its parent is unable to forward its traffic. However, this mechanism is only evaluated over one particular tree topology.

Both CODA [47] and Fusion [17] provide congestion-mitigation strategies. CODA senses both channel occupancy and queue length for measuring congestion levels, while Fusion uses only queue length. Both CODA and Fusion combine a hop-by-hop flow-control mechanism with a source-rate-limiting mechanism. The hop-by-hop flow control can provide fast feedback on congestion, while source rate control benefits fairness significantly, as Fusion demonstrated. To improve the network efficiency, Fusion also developed a prioritized medium access layer that allows congestion at local nodes to drain quickly. CODA's focus is congestion control, not fairness. Therefore, their evaluation only focuses on the aggregate network efficiency. Fusion suggests that their mechanism

improves both fairness and efficiency of the network. However, they only evaluated one tree topology using their testbed.

The most recent works are the Aggregate Fairness Algorithm (AFA) [8] and Interference-aware Fair Rate Control (IFRC) [38]. Both implement rate-control mechanisms. AFA proposes a localized algorithm for Aggregate Fairness that can be applied over any routing algorithms. It limits rate both at the source and the forwarding nodes. It observes that the main forms of congestion are caused by radio collision and buffer overflow. The basic idea is for a forwarding node to estimate the number of flows coming from each neighbor (also called the link's aggregate flow weight) and allocate bandwidth proportional to that number. The congestion information is exchanged in the network by piggybacking the buffer state in the frame header. The actual rate from an upstream link should be proportional to the link's aggregate flow weight. The rate limit enforced on upstream links by congested nodes is proportional to the link's aggregate flow weight. AFA is evaluated by simulating a random topology that consists of 500 sensors within a $1000 \times 1000 \text{ m}^2$ area and demonstrating that AFA is able to achieve much better end-to-end fairness than other schemes.

IFRC [38] applies rate control to each node based initially on estimation. The rate constraint is that the sending rate of a flow be no greater than the sending rate of the most congested node in the neighborhood. IFRC is similar to Woo and Culler's approach [52] in that it also implements an AIMD control law. However, instead of explicit hop-by-hop backpressure, nodes in IFRC exchange congestion indicators and converge in a distributed fashion to a fair and efficient rate. IFRC also applies aggressive rate reduction to avoid dropping packets. While IFRC appears effective as a fairness mechanism, it is only evaluated with one topology on a 40-node testbed.

Tang *et al.* [45] and Lee [27] have developed algorithms to compute the max-min fair-share rate in multi-channel wireless networks. The former algorithm is only evaluated

over several topologies, while the latter focuses on the accuracy of the computation of fair-share rate.

Jamshaid *et al.* [21] propose an implicit source-rate-control algorithm to achieve fairness. In this scheme, rate limiting is explicitly applied only on the gateway node. Other nodes do not have to be aware of the algorithm. When a flow transmits faster than the fair-share rate, the gateway node indicates congestion by throttling the traffic of that flow. This scheme relies on the source node limiting its own traffic when detecting packet losses. TCP traffic is the primary traffic of this type. When the source node does not get TCP ACKs for the dropped packets, the TCP protocol will act on the feedback and invoke the congestion-control algorithm to limit the source rate. This work evaluates limited chain, grid and random topologies.

2.5.3 MAC-LAYER COORDINATION PROBLEM

Prior work suggests that source rate control is sufficient to achieve network-layer fairness in a WMN. However, Rao and Stoica [39] demonstrate that in certain situations, source rate control is not enough. They perform experiments on a six-node wireless testbed using 802.11a radios. Each experiment consists of two simultaneous 1-hop UDP flows. They observe that when the interference between flows is asymmetric, rate-limiting the flow with the higher throughput does not substantially improve the throughput of the other flow. These asymmetric interactions between nodes are variations of the hidden-terminal and exposed-terminal problems. Rao and Stoica suggest that there are two requirements to achieve fairness: (1) each node should only access the medium for its fair share of time; (2) no other nodes should interfere with the node that is transmitting. They state that source rate control can only satisfy the first requirement and the 802.11 MAC fails to satisfy the second requirement. To solve the problem, they propose an Overlay MAC Layer (OML), loosely synchronizing clocks among nodes and allocating

time slots to no more than one node in the same interference region. This solution is a form of TDMA over 802.11, with large time slots, suggesting it will only be useful for networks where all nodes wish to do bulk data transfer. In addition, it is very challenging to estimate interference accurately; their simple two-hop neighborhood interference assumption can easily underestimate or overestimate interference.

Garetto *et al.* [13] study the generic coordination problem in CSMA-based single-channel wireless networks. They identify four categories of packet loss due to the MAC-coordination problem. The first is loss due to collisions between coordinated stations. This happens only if two contending nodes finish their random backoff at the same time. This is rare. The second category is loss due to an asymmetric view of channel status among nodes. The third is loss due to near hidden terminals, in which the receivers of the contending flows are each within interference range of the other receiver's sender. This decreases the throughput of both flows. The last category is loss due to far hidden terminals, in which the two competing flows are only connected by the receivers and the control packet of one flow can interfere with the data packet of the other. They realize the difference between congestion-induced collision and interference-induced collision by separating "topology-induced imbalance" from "MAC starvation." They develop a new model that captures the CSMA-induced coordination problem. We show in Chapter 4 that their study is similar to ours. However, their work only studies single-hop flows in multi-hop networks, which does not capture the characteristics of multi-hop flows. Therefore, the causes they identify for the unfairness of single-hop flows cannot cover all cases of unfairness in multi-hop flows. In addition, they do not rate-limit the network. Rather, they over-drive it. Our work is intended to determine the limitations of source rate control even when there is sufficient wireless capacity.

The Asynchronous Multi-channel Coordination Protocol (AMCP) [41] addresses the same problem discussed above in multi-channel networks by using a dedicated control

channel. Nodes contend in the control channel and use RTS/CTS to piggy-back channel utilization information in the network and negotiate data channels. This method wastes a whole channel for control packets which need little bandwidth. It is only evaluated with a few networks using 12 channels. It is not clear how efficient this method is in an 802.11b network, where there are only 3 orthogonal channels.

Wang and Garcia-Luna-Aceves [49] study single-hop flows in 802.11-based multi-hop wireless networks both analytically and experimentally. They show that the 802.11 MAC cannot guarantee collision-free transmissions of data packets because of variations of the hidden-terminal problem. They propose to use a fixed CW instead of BEB, so that the node that is shut off by a hidden terminal will not build up its CW value, getting fewer chances to compete for the channel, leading to severe unfairness. However, their analysis shows that while a fixed CW provides fairness, it also yields much worse throughput. The problem, as we show in Chapter 4, is that the 802.11 MAC tries to address both congestion-induced collision and non-congestion-induced collision at the same time. A fixed CW increases collisions when congestion occurs. Moreover, their work only studies the 802.11 MAC with RTS/CTS turned on. In Chapter 3 we show that performance is better without RTS/CTS.

Heusse *et al.* [16] propose an access method called “idle sense” to improve throughput and fairness in WLANs where all the nodes are within transmission range of each other. Their method also does not use BEB. Instead, they propose to dynamically adjust the CW value by comparing the average number of idle slots between transmission attempts with the optimal value. They observe that packet loss is not only due to contention, it can also be because of poor link quality. Therefore, the BEB mechanism in 802.11 can be very inefficient if most packet losses are due to radio errors. Link error is not the current focus in our work; rather, we restrict ourselves to dealing with congestion-induced loss and non-congestion-induced loss. We assume that all packet losses are due

to contention. We will show in Chapter 4 that BEB is inefficient even when the link error is not a factor. We adopted their idea of using the same CW values for all nodes in one of our solutions, though in our solution it is strictly fixed.

2.5.4 SUMMARY

Prior work has proposed source rate control as a mechanism for achieving per-node fair throughput in WMNs. It has also been observed that MAC-layer coordination causes unfairness among competing flows. However, there has been no extensive study evaluating how effective source rate control is in achieving fairness over many topologies and varied flows. In this thesis, we study the efficacy of the source-rate-control mechanism, showing where it works, where it fails, and why; we also propose and evaluate several techniques to enable fairness in cases where source rate control is insufficient.

3 SOURCE RATE CONTROL

In this chapter we study the effectiveness of source rate control in achieving fairness in WMNs by running simulations using the Network Simulator, ns2 [1]. We want to answer three questions:

1. How well does the source-rate-control mechanism achieve fairness among network flows?
2. Is the clique model accurate in estimating network capacity?
3. If the clique model is mostly accurate, how often and severely do some flows not achieve the throughput computed by the clique model?

We first describe the simulation setup and experiment measurements we use to evaluate the simulation results. Next, we conduct experiments and that show that source rate control can alleviate the fairness problem in many cases, but cannot completely achieve fairness. We quantify the severity of its limitation.

3.1 EXPERIMENT DESIGN

We describe the experiment design in the following four categories: ns2 setup, generation of topologies, flow generation, and measurements.

3.1.1 NS2

Ns2 has implemented various radio propagation models, MAC protocols and routing protocols. The radio propagation models are used to predict the received signal power of each packet. A packet is received successfully by the receiver when the power of the signal is higher than the receiving threshold at the receiver. Ns2 has implemented a free-space model, a two-ray ground model and a shadowing model. The free-space model and the two-ray ground model assume the transmission area is an ideal circle, while the shadowing model extends it to a probabilistic model that incorporates multipath fading effects. For the purpose of our study, we want to avoid the complication introduced by a probabilistic model. On the other hand, the free-space model cannot predict the signal power well for long distances because it assumes a simple line-of-sight path only. Therefore, we choose the two-ray ground model that considers both the direct path and a ground-reflection path, and thus can better predict the power when the distance is large. The transmission range is 250 meters, and the interference range is 550 meters. Ns2 implements a simplified “power-capture” model, in which, when multiple packets arrive at the receiver at the same time, only the first packet can be captured if its signal power is higher than any of the other packets by at least 10dB. We use the 802.11 MAC protocol, and set the base transmission rate to 1 Mbps. We limit the queue (i.e., IFQ) size at each node to 20 packets. In the ns2 simulator, we implemented a static shortest-path routing protocol and the clique model (collision-domain model in the case of multi-channel WMNs) to estimate network capacity.

3.1.2 TOPOLOGY GENERATION

We run our simulations over chains, grids and random topologies. In this section, we describe how these topologies are generated . In all the topologies, we always assign




Figure 3.1: Chain Topology

node number 0 as the gateway. The gateway may be placed either in the corner or in the middle of a chain or a grid.

Chain In our chain topologies, as shown in Figure 3.1, D_1 denotes the one-hop distance, D_2 the two-hop distance, and D_3 the three-hop distance. The constraints are as follows:

$$\begin{aligned} D_1 &\leq 250 \text{ meters} \\ D_2 &> 250 \text{ meters} \end{aligned} \quad (3.1)$$

This implies that $D_2 \leq 500 \text{ meters}$; *i.e.*, nodes that are two hops away are within interference range of each other. In the common chain topologies studied in other literature, the distances between adjacent nodes are often equal, and nodes that are three hops apart are out of interference range. This can be represented as:

$$\begin{aligned} D_3 &> 550 \text{ meters} \\ D_2 &= 2D_1 \\ D_3 &= 3D_1 \end{aligned} \quad (3.2)$$

This results in $183 \text{ meters} \leq D_1 \leq 250 \text{ meters}$. Therefore, in our simulations with common chains, we use equal-distance chains with $D_1 = 200 \text{ meters}$.




Figure 3.2: 3x3 Grid

Grid In grid topologies, the distances between adjacent nodes are the same. For example, in the 3x3 grid shown in Figure 3.2, the distances between nodes 1 and 2, or 0 and 1, or 0 and 3 are all equal. If we denote the distance to be X , grid topologies are generated with the following constraints:

$$\begin{aligned}
 X &\leq 250 \text{ meters} \\
 2X &> 250 \text{ meters} \\
 3X &> 550 \text{ meters}
 \end{aligned}
 \tag{3.3}$$

This results in $183 \text{ meters} \leq X \leq 250 \text{ meters}$.

Random Topology We generate random topologies by placing nodes randomly in an area of $1000 \times 1000 \text{ m}^2$ and ensuring that the network is connected; *i.e.*, every node in the network is within transmission range of at least one other node. We simulate

various network densities by placing 15 to 30 nodes in a network.

3.1.3 FLOW GENERATION

We vary the numbers of flows in each topology. For a network with N nodes, each node has a 50% probability of an upstream flow and a 50% probability of a downstream flow. To avoid TCP complications, UDP data was simulated using Constant Bit Rate (CBR) traffic. Each packet is 1500 bytes long.

3.1.4 MEASUREMENTS

The experiments we conducted are as follows. For a given set of flows in a given network topology, we compute the fair-share rate of the flows using the clique model. Since the bandwidth efficiency at a data rate of 1 Mbps is 90% when RTS/CTS is not used and 86% otherwise [23], the capacity of a clique $B_{(u)}$ is 900 kbps when not using RTS/CTS and 860 kbps otherwise. We run the simulation by limiting the source of each flow to a certain rate, which we call the “input rate,” and measure the throughput of each flow at the destination. We choose 50 input rates ranging from 50% of the computed fair-share rate to 150% of the computed fair-share rate.

The first 1000 seconds of each experiment allows routing to be established and no data is collected in that period. For each input rate, we have a warm-up time that is long enough to transmit at least 200 packets, so that the network stabilizes. We then measure the total number of packets received at each destination of each flow for a long-enough time period that ensures at least 2000 packets are transmitted. The throughput for each flow in the experiment is then *total packets received/measured time*. Since the packets can be received in bursts at times, the throughput we measure can be slightly higher or lower than the real value. However, this effect is not significant. We repeat each experiment 5 times and average the 5 results to give an expected throughput for

each flow for any given input rate.

To evaluate the efficacy of source rate control in achieving fairness, we need to measure the fairness and efficiency of the results. For each topology and set of flows, we measure the throughput of each flow at an offered per-node load of the computed fair-share rate (FSR). We refer to the flow with the least throughput as the minimum throughput flow (MTF) for the given topology and flows. Since the throughputs and FSR values vary for each experiment, we normalize by computing $throughput/FSR$ for each throughput. We then compute the following metrics over each type of topology with multiple experiments.

1. JFI of throughputs: This is Jain's fairness index [20] that represents the overall fairness of the network.
2. Average (Avg.) of throughputs/FSR: This value reflects the average aggregate throughput over all topologies at the computed fair-share rate. Aggregate throughput represents the efficiency of the network.
3. Standard Deviation (σ) of throughputs/FSR: This value shows the fairness among all the flows. The lower this value is, the fairer the flows are.
4. Average (Avg.) of MTF/FSR: This value demonstrates the degree of the starvation experienced by the minimum throughput flow in general over all topologies.
5. Standard Deviation (σ) of MTF/FSR: This value demonstrates how significant the starvation can be in some topologies.

The JFI value reflects the overall fairness among all the experiments. The combination of the average of throughputs/FSR and the σ of throughputs/FSR represent whether the network is in general fair and efficient. If the average of throughputs/FSR is high, and the standard deviation is low, then the network is generally fair and efficient. However, one or a few flows could still be suffering from unfairness. The severity and pervasive-

ness of the problem with the minimum throughput flow is reflected by the average of MTF/FSR and σ of MTF/FSR.

3.2 EXPERIMENTAL RESULTS

In the experiment over a 7-hop chain topology shown in Figure 3.3 there are seven flows, shown as the arrowed lines indicating the source and destination of each flow. Figure 3.4 is the graph generated by plotting the simulation results. The horizontal axis is the input rate to which each source is limited; the vertical axis is the actual throughput achieved by the source. Each data line in the graph represents a given flow, with the legend identifying the source and destination of the flow. For example, the legend 2- > 0 means the flow starts at node 2 and sinks at node 0. The vertical line with label “Fair Share” beside it marks the fair-share rate computed by the clique model. The 802.11 MAC protocol is not using the RTS/CTS protocol in this simulation.

The computed fair-share rate for this experiment is 39,130 bps. We can see from the graph that the maximum data point for which the throughputs of all the flows equal to their input rate is approximately 39,000 bps, which is fairly close to the computed fair-share rate. We can see that when the sources are limited to rates lower than the fair-share rate, the throughput of each flow is approximately the same as the input rate. When the input rate of each flow exceeds the fair-share rate, the throughput of some flows drops dramatically. In particular, flows 5-to-0 and 7-to-0 get the lowest throughputs.

The second experiment is a random 30-node topology in a 1000x1000 m^2 area with about 25 flows. The topology is shown in Figure 3.5 and Figure 3.6. In Figure 3.5, the lines between nodes represent transmission links. In Figure 3.6, a dashed line between any two nodes indicating that those nodes are within interference range. RTS/CTS is also turned off in this simulation.




Figure 3.3: 7-Hop Chain




Figure 3.4: Throughput vs. Input Rate for 7-hop chain, without RTS/CTS




Figure 3.5: Transmission Links in Sample Random Topology




Figure 3.6: Nodes within Interference Range for Sample Random Topology

Figure 3.7 shows the simulation result for this topology. While this is not exactly fair sharing of the network, at the computed fair-share rate point, 12,100 bps, the throughputs of most of the flows are approximately the same as the input rate. Three flows deviated from the fair-share rate point early. The computed fair-share rate again matches the experimental fair-share rate of most of the flows quite well.

We also performed the same experiment with the RTS/CTS protocol turned on. The effect of RTS/CTS is shown in Figure 3.8. We can tell from the graph that RTS/CTS makes the performance worse. Not only do more flows deviate early, but some flows only receive one-third of the input rate at the computed fair-share rate point.

Analyzing these experiments, we observe the following patterns:

1. The source-rate-control mechanism is able to achieve fairness among most flows when RTS/CTS was not used.
2. The clique model appears to be fairly accurate in computing the fair-share rate for most of the flows, again, when RTS/CTS was not used.
3. In the random-topology experiment, some flows get much lower throughput than others at the computed fair-share rate point.

We therefore wished to continue our study by collecting data statistically for various topologies and flows to see if these patterns also occur in other scenarios.

3.2.1 ACCURACY OF THE CLIQUE MODEL

We first attempt to answer whether the clique model is accurate in estimating the network capacity by running simulations over various topologies and flows. To determine the accuracy of the model we need to measure the fair-share rate in each simulation




Figure 3.7: Throughput vs. Input Rate for Sample Random Topology, without RTS/CTS




Figure 3.8: Throughput vs. Input Rate for Sample Random Topology, with RTS/CTS

and compare it with the fair-share rate computed by the clique model. We measure the experimental fair-share rate in each simulation using the following procedures.

1. At each input-rate data point (*i.e.*, the rate that source nodes are limited to), we measure the throughput of each flow.
2. We record all the flows whose throughput is 5% lower than the input rate.
3. We repeat the above steps until the point at which the throughput of more than one-third of the flows is less than the input rate by more than 5% for at least four successive data points.
4. We then take the one point prior to the first of those four points as the measured fair-share rate.

While this measurement may seem somewhat arbitrary, visual inspection of plots suggested it is fairly reasonable.

To show the accuracy of the clique model, we executed the experiment in single-channel WMNs, over 50 random topologies that contain 30 nodes in a 1000×1000 m^2 area with between 25 and 40 flows for each topology. We did the experiments both with and without RTS/CTS. We then computed the average error in the clique model compared with the measured fair-share rate, together with the standard deviation.

Results are shown in Table 3.1. The value “FSR” is the computed value of the clique model, and “fp” is the measured fair-share point. We can see that the average error of the clique model, when RTS/CTS is turned off, is about 0.008%, which is fairly low. However, the standard deviation of the error is 11%, which means that in about 32% of the topologies, the clique model either over-estimates or under-estimates the fair-share rate by more than 10%. This is because the clique model may over-estimate the interference, hence, under-estimate the capacity; on the other hand, it may over-estimate the capacity

Measured Entity	No RTS/CTS	With RTS/CTS
Avg (FSR-fp)/fp	0.008%	18%
Std. Dev. of (FSR-fp)/fp	11%	36%

Table 3.1: 30-node Random Topology Accuracy Results

because it assumes perfect MAC scheduling. We discuss this in more detail about these issues in Chapter 4. Nevertheless, the clique model is still fairly accurate in general. On the contrary, when RTS/CTS is turned on, the clique model over-estimates the network capacity by 18%, with a standard deviation of 36%. This suggests that the flows often achieve fairly low throughput when using RTS/CTS because the clique model is not able to capture the characteristics of the RTS/CTS effect.

The accuracy of the clique model in multi-channel WMNs is evaluated in Lee’s thesis [27]. It suggests that the clique model is not very accurate for multi-channel networks. Instead, a *realistic collision domain* model is more accurate. Therefore, for multi-channel experiments, we use Lee’s multi-channel collision-domain model, in which, “two links contend if one endpoint of one link is within transmission range of one endpoint of the other link” [27]. Each link has a collision domain that contains all the links contending with it. Moreover, only one link in a collision domain can transmit at the same time.

3.2.2 EFFECTIVENESS OF SOURCE RATE CONTROL

In this section, we show how effective source rate control is in achieving fairness. Since we demonstrated that the clique model is fairly accurate in single-channel WMNs, we evaluate the fairness and efficiency of the network while sources are rate-limited to the fair-share rate computed by the clique model in single-channel WMNs. We use the same set of experiments conducted in the previous section and compute the fairness and efficiency metrics described in Section 3.1.4.

Measured Entity	No RTS/CTS	With RTS/CTS
Avg throughputs/FSR	95%	88%
σ throughputs/FSR	12%	17%

Table 3.2: Fairness and Efficiency with Source Rate Control in Random Topology

The results for the 30-node random topologies, as shown in Table 3.2, show that with source rate control, the aggregate throughput of the network is, on average, 95% of the theoretical value when RTS/CTS is not used. Note that the value is 95%, not 100%, which seems inconsistent with the average $(FSR-fp)/fp$ value we got in Table 3.1 (0.008%). This is because when we measure the accuracy of the clique model, at the measured fair-share rate, one-third of the flows were getting throughput less than the fair-share point. If we assume that the throughputs are normally distributed, about 68% of the data are within one σ of the mean; *i.e.*, 16% of the data are below $mean - \sigma$ and 16% of the data are above $mean + \sigma$. The table shows that half of the flows examined get throughput that is higher than 95% of the fair-share rate. At the same time, more than 15% of the flows get throughput 17% less than the fair share. When RTS/CTS is used, the results are much worse; the aggregate throughput is 7% lower than without RTS/CTS. Moreover, the network is less fair, with more than half of the flows getting throughput at least 12% below their fair-share rate, and more than 15% of the flows getting 30% less. In summary, the results demonstrate that source rate control does provide fairness to most of the flows in the network when RTS/CTS is not used. However, it cannot achieve complete fairness.

3.2.3 SEVERITY OF UNFAIRNESS WITH SOURCE RATE CONTROL

In this section, we continue the study by trying to understand how often and severely some flows cannot achieve the throughput computed by the clique model.

We first analyze in more detail the flows that did not get their fair share in our earlier

experiment on the sample random topology, as shown in Figure 3.7. The throughput achieved by most flows is approximately the rate to which they are limited, up to about 12 kbps, which is the theoretical capacity computed for this particular network and set of flows. However, three flows are falling short, two significantly. At 12 kbps flows 13-to-0 and 22-to-0 receive only three-quarters of their rate-limited bandwidth. Even at only 80% of theoretical network capacity, the throughput of these flows is 10% below their input rate. At such a low load this is not caused by a lack of wireless capacity, as shall be seen in Chapter 4. From now on, we name this type of unfairness *Structural Unfairness*. We define structural unfairness as follows:

Definition 1 (Structural Unfairness) *Structural unfairness is unfairness experienced by one or more flows within a wireless mesh network when the sources of all flows are rate limited to the computed fair-share rate.*

Now we wish to determine if the structural unfairness (SU) is a general problem across many topologies and networks flows. We take the same set of simulation results we used to analyze the accuracy of the clique model for the evaluation.

We first evaluate how many topologies and flows have this problem. Since we did not know how severe the problem might be, we based our SU evaluation on a throughput drop of between 5% and 25%. We consider a flow to be experiencing structural unfairness if the throughput of that flow is more than a certain percentage lower than the computed fair-share rate when the input rates of the sources are at the fair-share rate. If there is at least one flow experiencing structural unfairness in a topology, we count that topology as a topology with an SU problem. The results are evaluated for each type of topology in the following table, in which “%SU topologies” means the percentage of topologies that experience the SU problem and “%SU flows” means the percentage of the total flows examined that experience the SU problem.

Benchmark	Measured Entity	30-node random topologies	5x5 grid
5%	%SU topologies	94%	95%
	%SU flows	12%	11%
10%	%SU topologies	82%	35%
	%SU flows	8%	4%
15%	%SU topologies	72%	25%
	%SU flows	6%	3%
20%	%SU topologies	68%	5%
	%SU flows	5%	0.5%
25%	%SU topologies	60%	0%
	%SU flows	4%	0%

The results show that significant unfairness experienced by one or more flows is a widespread phenomenon. In the random topologies we examined, 82% of the topologies have flows that get throughput that is 10% lower, and 60% of the topologies have flows that get one-quarter less. The SU problem is less severe in our grid topologies. However, there are still one-third of the topologies that have flows getting 10% less throughput and one-quarter that have flows with throughput 15% lower.

Now that we know that SU is a general problem, we wish to study in more detail how severe the structural unfairness can be when a network experiences it. For the same set of experiments, we compute the metrics to indicate the degree and significance of the starvation described in Section 3.1.4. The results for the 30-node random topologies are shown in Table 3.3.

Assuming normal distribution, the results show that half of the topologies examined had flows that failed to achieve even two-thirds of their input rate at the fair-share rate. More than 85% of topologies had flows that experienced a degradation of more than 10%, with the degradation being more than 40% when the RTS/CTS protocol was used.

Measured Entity	No RTS/CTS	With RTS/CTS
Avg of MTF/FSR	62%	44%
Std. Dev. MTF/FSR	25%	15%

Table 3.3: Severity of SU in 30-node Random Topology

Measured Entity	No RTS/CTS
Avg of MTF/FSR	88%
Std. Dev. MTF/FSR	5.9%

Table 3.4: Severity of SU in 5x5 GRID Topology

While unfairness is somewhat lessened when RTS/CTS is turned off, it is still the case that the minimum throughput flow in 15% of the experiments is less than 37% of its input rate. RTS/CTS makes the problem more severe. In particular, not only is the minimum throughput flow worse, but the deviation is smaller, and thus the results are *consistently* worse! Further, in 15% of the experiments RTS/CTS causes the minimum throughput flow to achieve less than one-fifth of its input rate.

In a 5x5 grid topology, the starvation evaluation results are shown in Table 3.4.

We can see from the table that the SU problem in the grid topologies is not as severe as in the random topologies. Assuming normal distribution, in half of the topologies examined, even the worst-case flow gets more than 88% of the fair share. However, there are still 15% of the flows that get throughput 18% less than the fair share.

3.3 SUMMARY

As a result of these experiments, we draw three conclusions.

1. Source rate control provides fairness among most network flows. However, it still

leaves some flows suffering from poor throughput.

2. The clique model is fairly accurate in estimating network capacity in single-channel WMNs.
3. It is fairly common that some flows in a network cannot get their fair share computed by the clique model, even when the sources are rate-limited to below the network capacity. We call this phenomenon structural unfairness. When this happens, some flows can get very low throughput, to the point of starvation.

Having understood that structural unfairness is a general problem, we wish to study the causes of the problem. Since all the experiments conducted above show that RTS/CTS is very inefficient in WMNs, in the following experiments, we always turn off RTS/CTS unless specified.

4 CAUSES OF STRUCTURAL UNFAIRNESS

The first question we needed to answer was whether structural unfairness is caused by lack of wireless network capacity. This is essentially asking if the fair-share rate computed by the clique model accurately represents the capacity of the network.

Jain *et al.* [19] investigate the maximum throughput in wireless WMNs by modeling networks using the clique model. They prove that the upper bound based only on the clique model constraints is tight only for a special class of conflict graphs called perfect graphs. Perfect graphs are the graphs for which the chromatic number equals the maximum clique size.

According to the statistical simulation results shown in Table 3.1, we see that over the 50 random topologies, the computed fair-share rate based on the clique model is fairly accurate, with an average error of 0.007% compared to the measured rate. This further suggests that the problem is not capacity.

However, the clique model has certain assumptions. First, it implicitly assumes a perfect global scheduler. In practice, a perfect global schedule is not easy, if even possible, to achieve in a multi-hop wireless network. In particular, it requires, at a minimum, accurate clock synchronization and distribution of current traffic usage. Second, the clique model assumes that flows are bidirectional over links. That is, it treats the sender and receiver as though both were simultaneously sending and receiving. In practice, 802.11 communication consists mostly of the transmitter sending, with just a brief ACK transmission from the receiver. Therefore, the clique model could over-estimate the interference of the network and under-estimate the capacity. Third, it assumes that there

is either interference between two links or no interference at all, which is not always the case in reality. For example, as described in Section 3.1, if the power signal of the sender is 10dB higher than that of the interfering node, the receiver is able to capture the packet from the sender if the sender starts transmission before the interfering node.

Extant MAC protocols, including 802.11, make MAC-scheduling decisions based solely on locally available information. In particular, 802.11 uses carrier sensing and collision avoidance. Our simulations, and most current approaches to achieving fairness, add in source rate limiting, but the MAC scheduling decisions remain fundamentally based on local knowledge, rather than global knowledge.

Given that the problem is not network capacity, but poor MAC scheduling, we wished to determine the precise reasons for non-optimal local decision making. Our initial investigation suggested that it was simply a case of the well-known and understood interaction between the hidden-terminal problem and the 802.11 binary-exponential-backoff algorithm [3, 4, 11, 32, 34, 48]. However, this explanation fails because in many instances where there is a hidden terminal the problem does not occur. For example, in experimenting with various length chains, where nodes are 200 meters apart, and thus there are multiple hidden terminals, application of source rate limiting is sufficient to provide flow fairness. This is consistent with results from Gambiroza *et al.* [12] where source rate limiting worked without flaw for a four-node chain, even though the two end nodes are hidden terminals with respect to each other.

In addition to the MAC scheduling problem, further investigation showed that there are another three criteria that are necessary factors in structural unfairness: topology, link utilization and wireless hop distance. We analyze each category in the next sections.




Figure 4.1: Topology Requirements

4.1 TOPOLOGY REQUIREMENT

The general topology necessary for structural unfairness to occur is illustrated in Figure 4.1.

1. Node S_1 is within transmission range of R_1 , and node S_2 is within transmission range of R_2 . Nodes S_1 and S_2 are out of carrier-sense range of each other. Therefore, when node S_1 wishes to send a packet to R_1 while S_2 is transmitting to R_2 , node S_1 has no way of knowing that S_2 is transmitting, and *vice versa*.
2. At least one of the following three distances are within interference range of each other: $D_{S_1R_2}$, $D_{S_2R_1}$ or $D_{R_1R_2}$.
 - (a) If $D_{S_2R_1}$ is less than interference range, while the distance $D_{S_1R_2}$ is greater, then we refer to it as an *asymmetric* case, as the transmission from S_1 to R_1 will be affected by S_2 's transmission, while that from S_2 to R_2 will not be affected by S_1 's transmission.
 - (b) When both the distance $D_{S_1R_2}$ and $D_{S_2R_1}$ are less than interference range, we refer to it as a *symmetric* case.

- (c) If $D_{R_1R_2}$ is less than interference range, the ACK packet sent from the receiver R_1 to S_1 could potentially collide with S_2 's data packet at R_2 . We call this *receiver-induced* structural unfairness.

We first examine the asymmetric case. When node S_2 is transmitting, node S_1 senses an idle medium and attempts to transmit. If it is using the RTS/CTS protocol, it will issue an RTS, and R_1 will not respond, since it senses the medium is busy. Node S_1 will therefore double its contention window, select a new random delay, and count down (since it perceives an idle medium, nothing stops it from counting down). It then retries, with the same effect. By the time S_2 finishes transmitting its message, S_1 has built up a very large contention window. Therefore, even when the medium is idle, S_1 will be busy counting down its backoff counter and waste the opportunity to transmit. In particular, source rate limiting will be of limited use, since when S_2 is not sending, S_1 will be in backoff.

Without RTS/CTS the problem still exists, but is less severe, since S_1 will attempt to send a long message, rather than a short RTS before discovering the problem, and thus its contention window builds up more slowly, though it does still build. In the presence of source rate limiting, node S_2 may no longer be sending, allowing S_1 to transmit successfully. Node S_1 is at an asymmetric structural disadvantage since R_2 is not within interference range of S_1 , and thus it always receives the message from S_2 correctly. When R_1 and R_2 both receive a packet from their respective senders correctly, the ACK packet sent by R_1 could also collide with the ACK packet from R_2 to S_2 . This effect of the ACK packets is negligible, for reasons we discuss in Section 4.3.1.

The symmetric case is similar, except that since $D_{S_1R_2}$ is also less than interference range, it is as probable that the S_2 -to- R_2 transmission will be affected as that from S_1 -to- R_1 . Thus, in the long term, the effect will be equal on the two receivers.

In the receiver-induced case, since node R_2 is within interference range of R_1 , when R_2 sends an ACK packet after successfully receiving a packet from S_2 , it could collide with the packet transmitted from S_1 to R_1 , because the receiver does not check the status of the channel before sending out an ACK. Likewise, an ACK from R_1 can interfere with reception of a packet at R_2 . As mentioned above, we describe the ACK effect in more detail in Section 4.3.1.

This understanding of the problem, thus far, is reasonably studied (see, *e.g.*, [11, 13, 32, 39]), though in the context of single-hop wireless networks in all the prior literature but one. Garetto *et al.* [13] focus their study on identifying the reasons for packet losses in single-hop flows and categorize them as information asymmetry (similar to our asymmetric case), near hidden terminal (similar to our symmetric case), and far hidden terminal (related to our receiver-induced case). Their study is based on networks without source rate limiting, and focuses on modeling the effects in terms of collision probability. As we will see, collision probability is a poor predictor of throughput or unfairness. Our study focuses on the problem after eliminating congestion by using source rate limiting. The OML [39] proposed by Rao and Stoica is the only paper that has studied this problem in multi-hop networks. However, their study is limited to a six-node wireless testbed. What has yet to be studied is why structural unfairness occurs in some of these topologies, such as the 30-node random topology case in Figure 3.7 and not others, such as the chain topology case in Figure 3.4, even though they both satisfy the generic topology in Figure 4.1. We observe that while the topology requirement is necessary, it is not sufficient to cause structural unfairness.

4.2 LINK UTILIZATION

Utilization of the two competing links is a factor that affects the collision probability. The reason is that the two senders in Figure 4.1 are effectively operating an Aloha protocol [44], since neither can sense the other. Therefore, they both transmit at will. If the acknowledgement from the receiver is not received within a timeout period. The sender waits for a random period and retransmits. The collision probability, P_c , as a function of nominal load, G , is:

$$P_c = 1 - e^{-2G} \quad (4.1)$$

If there is spare capacity in the system given the offered load, this can be used to provide for the retransmissions when collisions occur, and no unfairness will be observed. If there is no capacity left in the system, then unfairness will occur immediately for the interfering links. The degree of unfairness depends on the ratio of unused capacity to the capacity needed to compensate for the collisions.

Consider the single-channel 5-node chain shown in Figure 4.2. The distance between nodes 4 and 3 is 249 meters, while other nodes are 185 meters apart. Node 0 is the gateway and there is one flow from every other node to the gateway. This topology is a case of asymmetric structural unfairness, with node 1 being out of node 4's carrier-sense range, but in interference range of node 3. As we can see from Figure 4.3, flow 4-to-0 receives less than 76 kbps, more than 15% less than the computed fair-share rate, 90 kbps. We analyzed the trace file at the computed fair-share rate. Node 1 caused 75827 collisions at node 3 for 26402 packets sent from node 4, causing the MAC of node 4 to discard 4039 packets after the maximum number of retries, which amounts to about 15% drops. This analysis show that the unfairness at node 4 was indeed because of the collisions at node 3 caused by node 1.




Figure 4.2: Single-channel uneven-distance chain




Figure 4.3: Asymmetric unfairness in single-channel chain

Doing the same experiment on a 3-channel chain, as shown in Figure 4.4, where links 2–1 and 3–2 use different channels from links 1–0 and 4–3, the fair-share rate is much higher than in the single-channel case. As with the single-channel experiment, flow 4-to-0 fell off before the computed fair-share rate (see Figure 4.5). The structural unfairness is far more severe in this experiment than in the single-channel case. Flow 4-to-0's throughput is less than 20% of the fair-share rate. This fair-share rate is computed using Lee's multi-channel collision-domain model. However, even at the fair-share rate computed using the more conservative clique model, 180 kbps, flow 4-to-0 only gets one-third of the fair-share rate. The trace file analysis show that node 1 caused 78879 collisions at node 3 for 26438 packets sent from node 4, causing the MAC of node 4 to discard 9564 packets after the maximum number of retries. Moreover, 12563 packets are dropped at node 4 because the queue is full, before they have a chance to be transmitted.

The difference between the single-channel and multi-channel results is because, in the multi-channel scenario, the link utilization of the competing links is much higher than in the single-channel chain. In the single-channel chain, the fair-share rate is $B/10$. Therefore, link 4–3 needs to carry 10% of the capacity, 3–2 20%, 2–1 30% and 1–0 40%. However, in the multi-channel case, the fair-share rate is $B/4$ according to the collision-domain model. This means that link 1–0 carries 100% of the wireless capacity and link 4–3 needs to carry 25%. Since links 1–0 and 4–3 are the competing links in both scenarios, the link utilization in the multi-channel chain is much higher, hence causing much more collisions at node 3, to the point that few packets can get through even after retransmission. This causes the IFQ to be filled quickly, and half of the packets are dropped at the queue.




Figure 4.4: Multi-channel uneven-distance chain




Figure 4.5: Severe unfairness in multi-channel chain

4.2.1 EFFECT OF MULTI-HOP FLOWS

Since flows typically traverse multiple hops in WMNs, it is possible that two competing links belong to the same flow. For example, in the chains we studied above, the competing links 4–3 and 1–0 are both on the route of flow 4-to-0. We therefore wish to study whether this factor has any effect on structural unfairness.

To avoid any complication caused by other flows, we reduce the flows to only flow 4-to-0. The fair-share rate in the single-channel case is $B/4$; hence, both links 4–3 and 1–0 need to carry 25% of the capacity. The collision probability is comparable to the case with all four flows. The result is shown in Figure 4.6. When the input rate equals to or exceeds the fair-share rate, 225 kbps, flow 4-to-0 gets 223 kbps, about 99% of the fair-share rate. The trace file show that at the fair-share rate, node 1 caused 25875 collisions at node 3 for 26374 packets sent. In addition, there are 168 packets that were dropped by IFQ at node 4. The number of collisions has dropped significantly compared to the four-flow case. This is because node 1 does not have its own traffic. Instead, it only forwards node 4's traffic. Therefore, at the beginning, when node 4 transmits, node 1 does not have data to send and cannot interfere with node 4's transmissions. When node 1 receives packets for flow 4-to-0, it forwards them to node 0. Its transmissions will collide with node 4's transmissions and silence node 4, which will soon drain the queue at node 0 and the link utilization on link 1–0 drops to 0% until more packets are transmitted successfully from node 4. Therefore, even though the average link utilization on link 1–0 is 25%, the oscillation of the load on link 1–0 makes it possible for node 0 and node 4 to alternate their transmissions and reduce the collision probability. The reason that flow 4-to-0 can keep its throughput the same as the fair-share rate even when the load is much higher is that the source node 4 cannot inject more packets than the network can handle. This rate cannot be higher than 900 kbps and the overloaded traffic is simply dropped at the IFQ of node 4. Therefore, the result is the same as if the load is

exactly the fair-share rate.

The same experiment in the multi-channel chain is shown in Figure 4.7. It also shows that flow 4-to-0 is not as starved as in the experiment with four flows. However, it is capped to 380 kbps, much smaller than the fair-share rate, 900 kbps, not even reaching the more conservative fair-share rate computed by the clique model, 450 kbps. The trace file shows that at the fair-share rate, 900 kbps, nearly 60% of the packets are lost due to IFQ drops. We wondered whether it was because the queue length, 20, was too small. Therefore, we did the same experiment with queue length of 50, 150 and 1500. However, the results are the same for all instances. The trace files show a similar number of IFQ drops and collision numbers. We found that at the beginning of the experiment, whenever a packet was sent through successfully from node 4 to node 3, the next packet would collide while node 1 tried to forward the packet to node 0, which throttled transmissions from node 4. The IFQ fills at node 4, since the load exceeds the service rate, with excess packets dropped from the queue. The reason transmissions at node 4 collide more quickly than in the single-channel case is that the link utilization on the two links is much higher; even at 380 kbps, link 1-0 is carrying 42% of the capacity and 4-3, because of retransmissions, is carrying (close to) 100% of the capacity. This is consistent with the analysis that higher link utilization causes higher collision probability. On the other hand, we also notice that flow 4-to-0 gets about 42% of the fair-share rate, which is much higher than in the four-flow scenario. The reason is the same as for the single-channel experiment, *i.e.*, a multi-hop flow can be disadvantaged by other flows, but cannot starve itself.




Figure 4.6: Multi-hop flow effect in single-channel chain




Figure 4.7: Multi-hop flow effect in multi-channel chain

4.3 HOP DISTANCE

The distance between wireless transceivers determines the power of the signal at the receiver, or equivalently, the power of an interfering signal needs to disrupt reception. If the receiver is receiving a packet, and is close enough to the sender and far enough from the interfering node, the power of the sender's signal will be high enough at the receiver that the receiver can correctly decode the packet despite the interfering node. However, if the interfering node initiates transmission first, the receiver will not be able to capture the packet from the sender. This effect is referred to as power capture.

In ns2, a receiver can receive a packet successfully if the received signal power of the sender is 10 dB higher than that of the interfering node. Denote by P_s the signal power of sender s at the receiver, and by P_i the signal power of interfering node i at the receiver. When $\frac{P_s}{P_i} \geq 10$ the power capture succeeds. If we denote the distance from the sender to the receiver by d_s , and the distance from the interfering node to the receiver by d_i , the received signal powers of the sender and the interfering node can be written as a function of distance:

$$\begin{aligned} P_s &\propto \frac{1}{d_s^4} \\ P_i &\propto \frac{1}{d_i^4} \end{aligned} \tag{4.2}$$

Therefore, we get:

$$\begin{aligned} \frac{P_s}{P_i} &= \frac{d_i^4}{d_s^4} \\ \frac{d_i^4}{d_s^4} &\geq 10 \\ \frac{d_i}{d_s} &\geq 10^{\frac{1}{4}} \end{aligned}$$

$$d_i \geq 10^{\frac{1}{4}} d_s$$

Therefore, if $d_i < 10^{\frac{1}{4}} d_s$, power capture at the receiver will not happen.

Consider a 5-node chain again, except this time the nodes are each 200 meters from their neighbors. The single-channel and multi-channel chains are shown in Figure 4.8 and Figure 4.9. We study the case where there is one flow from every other node to the gateway. Note that these two topologies have the same contention graphs as when the nodes have uneven distances between them as in Figure 4.2 and Figure 4.4. The link utilization on each link is also identical. However, all nodes achieve their respective fair-share rates this time, 90 kbps for the single-channel case (see Figure 4.10), and 225 kbps for the multi-channel case (see Figure 4.11). The reason is that this time node 3 is closer (200 meters) to its sender, node 4, and farther (400 meters) from the interfering node 1. According to the equations above, node 3 is able to power capture as long as the interfering node is further than 355 meters. Therefore, if node 4 starts transmitting before node 1, node 3 is able to power capture the packet from node 4. In the uneven-distance chains, node 3 is further (249 meters) from the sender and closer (370 meters) to the interfering node. Since node 3 can only power capture if the interfering node is further than 443 meters, even when node 4 starts transmitting first, the signal-to-noise ratio (SNR) is not high enough for node 3 to decode the packet correctly.

However, power capture alone is insufficient to explain the unfairness. After all, power capture does not allow a node to start reception, if an interfering node has started first. According to the trace file, in the single-channel case, when 22213 packets are sent from node 4 to node 3, 27437 packets collide with packets from node 1. Compared with 57264 packet collisions when 20724 packets are transmitted in the uneven chain, the number of collisions has been reduced by more than half because of the power-capture effect. However, this is still an extremely high collision rate. There are, on average, 1.24




Figure 4.8: Single-channel even-distance chain




Figure 4.9: Multi-channel even-distance chain




Figure 4.10: Fairness in single-channel chain

collisions for each successful transmission. The reason this does not affect the overall throughput is that the link utilization on the 4–3 link is 25% of the capacity; the spare capacity is enough to compensate for the retransmissions.

The trace file for the multi-channel case is surprising. Out of 22076 packets sent from node 4 to node 3, there are 19985 collisions at node 3. While still very high, fairness is maintained for the same reason as in the single-channel case. What is intriguing, however, is that there are fewer collisions than in the single-channel case. This is counter-intuitive since we already know that the link utilization in the multi-channel case is much higher than in the single-channel case, which should cause more collisions. After further investigation, we understand that since in the single-channel case all the nodes share the same channel, the transmissions from nodes 2 and 3 have synchronized the transmissions from nodes 1 and 4. For example, assume node 1 transmits first. Since nodes 2 and 3 are in carrier-sense range, they will not transmit at the same time. However, node 4 is out of carrier-sense range of node 1, so it could transmit while node 1 is transmitting, in which case node 1 transmits successfully and node 4's packet collides at node 3. Now it is the turn for either node 2 or node 3 to transmit. Node 1 picks a random backoff counter and node 4 doubles its contention window and picks a random backoff counter from the new contention-window size. It is more likely that node 4 picks a bigger backoff counter than node 1. Then after node 2 or 3 finishes its transmission, assume node 1 finishes its backoff period and gets access to the channel again. For the same reason as described above, nodes 2 and 3 will not transmit. Node 4, on the other hand, is likely to finish its backoff period while node 1 is still transmitting and so starts to retransmit. Obviously, the collision happens again at node 3. The same phenomenon is likely to happen again and again. However, in the multi-channel case, because nodes 2 and 3 are on different channels, nodes 1 and 4 will keep transmitting again and again after their random backoffs. Their transmissions are not synchronized by the middle

nodes. Therefore, the collisions are rarer than in the single-channel case. Note that, since the link utilization in the single-channel case is relatively low, even though there are so many collisions, flow 4-to-0 still gets its fair share. This is consistent with the analysis in Section 4.2.

4.3.1 EFFECT OF ACK

Until now, we have been ignoring the effect of ACK packets and the receiver-induced structural-unfairness topology. Let us first assume that the receivers are out of carrier-sense range of each other; *i.e.*, the non-receiver-induced case. In an asymmetric structural-unfairness case, where $D_{S_1R_2}$ is greater than interference range, transmissions from S_2 will collide with the packet sent from S_1 to R_1 if they are sent simultaneously. However, if S_1 starts to transmit before S_2 , and R_1 successfully receives the packet from S_1 as a result of power capture, it will reply with an ACK packet after SIFS, without checking whether the channel is idle. Suppose R_2 also receives a packet from S_2 at this point and replies with an ACK to S_2 . This ACK could collide at S_2 with the ACK sent by R_1 . This is similar to the symmetric case, except ACKs from both receivers could potentially collide at respective senders. We refer to this as ACK-ACK collision.

In the four-flow multi-channel chain experiments, as shown in Figure 4.9, most of the collisions occur at node 3 when node 1 is transmitting at the same time. However, there are also 771 ACK packets from node 0 that collide at node 1 with ACK packets sent from node 3 to node 4. This is out of 88376 packets sent by node 1. The effect of ACK-ACK drops here is less than 1%. Compared with the collisions at node 3 from packets sent from node 1 to node 0, this is negligible.

This situation will always be the case because if a receiver can send ACKs that collide with the ACK packet sent to another sender, the receiver has to be within interference range of that sender. Therefore, that sender's data packets could also collide with

the packets from the receiver's sender. Since we assume that the data packet size is 1500 bytes, much bigger than the ACK packets, the link utilization is due mostly to normal packets, not ACK packets. Therefore, the probability that a collision is caused by normal packets is much higher than the probability of ACK-ACK collisions. However, if the data packet size is very small, intuitively, the effect of ACK-ACK collision should be greater. We leave this problem to future work and focus only on a packet size of 1500 bytes.

Now consider the receiver-induced case, where the two receivers are within interference range of each other. The ACK packets sent by the receivers could collide with packets sent by the other sender. Note that this interference is necessarily symmetric between the two competing links. To study the effect of these ACK packets, we designed the following experiments (see Figure 4.12). To avoid complications from other sources, such as multi-hop flows and collisions due to other packets, the receivers are not within interference range of the other sender, and we use one-hop flows 3-to-1 and 4-to-0. They both use channel 0. We perform the two simulations by varying the distances between nodes. In the first simulation, all nodes are 200 meters apart to allow power capture, giving the system more spare capacity. The second simulation makes the distances between nodes 1 and 2, and 2 and 0 150 meters, while nodes 3 and 1, and nodes 4 and 0 are 249 meters apart. This is to ensure that the signals from the senders are weak, but the interfering signals are strong. As such, there is no power capture at nodes 1 and 0. This means that if node 4 is transmitting, an ACK from node 1 to 3 will interfere with it, and *vice versa*. The fair-share rate computed by Lee's collision-domain model is 900 kbps for each flow; hence, both links 3-1 and 4-0 carry 100% of the capacity. We expect this experiment to tell us the near upper-bound of how much ACK packets can collide with normal packets.

Figures 4.13 and Figure 4.14 show the results for the two simulations. In the first,




Figure 4.11: Fairness in multi-channel chain




Figure 4.12: Multi-channel chain with ACK interference




Figure 4.13: Effect of ACK interference with Power Capture




Figure 4.14: Effect of ACK interference without Power Capture

Measured Entity	With Power Capture	Without Power Capture
Fair-share rate	900 kbps	450 kbps
Total CBR packets sent	52408	52521
Total CBR packets received	51795	49279
Total collisions	890	38781
Total CBK drops	0	3216
Total IFQ drops	546	0

Table 4.1: Collision Results on Effect of ACK packets

each flow gets 890 kbps, about 99% of the fair-share rate. In the second, each flow achieves only 450 kbps, which is the fair-share rate computed by the clique model. In the second scenario, since there is no power capture, links 3–1 and 4–0 cannot operate at the same time. In this case, the collision-domain model clearly over-estimates the fair-share rate. Therefore, we analyze the collisions at the fair-share rate computed by the clique model, 450 kbps for this instance. The collision results for both simulations are shown in Table 4.1. The total collisions is the total number of collisions at the MAC layer; the CBK drops is the number of packets that are dropped by the MAC layer after the maximum number of retries, which is seven times in our simulation; the IFQ drops is the number of packets dropped from the IFQ because it is full.

It is clear that when power capture does not work, the ACK packets could cause significant collisions. However, in scenarios where there is power capture, the ACK effect on collisions can be safely ignored.

4.3.2 EFFECT OF DOMINATING NODES

If S_1 and S_2 are the only two nodes that are out of carrier-sense range of each other and all the other nodes are within interference range of each other, we call S_1 and S_2 “dominating

nodes.” In this section we study how dominating nodes affect structural unfairness. We derive various scenarios from the same single-channel 7x7 grid, as shown in Figure 4.15.

First, let us look at the two scenarios shown in Figure 4.16. The nodes are 180 meters from their neighbors in the first case and 200 meters in the second. The arrows in the graph represent the routing paths of the flows. In the first, we simulate with two flows 42-to-0 and 15-to-0. In the second, we simulate with flows 28-to-0 and 7-to-0. Both topologies are symmetric structural-unfairness cases. The structural-unfairness topologies are drawn in dashed lines. Both topologies share the same clique graph and the fair-share rate of each flow is $B/4$, 225 kbps.

Figure 4.17 and Figure 4.18 are the results of the two simulations. The flows in the second simulation not only get their fair share at the computed fair-share rate, their throughputs are the same as the input rate until the rate is higher than 250 kbps. However, the flows in the first simulation only achieve 95% of the fair-share rate. The collision results are in Table 4.2.

The number of collisions in the first simulation is twice as many as in the second. This is counter-intuitive. In the first scenario, the senders are 180 meters from the receivers and the interfering nodes are 402.5 meters; while in the second scenario, the senders are 200 meters (further) from the receivers and the interfering nodes are 400 meters (closer) from the receivers. According to the power-capture theory, there should not be more collisions in the first scenario. After further investigation, we found that nodes 15 and 42 are not only competing nodes, but also dominating nodes in the first scenario. Therefore, when node 15 transmits, only node 42 will transmit at the same time. Assume node 15 finishes transmitting; if node 42 is still transmitting, all the other nodes will sense the channel busy, except node 15. Node 15 then starts transmitting again. The same happens when node 42 finishes transmitting. Nodes 15 and 42 dominate this scenario, which increases their chances to collide with each other. The situation in the second scenario




Figure 4.15: 7X7 GRID




Figure 4.16: Scenarios with/without Dominating Nodes




Figure 4.17: Effect of Dominating Nodes




Figure 4.18: Without Dominating Nodes

is different. Nodes 7 and 28 are the competing nodes. However, both nodes 42 and 28 are out of carrier-sense range of node 7. Therefore, when node 7 is transmitting, node 28 can also transmit and leverage the probability that node 35 transmits simultaneously with node 7.

To prove our point, we extended the first scenario by one hop and reduced the second scenario by one hop. The new topologies are shown in Figure 4.19. The structural-unfairness topologies are not changed in the new scenarios. However, in the first, nodes 15 and 42 are no longer dominating nodes, while in the second scenario, nodes 7 and 35 become dominating. Therefore, we expect the results to be the opposite of the original scenarios.

The results in Figure 4.20 and Figure 4.21 show that our analysis is correct. The flows achieve their fair-share rate, 180 kbps, after removing the dominating-node situation in the first scenario. In the second scenario, after creating the dominating-node situation, each flow gets 7% less than the fair-share rate.

4.4 SUMMARY

In summary, the situations that cause multi-hop flows to suffer from structural unfairness are very complex. The following factors are necessary, but likely not exhaustive:

Proposition 1 *Structural unfairness of one or more flows within a WMN requires that there exist a link S_1-R_1 in one flow and a link S_2-R_2 in another flow that satisfy all of the following four conditions:*

1. *The MAC layer makes scheduling choices based on local information only (no perfect global scheduler).*
2. *The sender S_1 must be beyond carrier-sense range of a second sender S_2 operating*

Measured Entity	First Scenario With Dominating Nodes	Second Scenario Without Dominating Nodes
Fair-share rate	225 kbps	225 kbps
Total CBR packets sent	52740	52646
Total CBR packets received	49165	52227
Total collisions	38460	19010
Total CBK drops	2810	1
Total IFQ drops	3	0

Table 4.2: Collision Results on Effect of Dominating Nodes




Figure 4.19: Demonstration of Dominating Nodes




Figure 4.20: Effect of Removing Dominating Nodes




Figure 4.21: Effect of Adding Dominating Nodes

on the same channel. If the receivers are out of interference range of each other, one receiver must be within interference range of the other sender.

- 3. The link utilization of S_1-R_1 and S_2-R_2 must be sufficiently large that the probability of packet collision is non-trivial. The link utilization must also be sufficiently high that spare capacity cannot compensate for retransmissions.*
- 4. The physical distance between S_1 and R_1 must be large enough and the physical distance between S_2 and R_1 must be small enough that power capture of the packet from S_1 by R_1 is not possible in the presence of a competing transmission from S_2 .*

Therefore, if we can eliminate any of these conditions (or reduce their effects) we can eliminate (or reduce) the problem of structural unfairness. We now examine solutions based on addressing these issues.

5 IMPROVING MAC SCHEDULING

In this chapter we study practical solutions to ameliorate the problems of unfairness by improving MAC scheduling, while remaining within the constraints of 802.11 hardware.

5.1 FIXING THE CONTENTION WINDOW

The first solution is based on our observation that the buildup of the contention window, as described in Section 4.1, is caused by the window serving two purposes. It acts as a mechanism to reduce the likelihood of collisions caused by congestion, as well as by poor MAC scheduling. The BEB mechanism doubles the contention window (CW_{max}) when a collision occurs, giving the successful node more chances to access the channel than the failed node. An 802.11 node believes that there is congestion whenever there is a collision, and it uses the BEB mechanism to reduce the load of the failed node. However, when the collision is not caused by congestion, the failed node will be disadvantaged by having fewer opportunities to access the network. If the collision probability is high enough that the spare capacity cannot compensate for the retransmissions, structural unfairness will occur.

This 802.11-MAC-scheduling problem is fundamentally because the scheduling is based on local knowledge only. In this section we propose to alleviate this problem by improving the 802.11 MAC scheduling decision-making.

We address this problem by separating the two functions that the contention window serves. First, we presume the congestion problem is dealt with by limiting the source rate to below the network capacity in a higher layer of the network. In other words, each source node only occupies the channel for its fair share of time. Then the 802.11

MAC only needs to schedule the nodes to access the channel without interfering with each other. It no longer needs to deal with congestion, but merely with the possibility of collisions. Therefore, the CW_{max} need not grow. We propose fixing it to a value that is sufficient to reduce the probability of collisions to an acceptably small level. We further set the CW_{max} equal for all nodes. This alleviates the problem that one node is given more privilege than the others for channel access when collisions happen. Note that this is only relevant for nodes within interference range, and does nothing to affect nodes beyond that range.

Now the problem is to compute an appropriate value for CW_{max} so that the probability that nodes within interference range will pick the same random backoff time and transmit at the same time is sufficiently small. We observe that this problem is a simple variation of the birthday paradox [50]. If there are n nodes in range that have data to transmit, and nodes randomly select a delay time from 0 to CW_{max} slot times (per the 802.11 standard), then the probability of two or more nodes picking the same delay time is:

$$p_c(n) = 1 - \prod_{i=0}^{n-1} \left(1 - \frac{i}{CW_{max} + 1} \right) \quad (5.1)$$

for $n > 1$. To make this probability sufficiently low, we need to know how many nodes within range of each other might transmit at any given time, and set the value of CW_{max} appropriately. Unfortunately, we cannot easily determine this number. However, we can reasonably expect it to be low, based on the fact that we are rate-limiting the sources. If we presume that this is in the 2-to-3 node range, or less, and fix the value of CW_{max} at 31, its default initial value, we expect reasonable results. This yields a collision probability of less than 10%. Conversely, if we assume that n is 3, for a collision probability p_c we

can calculate CW_{max} :

$$\begin{aligned}
p_c &= 1 - \left(1 - \frac{1}{CW_{max} + 1}\right) \left(1 - \frac{2}{CW_{max} + 1}\right) \\
&= 1 - \left(\frac{CW_{max}}{CW_{max} + 1}\right) \left(\frac{CW_{max} - 1}{CW_{max} + 1}\right) \\
&= \frac{(CW_{max} + 1)^2 - CW_{max}(CW_{max} - 1)}{(CW_{max} + 1)^2} \\
(CW_{max} + 1)^2 p_c &= 3CW_{max} + 1 \\
(CW_{max}^2 + 2CW_{max} + 1)p_c &= 3CW_{max} + 1 \tag{5.2} \\
CW_{max}^2 + \left(2 - \frac{3}{p_c}\right)CW_{max} + \frac{p_c - 1}{p_c} &= 0 \\
CW_{max}^2 + \frac{2p_c - 3}{p_c}CW_{max} &= \frac{1}{p_c} - 1 \\
\left(CW_{max} + 1 - \frac{3}{2p_c}\right)^2 &= \frac{1}{p_c} - 1 + \left(1 - \frac{3}{2p_c}\right)^2 \\
CW_{max} &= \sqrt{\frac{9}{4p_c^2} - \frac{2}{p_c} + \frac{3}{2p_c}} - 1
\end{aligned}$$

When p_c is small, CW_{max} is approximately $\frac{3}{p_c} - 1$. Therefore, if we want the p_c to be 5%, the value of CW_{max} is approximately 59. This technique is implementable using commodity 802.11 hardware.

Our idea was first inspired by Heusse *et al.* and Wang and Garcia-Luna-Aceves. Heusse *et al.* [16] proposed to adjust the CW_{max} of all nodes dynamically to an equal value based on the load of a WLAN. They aimed to solve the short-term unfairness caused by the BEB mechanism in reacting to poor link quality in the same way as it reacts to collisions. We borrow the idea of setting the CW_{max} value equal on all nodes. However, since we use source rate control to deal with the load issue in the network, the value need not be changed. Moreover, even though the link quality of WMNs is not the current focus of this work, our technique of fixing the CW_{max} should also help deal with

unfairness caused by poor link quality.

Wang and Garcia-Luna-Aceves [49] also investigated the technique of fixing the CW_{max} , but in a congested network. They observed that even though a fixed CW_{max} improves fairness, it yields much worse throughput because collisions increase when congestion occurs. Again, since we deal with the congestion of the network by source rate control, fixing CW_{max} can improve fairness without sacrificing the throughput.

5.2 PER-NODE RATE LIMITING

Our second method is based on our study in Section 4.2. According to Equation 4.1, the higher the link utilization, the higher the collision probability of two competing links. Therefore, if we can reduce link utilization we reduce the collision probability of two competing links.

We observe that with the 802.11 MAC, each mesh router tries to compete for the channel as long as there is more than one packet in its queue. Because access to the channel is random, it is possible that a router receives a burst of packets within a short period of time and attempts to send them out all at once. This will cause the link utilization to be high for a short period of time and low at other times. If the senders of the competing links are within carrier-sense range of each other, the burstiness of one sender will delay the transmission of the other for a short while. Over a long period of time, this is not a problem. However, if the senders are out of carrier-sense range, since source rate limiting does not limit the rate on the intermediate router, when a collision occurs at the router, it will retransmit after the random backoff. This causes the link utilization to be temporarily higher, and increases the collision probability of the retransmissions, which in turn will further increase the link utilization.

We therefore propose rate-limit each mesh router to the sum of the rates of the traffic

flowing through or originating at that router. This will spread out the packet deliveries and reduce the short-term high link utilization as a result.

We implemented this technique at each router in ns2 as shown in Figure 5.1. In the original ns2 implementation, each mesh router analyzes both forwarded traffic and traffic it originates. If the traffic is to be sent, it goes through the link layer and is inserted into a FIFO queue, called the IFQ. It waits in the IFQ until the MAC layer gets access to the channel. We implemented a token bucket and inserted it between the IFQ and the MAC layer, as shown in the dashed rectangle. The token-bucket rate is adjusted dynamically to the sum of the fair-share rate of all the flows going through the router. The token bucket is equivalent to a leaky bucket if the size of the bucket is set to one packet. To allow a little fluctuation of the traffic, we set the bucket size to two packets.

5.3 EXPERIMENTAL EVALUATION

We now evaluate our mechanisms using simulations with the same basic experiment setup as described in Section 3.1. First we demonstrate our improvements using the sample random topology (see Figure 3.5) whose standard 802.11 behaviour was illustrated in Figure 3.7. We see that the behaviour with $CW_{max} = 31$, shown in Figure 5.2, has clearly reduced the structural-unfairness problem significantly, while maintaining similar aggregate throughput. Specifically, one of the three problem flows is no longer a problem at all, while the other two are now getting about 10,000 bps throughput, reaching 83% of their fair share at the fair-share rate, where without this mechanism they were only at 74%, getting about 9,000 bps throughput.

The behaviour with per-node rate limiting is shown in Figure 5.3. This approach is not as effective as fixing the contention window. It improved the throughput of two flows marginally. However, we notice that this improvement is with a small sacrifice of

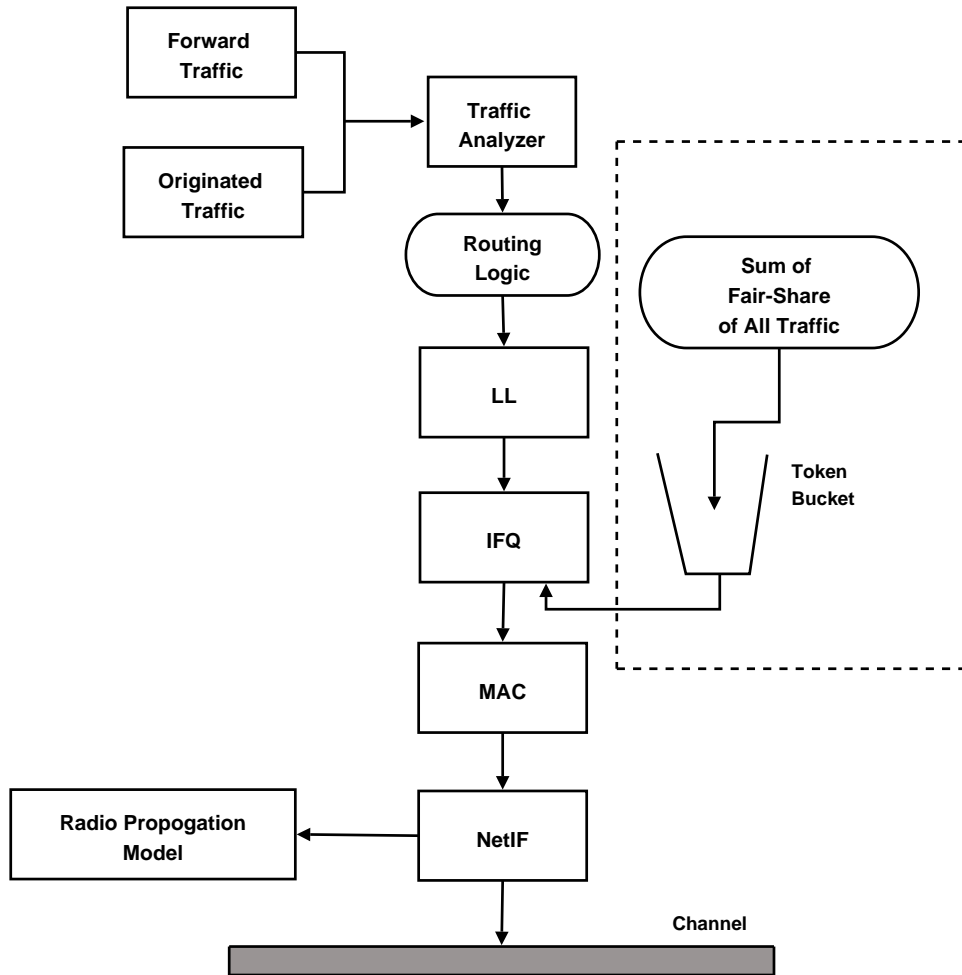


Figure 5.1: Implementation of Per-Node Rate Limiting in ns2

aggregate throughput.

Figure 5.4 shows the combined effect of the two approaches. Even though some of the flows still have slightly reduced throughput, the overall aggregate throughput and fairness are both improved. The two worst flows have improved significantly with this approach. They are achieving 104 kbps and 106 kbps, respectively, reaching 86% to 88% of the fair-share rate.

To study the effectiveness of these mechanisms we collected statistical results over




Figure 5.2: Fixed $CW_{max} = 31$




Figure 5.3: Per-node Rate Limiting

	With RTS CTS	No RTS/CTS			
		As Is	Fixed CW	Per Router Rate Control	Both
Avg of MTF/FSR	44%	62%	69%	69%	73%
Std. Dev. of MTF/FSR	15%	25%	25%	23%	23%
Avg of (thruputs/FSR)	88%	95%	95%	93%	94%
Std. Dev. of (thruputs/FSR)	17%	12%	10%	11%	10%

Table 5.1: Improvement of Fixed CW and Per-node rate limiting

the 50 random topologies with 30 nodes used in Chapter 3. We also use the same metrics to demonstrate improvement in fairness and structural-unfairness severity. Results for the use of fixed CW_{max} as well as per-node rate limiting are presented in Table 5.1. In addition, we repeat the data from Section 3.2.3 for comparison. As the average of MTF/FSR shows, both fixed CW_{max} and per-node rate limiting improve the average of the worst-case flows by more than 11%. The fixed CW_{max} maintained the same aggregate throughput as when no mechanism is used. Per-node rate limiting, while sacrificing aggregate throughput by 2%, has its own benefits. In particular, it makes the standard deviation of the worst-case flows about 10% less than that of the fixed CW_{max} , which means a better fairness index. The combination of the mechanisms has further improved the average of the worst-case flows by 6% compared to each individual method, and with a lower standard deviation, 23%. The aggregate throughput remains the same. The standard deviation of all flows is reduced to 17% less than when the mechanisms are not used. These results demonstrate that the combination of the two techniques has improved the fairness without sacrificing the efficiency of the network.

Since the choice of CW_{max} was determined by presumption rather than clear knowledge, we wished to study the optimal value for CW_{max} . We evaluated CW_{max} values

ranging from 2 to 50 over the 30-node random topologies. To limit the total simulation time, we randomly picked 20 topologies out of the 50 for this evaluation. Then we calculated the throughput of the worse-case flow, the average throughput of all flows, and the standard deviation of all the flow throughputs for each CW_{max} . We present the results in Figure 5.5.

All the fairness curves flatten almost immediately, being more or less flat once CW_{max} has reached 10. According to Equation 5.2, only if the number of contending nodes is fewer than 3 will the collision probability be under 10%, in which case we can expect reasonable results. These results suggest two things:

1. Our presumption of 2 to 3 contending nodes is actually high. Source rate control has reduced the number of contending links to a fraction of what would otherwise occur.
2. There is no obvious optimal value of CW_{max} . It is the fact that CW_{max} is fixed and set equal at all the nodes that really improved the 802.11 MAC scheduling. This fact makes the method of fixing CW_{max} more practical as a solution.

5.4 SUMMARY

This chapter proposes two techniques, fixing CW_{max} and per-node rate limiting, to reduce the structural-unfairness problem. We performed extensive simulations and have demonstrated that these techniques substantially ameliorate the problem, providing on average 18% improvement for the worse-case flows.




Figure 5.4: Combined Effect

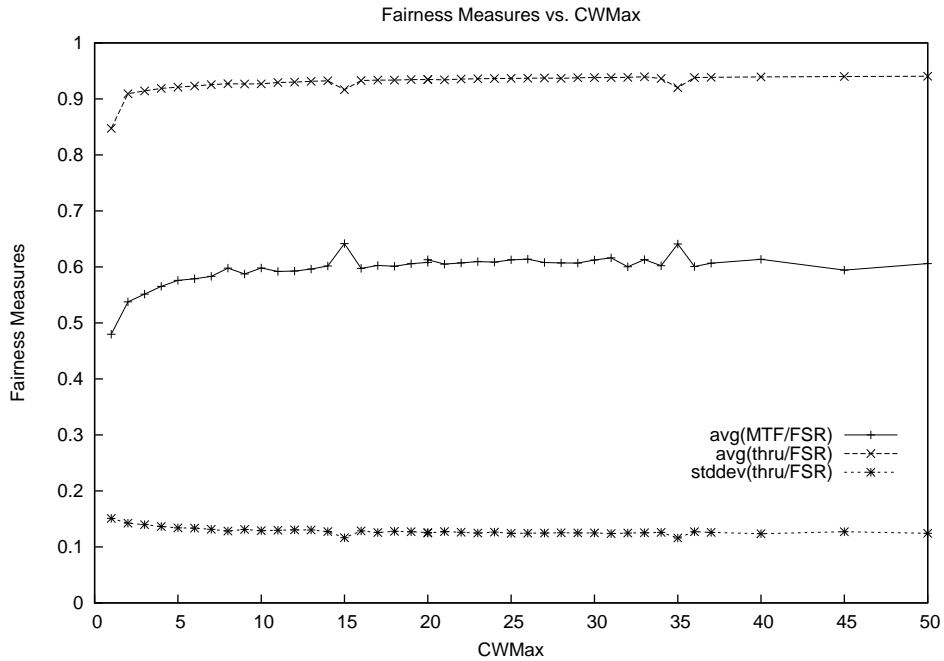


Figure 5.5: Fairness vs. CW_{max}

6 ELIMINATING STRUCTURAL UNFAIRNESS CONDITIONS

As stated in Chapter 4, if we can remove any of the structural-unfairness conditions, we can eliminate the problem. In this chapter, we provide case studies of eliminating one or more of the necessary conditions using three techniques: channel re-assignment, careful node placement, and re-routing. These must be done with three constraints in mind: mesh connectivity must be maintained, aggregate throughput must be maintained, and any new structural unfairness must be less than the case being resolved.

6.1 IDENTIFY THE PROBLEMATIC LINKS

The first step is to identify the SU topology; in essence, the problematic links. We then can apply appropriate techniques to remove the problem. We developed a three-step algorithm to identify the problematic links.

Identify the structurally unfair flows $F(su)$. Let F be the set of all flows in the network, $FSR(f)$ be the computed fair-share rate for flow f , $T(f)$ be the throughput of flow f , and p_d be the percentage of throughput drop that defines structural unfairness. Then:

$$\forall f \in F, f \in F(su) \iff \frac{FSR(f) - T(f)}{FSR(f)} > p_d \quad (6.1)$$

Identify candidate disadvantaged links $L(d)$. To do this, we first identify all links in the disadvantaged flows. We remove from this set any links that are part of flows

that receive their fair share. This creates a small set of candidate links, $L(d)$, typically one per disadvantaged flow.

Let L denote the set of all links in the networks.

$$\begin{aligned} \forall l \in L, l \in L(d) &\iff \exists \text{flow } f \in F(su), f \text{ traverses } l \\ &\wedge \exists \text{flow } f' \in F, f' \notin F(su), f' \text{ traverses } l \end{aligned} \quad (6.2)$$

Find potential partner links $L(p_i)$ for each disadvantaged link i . For each disadvantaged link i , we find all the links in the network that satisfy the topology requirement (Section 4.1), and identify them as topological partners of link i . This will create a set of links $L(p_i)$ for each link i . We then sort potential partners in $L(p_i)$ by link utilization, so as to consider higher-load links first.

6.2 CASE STUDY 1: CHANNEL RE-ASSIGNMENT

The first of our case studies uses channel reassignment to remove the problematic topology. Either of the contending links may have its channel reassigned. However, maintaining mesh connectivity is non-trivial. Most deployed multi-channel meshes use two interfaces on fixed channels. As such, any change in channel may preclude connectivity for other nodes. We therefore approach this problem by changing the channel on only one interface of one of the four nodes, so as to match its partners other channel. We consider the four possibilities iteratively, selecting the best choice based on the other two constraints.

Consider the 3-channel uneven-distance chain (see Figure 4.4) we discussed in Section 4.2. Links 4–3 and 1–0 both use channel 0 and they form an asymmetric structural-unfairness case. Flow 4-to-0 gets poor throughput because many packets sent by node 4

collide with the packets sent by node 1. If links 4–3 and 1–0 do not use the same channel, the asymmetric structural-unfairness topology disappears, hence the fairness problem disappears. Therefore, we change nodes 4 and 3 to communicate on channel 1 instead of channel 0, as shown in Figure 6.1. With the new channel assignment, the network is still connected, the fair-share rate remains $B/4$, 225 kbps, and there is no structural-unfairness topology created. Therefore, channel re-assignment has eliminated the structural unfairness in this scenario.

Figure 6.2 is the result after the channel re-assignment. We can see that all the flows get their fair-share rate, 225 kbps, which is consistent with our analysis.

6.3 CASE STUDY 2: RE-ROUTING

The second of our case studies uses re-routing. This is based on the same general principle of identifying the disadvantaged links and then considering all alternate equivalent-length routes that satisfy the three constraints.

Consider the scenario shown in Figure 6.3. It is also derived from the 7x7 grid (see Figure 4.15). The nodes are 200 meters from their neighbors. Flow 15-to-0 routes through nodes 8 and 1, and flow 20-to-0 routes through nodes 13 and 6. They have formed two symmetric structural-unfairness topologies. One is between links 15–8 and 6–0 and the other is between links 1–0 and 20–13. Both of the flows achieve less throughput than the fair-share rate. However, if we change the routing of both flows to go through nodes 14 and 7 instead, as shown in Figure 6.4, all the nodes involved now are in interference range of each other. The structural-unfairness topologies disappear.

Figures 6.5 and 6.6 are the results with the original routing and after re-rerouting respectively. In the first result both flows get 7% less than the fair-share rate. Re-routing improved the throughput of both flows to 99% of the fair-share rate in the second result.




Figure 6.1: Channel Re-assignment




Figure 6.2: Removed SU Topology After Channel Re-assignment




Figure 6.3: Original Routing With SU topologies




Figure 6.4: After Re-Routing Without SU topologies




Figure 6.5: Structural Unfairness With Original Routing




Figure 6.6: Fairness With Routing Changes

6.4 CASE STUDY 3: CAREFUL NODE PLACEMENT

Our third case study uses careful node placement to either remove the offending topology or maximize the chance of power capture. We attempt the first whenever the topological conditions are only just satisfied, such that a small movement can remove the condition. When this is not true, we move the receiver on the disadvantaged link so as to minimize its distance to its sender and maximize its distance to the offending sender.

In Chapter 4 we demonstrated, though not explicitly, how careful node placement can maximize the chance of power capture. In both the single-channel (see Figure 4.2) and multi-channel (see Figure 4.4) uneven-distance chains, we demonstrated that flow 4-to-0 experiences structural unfairness (see results in Figures 4.3 and 4.5). However, the unfairness disappears in the even-distance chains as shown in Figures 4.8 and 4.9) because, after careful node placement, the suffering node 3 is closer to the sender and further from the interfering node (see results in Figures 4.10 and 4.11). Therefore, the power-capture effect eliminated the unfairness condition.

6.5 SUMMARY

In summary, we demonstrated that the three techniques of channel re-assignment, re-routing and careful node placement are effective whenever applicable in removing structural-unfairness conditions and eliminating the unfairness problem. We identify the constraints when applying these techniques, but leave the detailed algorithm design to future work.

7 CONCLUSION AND FUTURE WORK

This thesis studies fairness schemes in WMNs; in particular, the source-rate-control mechanism. We demonstrated that source rate control, while effective in alleviating the fairness problem in many cases, cannot achieve fairness completely. We define the unfairness experienced with the presence of source rate control as “structural unfairness.” We show that the problem of structural unfairness in WMNs is a widespread phenomenon and studied in depth the four required conditions for structural unfairness. We then presented two methods to ameliorate the problem by improving 802.11 MAC scheduling. Our proposed mechanisms are feasible without alteration of commodity hardware. Even in the worst case, the worst-case flow is only at 30% below the average, though the large deviation suggests that this varies quite a bit by topology. The average throughput is slightly lower when using our approach than when just omitting the RTS/CTS protocol, though the deviation is smaller, meaning that our approach is objectively fairer. Finally, we presented case studies of using three techniques to effectively remove the required condition and completely eliminate the unfairness problem.

7.1 FUTURE WORK

When studying the causes of structural unfairness, we pointed out that the collision probability does not lead directly to unfairness. Whether or not unfairness happens also depends on the link utilization and spare capacity needed to compensate for the retransmissions. However, the quantitative relationships among the various factors have yet to be studied.

In our study, to avoid complications from multiple sources, we use a simplified

model, assuming perfect link quality and an ideal circle of transmission. In future work, we plan to extend our study to the effect of other sources of noise. We would like to use a more realistic radio-propagation model, the shadowing model, to incorporate the multipath-fading effect.

In Chapter 4 we pointed out that our study focuses on large data packets and all our simulations use a fixed packet size of 1500 bytes. We noticed that when the packet size is comparable to the ACK packet, the effect may vary a lot. We would like to extend our study to variable length packet sizes, as well as shorter packet sizes.

Finally, we would like to generalize the three techniques to remove the structural-unfairness conditions and design SU-aware protocols for channel-assignment, routing and node placement.

REFERENCES

- [1] The network simulator - ns-2. Available at <http://www.isi.edu/nsnam/ns/>.
- [2] J.C.R. Bennett and H. Zhang. WF2Q: Worst-case fair weighted fair queueing. In *INFOCOM '96: Proceedings of the 25th IEEE International Conference on Computer Communications*, pages 120–128, March 1996.
- [3] Brahim Bensaou, Yu Wang, and Chi Chung Ko. Fair medium access in 802.11 based wireless ad-hoc networks. In *MobiHoc '00: Proceedings of the 1st ACM International Symposium on Mobile Ad Hoc Networking & Computing*, pages 99–106, Piscataway, NJ, USA, 2000. IEEE Press.
- [4] Vaduvur Bharghavan, Alan Demers, Scott Shenker, and Lixia Zhang. Macaw: a media access protocol for wireless lan's. In *SIGCOMM '94: Proceedings of the Conference on Communications Architectures, Protocols and Applications*, pages 212–225, New York, NY, USA, 1994. ACM Press.
- [5] John Bicket, Daniel Aguayo, Sanjit Biswas, and Robert Morris. Architecture and evaluation of an unplanned 802.11b mesh network. In *MobiCom '05: Proceedings of the 11th Annual International Conference on Mobile Computing and Networking*, pages 31–42, New York, NY, USA, 2005. ACM Press.
- [6] Jean-Yves Le Boudec. Rate adaptation, congestion control and fairness: A tutorial, October 2003. Available at http://ica1www.epfl.ch/PS_files/LEB3132.pdf.
- [7] Bob Briscoe. Flow rate fairness: Dismantling a religion. Available at <http://tools-ietf.org/html/draft-briscoe-tsvarea-fair-00.pdf>.

- [8] Shigang Chen and Zhan Zhang. Localized algorithm for aggregate fairness in wireless sensor networks. In *MobiCom '06: Proceedings of the 12th Annual International Conference on Mobile Computing and Networking*, pages 274–285, New York, NY, USA, 2006. ACM Press.
- [9] Qunfeng Dong, Suman Banerjee, and Benyuan Liu. Throughput optimization and fair bandwidth allocation in multi-hop wireless lans. In *INFOCOM '06: Proceedings of the 25th IEEE International Conference on Computer Communications*, pages 1–12, Barcelona, Spain, 2006.
- [10] Mustafa Ergen. IEEE 802.11 tutorial, June 2002. Available at <http://citeseer.ist.psu.edu/536785.html>.
- [11] Zuyuan Fang and Brahim Bensaou. Fair bandwidth sharing algorithm based on game theory frameworks for wireless ad-hoc networks. In *INFOCOM '04: Proceedings of the 23rd Annual Joint Conference of the IEEE Computer and Communications Societies*, volume 2, pages 1284–1295, Hong Kong, 2004.
- [12] Violeta Gambiroza, Bahareh Sadeghi, and Edward W. Knightly. End-to-end performance and fairness in multihop wireless backhaul networks. In *MobiCom '04: Proceedings of the 10th Annual International Conference on Mobile Computing and Networking*, pages 287–301, New York, NY, USA, 2004. ACM Press.
- [13] Michele Garetto, Theodoros Salonidis, and Edward W. Knightly. Modeling per-flow throughput and capturing starvation in csma multi-hop wireless networks. In *INFOCOM '06: Proceedings of the 25th IEEE International Conference on Computer Communications*, pages 1–13, Barcelona, Spain, April 2006.
- [14] P. Goyal, H. M. Vin, and H. Cheng. Start-time fair queueing: A scheduling algorithm for integrated services packet switching networks. In *SIGCOMM '96:*

Proceedings of the Conference on Communications Architectures, Protocols and Applications, volume 5, pages 690–704, October 1996.

- [15] CDMA Development Group. CDMA. Available at <http://www.cdg.org/>.
- [16] Martin Heusse, Franck Rousseau, Romaric Guillier, and Andrzej Duda. Idle sense: an optimal access method for high throughput and fairness in rate diverse wireless lans. In *SIGCOMM '05: Proceedings of the 2005 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications*, pages 121–132, New York, NY, USA, 2005. ACM Press.
- [17] Bret Hull, Kyle Jamieson, and Hari Balakrishnan. Mitigating congestion in wireless sensor networks. In *SenSys '04: Proceedings of the 2nd International Conference on Embedded Networked Sensor Systems*, pages 134–147, New York, NY, USA, 2004. ACM Press.
- [18] IEEE. IEEE 802.11 wireless local area networks. Available at <http://www.ieee802.org/11/>.
- [19] Kamal Jain, Jitendra Padhye, Venkata N. Padmanabhan, and Lili Qiu. Impact of interference on multi-hop wireless network performance. In *MobiCom '03: Proceedings of the 9th Annual International Conference on Mobile Computing and Networking*, pages 66–80, New York, NY, USA, 2003. ACM Press.
- [20] R. Jain, D. Chiu, and W. Hawe. A Quantitative Measure Of Fairness And Discrimination For Resource Allocation In Shared Computer Systems. *ArXiv Computer Science e-prints*, September 1998.

- [21] Kamran Jamshaid, Lily Li, and Paul A.S. Ward. Gateway rate control of wireless mesh networks. In *WiMeshNets '06: First International Workshop on "Wireless mesh: Moving Towards Applications"*, August 2006.
- [22] Jangeun Jun. Capacity estimation of wireless mesh networks. Master's thesis, North Carolina State University, Raleigh, North Carolina, 2002.
- [23] Jangeun Jun, Pushkin Peddabachagari, and Mihail Sichitiu. Theoretical maximum throughput of ieee 802.11 and its applications. In *NCA '03: Proceedings of the Second IEEE International Symposium on Network Computing and Applications*, page 249, Washington, DC, USA, 2003. IEEE Computer Society.
- [24] Jangeun Jun and Mihail L. Sichitiu. Fairness and QoS in multihop wireless networks. In *VTC '03: Proceedings of the 58th IEEE Vehicular Technology Conference*, pages 2936–2940, Orlando, FL, October 2003.
- [25] Jangeun Jun and Mihail L. Sichitiu. The nominal capacity of wireless mesh networks. *IEEE Wireless Communications*, 10:8–14, October 2003.
- [26] R. Karrer, A. Sabharwal, and E. Knightly. Enabling large-scale wireless broadband: the case for taps. *SIGCOMM Computer Communication Rev.*, 34(1):27–32, 2004.
- [27] Ann Lee. Achieving fairness in 802.11-based multi-channel wireless mesh networks. Master's thesis, University of Waterloo, Waterloo, Ontario, 2006.
- [28] Haiyun Luo and Songwu Lu. A topology-independent fair queueing model in ad hoc wireless networks. In *ICNP '00: Proceedings of the 2000 International Conference on Network Protocols*, page 325, Washington, DC, USA, 2000. IEEE Computer Society.

- [29] Haiyun Luo, Songwu Lu, and Vaduvur Bharghavan. A new model for packet scheduling in multihop wireless networks. In *MobiCom '00: Proceedings of the 6th Annual International Conference on Mobile Computing and Networking*, pages 76–86, New York, NY, USA, 2000. ACM Press.
- [30] Laurent Massoulié and James Roberts. Bandwidth sharing: Objectives and algorithms. In *INFOCOM '99: Proceedings of the 18th IEEE International Conference on Computer Communications*, volume 3, pages 1395–1403, March 1999.
- [31] Microsoft. Microsoft mesh connectivity layer. Available at <http://research.microsoft.com/mesh/>.
- [32] Thyagarajan Nandagopal, Tae-eun Kim, Xia Gao, and Vaduvur Bharghavan. Achieving MAC layer fairness in wireless packet networks. In *Mobicom '00: Proceedings of the 6th Annual International Conference on Mobile Computing and Networking*, pages 87–98, 2000.
- [33] Nortel Networks. Wireless mesh network: Extending the reach of wireless LAN, securely and cost-effectively (white paper). Available at <http://www.bct-global.com/new/Nortel>
- [34] T. Ozugur, M. Naghshineh, P. Kermani, and J. A. Copeland. Fair media access for wireless LANs. In *GLOBALCOM '99: Global Telecommunications Conference*, volume 1B, pages 570–579, December 1999.
- [35] Cüneyt Özveren, Robert Simcoe, and George Varghese. Reliable and efficient hop-by-hop flow control. In *SIGCOMM '94: Proceedings of the Conference on Communications Architectures, Protocols and Applications*, pages 89–100, New York, NY, USA, 1994. ACM Press.

- [36] Peter Kirstein Panos Gevros, Jon Crowcroft and Saleem Bhatti. Congestion control mechanisms and the best effort service model. *IEEE Network Magazine*, pages 16–26, May 2001.
- [37] Bozidar Radunovic and Jean-Yves Le Boudec. A unified framework for max-min and min-max fairness with applications. In *Proceedings of the 40th Annual Allerton Conference on Communication, Control, and Computing*, 2002.
- [38] Sumit Rangwala, Ramakrishna Gummadi, Ramesh Govindan, and Konstantinos Psounis. Interference-aware fair rate control in wireless sensor networks. In *SIGCOMM '06: Proceedings of the 2006 conference on Applications, Technologies, Architectures, and Protocols for Computer Communications*, pages 63–74, New York, NY, USA, 2006. ACM Press.
- [39] Ananth Rao and Ion Stoica. An overlay MAC layer for 802.11 networks. In *MobiSys '05: Proceedings of the 3rd International Conference on Mobile Systems, Applications, and Services*, pages 135–148, New York, NY, USA, 2005. ACM Press.
- [40] Jochen H. Schiller. *Mobile Communications*. Addison-Wesley, 2000.
- [41] Jingpu Shi, Theodoros Salonidis, and Edward W. Knightly. Starvation mitigation through multi-channel coordination in csma multi-hop wireless networks. In *MobiHoc '06: Proceedings of the 7th ACM International Symposium on Mobile Ad Hoc Networking and Computing*, pages 214–225, New York, NY, USA, 2006. ACM Press.
- [42] Arun K. Somani and Jianwei Zhou. Achieving fairness in distributed scheduling in wireless ad-hoc networks. In *Conference Proceedings of the 2003 IEEE International Performance, Computing, and Communications Conference*, pages 95–102, April 2003.

- [43] Dimitrios Stiliadis and Anujan Varma. Efficient fair queueing algorithms for packet-switched networks. *IEEE/ACM Transactions on Networking*, 6(2):175–185, 1998.
- [44] Andrew S. Tanenbaum. *Computer Networks*. Pearson Education, 1996.
- [45] Jian Tang, Guoliang Xue, and Weiyi Zhang. Maximum throughput and fair bandwidth allocation in multi-channel wireless mesh networks. In *INFOCOM '06: Proceedings of the 25th IEEE International Conference on Computer Communications*, pages 1–10, April 2006.
- [46] L. Tassiulas and S. Sarkar. Maxmin fair scheduling in wireless networks. In *INFOCOM '02: Proceedings of the 21st Annual Joint Conference of the IEEE Computer and Communications Societies*, volume 10, pages 320–328, June 2002.
- [47] Chieh-Yih Wan, Shane B. Eisenman, and Andrew T. Campbell. CODA: Congestion detection and avoidance in sensor networks. In *Proc. of the First International Conference on Embedded Networked Sensor Systems*, pages 266–279, 2003.
- [48] Yu Wang and J. J. Garcia-Luna-Aceves. Collision avoidance in multi-hop ad hoc networks. In *Proceedings of the 10th IEEE International Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunications Systems (MAS-COTS'02)*, pages 145–154. IEEE Computer Society, 2002.
- [49] Yu Wang and J. J. Garcia-Luna-Aceves. Modeling of collision avoidance protocols in single-channel multihop wireless networks. *Wireless Networks*, 10(5):495–506, 2004.
- [50] Wikipedia. Birthday paradox. Available at http://en.wikipedia.org/wiki/Birthday_paradox.

- [51] Wikipedia. Wireless community network. Available at http://en.wikipedia.org/wiki/Wireless_community_network.
- [52] Alec Woo and David E. Culler. A transmission control scheme for media access in sensor networks. In *MobiCom '01: Proceedings of the 7th Annual International Conference on Mobile Computing and Networking*, pages 221–235, New York, NY, USA, 2001. ACM Press.
- [53] Yuan Xue, Baochun Li, and Klara Nahrstedt. Price-based resource allocation in wireless ad hoc networks. In *Quality of Service - IWQoS 2003: 11th International Workshop*, pages 79–96, 2003.
- [54] Yung Yi and S.Shakkottai. Hop-by-hop congestion control over a wireless multi-hop network. In *INFOCOM '04: Proceedings of the 23rd Annual Joint Conference of the IEEE Computer and Communications Societies*, volume 4, pages 2548–2558, March 2004.