Providing Fairness and Maximising Throughput in 802.11 Wireless Mesh Network

Qizhi Cao

Hamilton Institute

National University of Ireland, Maynooth

Head of Department: Prof. Douglas Leith

Research Supervisor: Prof. Douglas Leith

May 4, 2010

Contents

Abstract

The thesis addresses two questions: how to provide fairness in 802.11 wireless mesh networks and how to maximize the overall throughput in an distributed way. Fairness and efficiency are two fundamental issues in wireless networks, which are widely researched in both wired and wireless networks. We consider 802.11 wireless networks due to their ubiquity and practical importance. The likely trend of the wireless networks is toward the use of multi-hop wireless networks (so-called mesh networks), which provide the potential of serving the users in larger converage with higher throughput. However, achieving fairness and efficiency remain bottleneck issues in the rollout of production-quickly mesh networks.

Chapter 1

Introduction

1.1   802.11 Wireless LAN

Over the last decades, wireless networks have become more and more popular as the last hop due to their support for mobility and flexibility. Pervasive in the workplace, the home, educational institutions, cafes, airports, and street corners, wireless LANs based on 802.11 technologies are now one of the most important access network technologies in the Internet today. There are several 802.11 standards for wireless LAN technology, including 802.11b, 802.11a and 802.11g. Table 1.1 summarizes some of the main characteristics of these standards.

1.1.1   Contention Mechanism

The IEEE 802.11 MAC defines two different access mechanisms, the mandatory Distributed Coordination Function (DCF) which provides distributed channel access based on a CSMA/CA, and the optional Point Coordination Function (PCF) which provides centrally controlled channel access through polling.

In the DCF, all stations contend for access to the medium, in a distributed manner, based on a CSMA/CA protocol. In the PCF, a Point Coordinator (PC), which is most often collocated in the AP, controls the medium access based on a polling scheme. As

Standard	Frequency Range (United States)	Data Rate
802.11b	2.4-2.485 GHz	up to 11 Mbps
802.11a	5.1-5.8 GHz	up to 54 Mbps
802.11g	2.4-2.485 GHz	up to 54 Mbps

Table 1.1: Summary of IEEE 802.11 standards

distributed operation is the most widely used, we focus here on the DCF scheme.

In the DCF, time is slotted and the channel states can be divided into three classes: idle slot, successful transmission and collision (two or more stations attempt to transmit at the same time). If there is no frame to transmit, a station chooses a random backoff value and it then decrements this each idle slot. While the channel is sensed busy, the counter value remains frozen until the channel has been detected idle for a period, known as the Distributed Inter-Frame Space (DIFS). When the counter reaches zero, the station transmits the entire frame and then waits for an acknowledgement which is sent by the receiver after a Short Inter-Frame Space (SIFS) after receiving the frame. If an acknowledgement is received, the transmitting station knows that its frame has been correctly received at the destination station. If the station has another frame to send, it begins a new random backoff. If the acknowledgement is not received, the transmitting station chooses a random value from a larger interval and counts it down again.

The random backoff value is uniformly chosen in the interval $[0, CW]$, called the Contention Window. At the first transmission attempt or after a successful transmission or discard, CW is set to the minimum value, CW_{min}. CW is doubled after each unsuccessful transmission until it reaches the maximum value, CW_{max}. An unsuccessful transmission is determined if the sender station does not receive an ACK frame within a specified ACK timeout period. After the ACK timeout period, the station assumes that a collision has occurred and enters into the backoff period again after waiting for medium to be idle for DIFS, such that the new backoff values is chosen from the doubled CW. A retransmit limit, referred to as the retry limit, is also specified in the DCF. A frame is discarded on an unsuccessful transmission after the retry limit is reached.

An additional mechanism, RTS/CTS, is defined to mitigate the hidden terminal problem found in wireless networks that use CSMA. With RTS/CTS, the sender and receiver perform a handshake mechanism by exchanging RTS (Request To Send) and CTS (Clear To Send) control frames. After waiting the DIFS time, prior to transmitting the data frame, the sender sends a RTS frame to the receiver, and the receiver responds with a CTS frame after waiting a SIFS time. The CTS frame indicates that the handshake is successful and ensures that the medium has been reserved for the particular sender and receiver for the subsequent data transmission.

1.1.2 802.11e

The IEEE 802.11 group also provides an enhanced extension to the 802.11 MAC, known as IEEE 802.11e, which aims to support quality of sevice (QoS) by introducing priority mechanisms.

IEEE 802.11e defines a new coordination function called the Hybrid Coordination Function (HCF) that combines aspects of the DCF and the PCF with enhanced QoS mechanisms. Similarly to the DCF, the Enhanced Distributed Channel Access (EDCA) is the distributed contention-based channel access mechanism of the HCF. The EDCA defines four Access Categories (ACs) for different types of traffic and a different values of MAC parameters may be used for each AC.

MAC parameters which were fixed in the previous IEEE 802.11 protocols can be adjusted in 802.11e, for example, CW_{min} and CW_{max} and additional parameters, such as AIFS and TXOP. The Arbitration Inter-Frame Space (AIFS) plays the role of DIFS in the DCF and is time the medium must be sensed idle before the transmission or backoff is started. Another new parameter is Transmission Opportunity (TXOP), the time during which the station has the right to remain transmitting once it has started a transmission. When a station wins access to the channel, it can transmit as many packets as possible such that the transmission duration does not exceed the TXOP limit. A non-zero value of TXOP Limit indicates that multiple frames might be sent in a burst and this is referred to as a Contention Free Burst (CFB). A TXOP of zero indicates that packet bursting is disabled and only one frame can be transmitted after winning channel access.

1.1.3 Rate Adaptation

The current 802.11 specification mandates the availability of multiple transmission rates at the physical layer (PHY) that use different modulation and coding schemes. For example, the 802.11b PHY supports four transmission rates (1∼11Mbps), the 802.11a PHY offers eight rates (6∼54Mbps), and the 802.11g PHY supports twelve rates (1∼54Mbps). To exploit such multi-rate capability, a sender must select the best transmission rate and dynamically adapt its selection to the time-varying and location-dependent channel quality, without explicit information feedback from the receiver.

Rate adaptation is a link-layer mechanism critical to the system performance in IEEE 802.11-based wireless networks, yet is left unspecified by the 802.11 standards. Auto Rate Fallback (ARF)[41], the first documented bit-rate selection algorithm, has been paid much

attention because of its easy implementation. Following this research, many rate adaptation algorithms are proposed. Among these algorithms, AMRR[42], ONOE[46], SampleRate[43] are implemented in the MadWifi[45] Driver.

1.2 Mesh Networks

Wireless mesh architecture offer the potential of providing a high-bandwidth network over a large coverage area. Wireless mesh networks are composed of wireless access points (routers) that facilitate the connectivity and intercommunication of wireless clients through multi-hop wireless paths. Mesh architectures sustain signal strength by breaking long distances into a series of shorter hops. Intermediate nodes not only boost the signal, but cooperatively make forwarding decisions based on their knowledge of the network, i.e. perform routing. Such an architecture with careful design may provide high bandwidth, spectral efficiency, and economic advantage over the coverage area. Wireless mesh network are distinguished from ad hoc networks by their relatively stable topology.

1.2.1 Issues in Mesh Networks

Although mesh networks offer the potential for achieving large coverage and high throughput, there are still many outstanding issues. These include channel allocation[27], efficient MAC protocols[23], routing algorithms[12], and techniques for hidden terminal mitigation.

The well-known hidden terminal effects are illustrated in Fig. 1.1. The signal can be obstructed by the buildings or other obstacles, and the signal strength is decreasing through the medium due to fading, so that some stations can not hear the transmissions from other stations. In mesh networks, however, the situation can be even worse. In the example shown in Fig. 1.1(c) where station C is hidden to A, the throughput between A and B would be zero if a same channel is used between A/B and C/D. The best solution is to use orthogonal channels, then in the same example, the througput between A and B would be the same with that between C and D.

1.2.2 Multi-Ratio Mesh Network

As mentioned previously, one way to mitigate the hidden terminal problem in mesh network is to use multi-radio techniques. The mesh network is formed from a series of WLANs, eahc operating on an orthogonal channel from its neighbor to avoid hidden terminals. Mesh
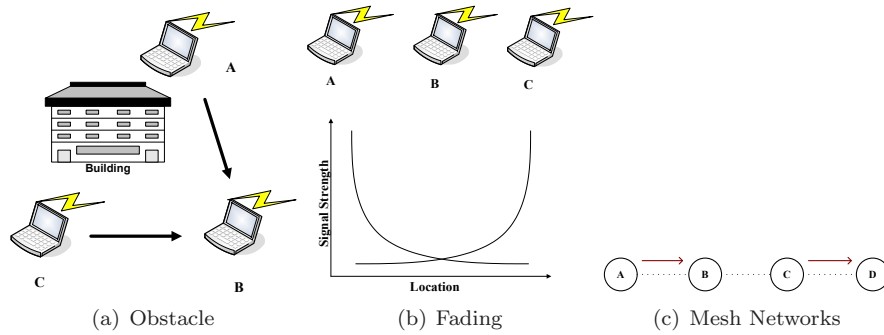
Figure 1.1: Hidden terminal problem

stations equipped with mutiple radios relay traffic between WLANs. This is illustrated for example in Fig. 1.2[38].
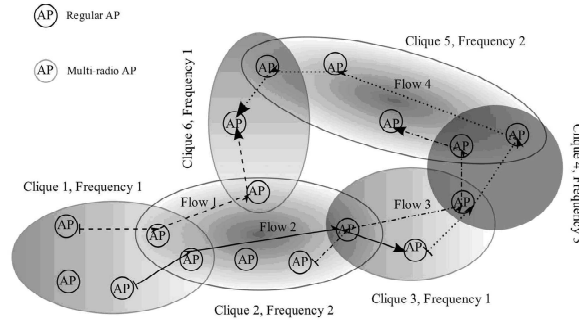

Figure 1.2: Illustrating class of mesh network.

1.3 Fairness

In most networks, there are circumstances where the externally offered load is larger than can be served by the network. Then, if no measures are taken to restrict the behaviour of traffic in the network, congestion occurs and the actual network throughput may decrease as the offered load increases due to repeated retransmission of lost packets. When the offered load must be cut back, it is important to do so fairly. Unfortunately, maximizing total network throughput is often incompatiable with fairness. Fairness can be defined in a number of different ways.

1.3.1 Max-Min Fairness and Proportional Fairness

Two important fairness concepts are max-min fairness and proportional fairness. Assume that each flow $r \in \mathcal{S}$ where $\mathcal{S}$ is the set of all the flows, has a data rate x_r :

- Max-min Fairness[4] [44]: a feasible allocation of rates $\{x_r\}$ is "max-min fair" if and

6

only if an increase of any rate within the domain of feasible allocations must be at the cost of a decrease of some already smaller rate. Formally, for any other feasible allocation $\{y_r\}$, if $y_r > x_r$ then there must exist some r' such that $x_{r'} \leq x_r$ and $y_{r'} < y_r$.

- Proportional Fairness[44]: an allocation of rates $\{x_r\}$ is "proportionally fair" if and only if, for any other feasible allocation $\{y_r\}$, we have:

$$\sum_{r=1}^{S} \frac{y_r - x_r}{x_r} \leq 0$$

In other words, any change in the allocation must have a negative average change.

Max-min fairness states that small flows receive what they demand and larger flows share the remaining capacity equally. Bandwidth is allocated equally to all flows until one is satisfied, then bandwidth is equally increased among the remainder and so on until all flows are satisfied or bandwidth is exhausted. Depending on the problem, a max-min fair allocation may or may not exist. However, if it exists, it is unique.

Proportional fairness is a compromise-based scheduling algorithm, which is based upon maintaining a balance between two competing interests: trying to maximize total wireless network throughput while at the same time allowing all users at least a minimal level of service. This is done by assigning each data flow a data rate or a scheduling priority (depending on the implementation) that is inversely proportional to its anticipated resource consumption.

1.3.2 Utility Fairness

Utility fairness[44] is a general fairness concept, which subsumes proportional fairness, minimum potential delay fairness, and max-min fairness as special cases.

Suppose that each flow r derives a utility (or benefit) of $U_r(x_r)$ when a data rate x_r is allocated to it. Then the utility fair approach is to allocate the network resources to solve the following optimization problem:

$$\max_{\{x_r, r \in \mathcal{S}\}} \sum_{r \in S} U_r(x_r)$$

subject to

$$\sum_{r:l \in \mathcal{R}_r} x_r \leq c_l, l \in \mathcal{L},$$

$$x_r \geq 0, r \in \mathcal{S},$$

where $\mathcal{R}_r$ is the set of links traversed by flow r, c_l is the capacity of link l, $\mathcal{L}$ is the set of all links and $\mathcal{S}$ is the set of all flows in the network.

When the utility function of flow r is of the form[25],

$$U_r = \begin{cases} \omega_r \frac{x^{1-\alpha_r}}{1-\alpha_r} & \alpha_r \geq 0, \alpha_r \neq 1, \\ \omega_r \log x_r & \alpha_r = 1. \end{cases}$$

Then this incorporates the following objectives:

- Rate maximization when $\alpha_r = 0$

- Proportional fairness when $\alpha_r = 1$

- Minimum potential delay fairness when $\alpha_r = 2$

- Max-min fairness when $\alpha_r \to \infty$

1.3.3   Fairness in 802.11 Wireless Mesh Networks

Fairness in wired networks has been extensively researched and globally fair rates can be attained by distributed approaches based on convex programming. However, these techniques can not be directly applied to CSMA/CA wireless networks, because the feasible rate region is nonconvex[37] [19].

Most of the literature on utility fairness in CSMA/CA networks realtes to Aloha networks rather than 802.11. It is shown in [37] that the Aloha rate region is log-convex and using this fact, distributed algorithm for achieving utility fair rate allocations are derived in [?](I guess it is not the paper about proportional fairness in aloha networks, and I don't know the paper you mention).

Recently, it has been established in [19] that the 802.11 rate region is also log-convex, and this is used in [38] to study max-min fairness in 802.11 mesh networks. We note that these results were obtained seperately from the work reported in this thesis, and are complementary to it. In [21],the fairness of TCP traffic in 802.11 WLANs is addressed and the use of 802.11e is proposed to solve the problem.

As noted previously, to counteract the loss due to noise and interference from other source, multiple data rates are implemented. However, a slow station may dominate the channel in a multi-rate network. [30] proposes the use of time-based fairness to improve performance of multi-rate WLANs.

1.4　Contribution

In the work, two fundamental issues, fairness and efficiency in 802.11 wirless networks are considered. The fairness is addressed and per-flow fairness is achieved by using TXOP. Meanwhile, the throughput efficiency is discussed by giving the theory proof and a distributed algorithm. Futhermore, part of the work is shown in [40] [39] and [38].

Chapter 2

MAC Unfairness in 802.11 Mesh Networks

2.1 Introduction

This chapter describes in more detail the MAC induced unfairness that can be found in 802.11 networks. In order to focus on MAC related issues, we consider topologies, where stations are equipped with multiple radios to avoid interference and hidden terminals. We also assume fixed routing.

2.2 Notation

Before proceeding we first describe the notation used in Fig. 2.1 and elsewhere. Client stations are marked by shadowed triangles, and mesh points (MPs) by circles. MPs are stations that relay traffic for client stations. There are 10 MPs in Fig. 2.1 where MP_9 acts as a gateway between the wireless multi-hop network and the wired Internet. Each MP has two radios that use channels in such a way that the channel in each hop is orthogonal to those in neighboring hops thereby avoiding interference between transmissions on different hops. Hence there are no hidden terminals. We assume that the set of routes from sources to destinations are already obtained by routing protocols such as those discussed in [11] and [12]. The routes are stable during the considered sessions' life time. We only consider single-path routing. We use *station* to refer to any wireless devices (both client stations and MPs). We say *client station* when referring to wireless devices other than MPs. Let
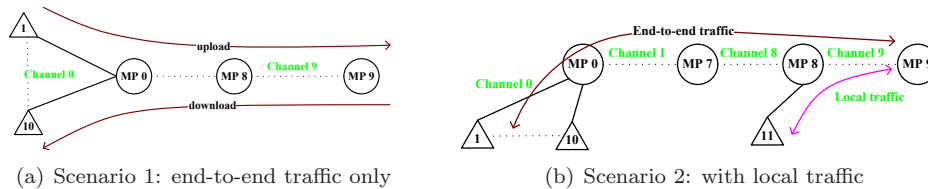
(a) Scenario 1: end-to-end traffic only   (b) Scenario 2: with local traffic

Figure 2.1: Illustrative wireless multi-hop scenarios.

n_i denote the number of *client station*s in channel i in the topologies in Fig.2.1.

2.3   802.11 Unfairness with UDP Traffic

We begin by considering UDP traffic in the topologies in Fig. 2.1, and then extend consideration to TCP traffic.

2.3.1   Without local flows

Even in the simple topology in Fig. 2.1(a), significant unfairness can exist between traffic flows. Each client station carries one upload and one download flow, and the upload and download traffic is aggregated at MP_9 which is the gateway between the wired and wireless networks. The MAC/PHY paremeters are given in Table. 2.1 and the rate of CBR traffic is 1 Mbps, so that all the stations are saturated. It can be seen that the throughput achieved by the upload flows is approximately an order of magnitude greater than that achieved by the download flows.

We can gain some insight into the source of this unfairness by looking at the corresponding per hop measurements shown in Fig. 2.2(b). It can be seen that on the relay hops, the aggregate throughput of the upload flows and of the download flows are approximately equal. However, at the left-hand hop, between the client stations and MP_0, the situation is very different.

Note that the 802.11 MAC ensures that roughly the same number of transmission opportunities are allocated to every station [21] [17], including the MPs. Thus, if there are n_0 client stations and all the stations (including MP_0) are saturated, we expect each of them to obtain roughly a $1/(n_0 + 1)$ share of the bandwidth, and similarly for MP_0 to obtain a $1/(n_0 + 1)$ share. The n_0 upload flows therefore together obtain an $n_0/(n_0 + 1)$ share whereas since all of the download flows must be transmitted via MP_0 and so they can only obtain approximately a $1/(n_0 + 1)$ share altogether. We can confirm this approximate reasoning by noting that the aggregate upload throughput at the left-hand hop in this

11

(a) Per flow throughput
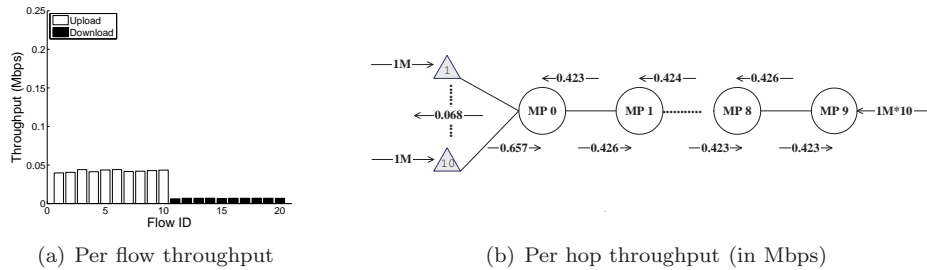

(b) Per hop throughput (in Mbps)

Figure 2.2: CBR results for scenario in Fig. 2.1(a). Per flow throughput is shown in Fig. 2.2(a). Per hop aggregate throughput (in Mbps) is plotted in Fig. 2.2(b).

T_{SIFS} (μs)	10
Idle slot duration (σ) (μs)	20
T_{DIFS} (μs)	50
CW_{min}	31
CW_{max}	1023
Retry limit	4
Packet size (bytes)	1000
PHY rate (Mbps)	1
PLCP rate (Mbps)	1

Table 2.1: MAC/PHY parameters used in simulations.

example is measured to be 0.657 Mbps while the aggregate download throughput is 0.068 Mbps. The ratio of upload to download throughput is thus 9.66, i.e. close to the value of $n_0 = 10$.

This type of unfairness is not new and has previously been observed in the context of single-hop WLANs (e.g., [21]). However, the impact of this unfairness can be far greater in a multi-hop context as we will see in the next section.

2.3.2 With local flows

To see the increased unfairness over multiple hops, consider the multi-hop network in Fig. 2.1(b) with one local client station at MP_8. End-to-end traffic from the left-hand client stations, numbered 1-10 in Fig. 2.1(b), now has to compete with the traffic from client station 11 at the MP_8 hop. The foregoing unfairness effect now acts multiplicatively at hops MP_0 and MP_8, greatly amplifying the level of unfairness. This effect is illustrated in Fig. 2.3(a). Here, stations 1-11 each carries one upload and one download flow, yielding 11 upload and 11 download flows in total. It can be seen from Fig. 2.3(a) that the upload flow at station 11 gains much greater throughput than the other flows.

What is happening is that at MP_8, each local upload flow obtains roughly a $1/(n_8 + 2)$ share of the bandwidth, where $n_8 = 1$ is the number of client stations associated with

12

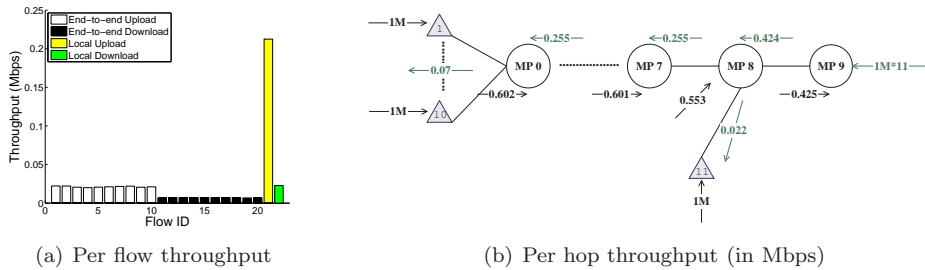(a) Per flow throughput   (b) Per hop throughput (in Mbps)

Figure 2.3: CBR results for scenario in Fig. 2.1(b). Per flow throughput is shown in Fig. 2.3(a). Per hop aggregate throughput (in Mbps) is plotted in Fig. 2.3(b). Note that 0.553 Mbps in Fig. 2.3(b) is the aggregate throughput of local upload originating from client station 11 and relay uploads from MP_7 to MP_8.

MP_8 and the 2 on the denominator accounts for end-to-end upload traffic from MP_7 and download traffic from MP_8. The aggregate upload traffic from client stations 1-10 also obtains a $1/(n_8+2)$ share (corresponding to the share of upload transmission opportunities allocated to MP_7). Thus each individual upload flow from client stations 1-10 obtains only a $1/10(n_8+2)$ share. In line with this analysis, Fig. 2.3(a) confirms that the upload flow from client station 11 obtains roughly an order of magnitude greater throughput than the upload flows from client stations 1-10.

The aggregate download traffic to client stations 1-11 also obtains a $1/(n_8+2)$ share at the MP_8 hop. The download traffic to client stations 1-10 then has to compete against the upload traffic from stations 1-10 for transmission opportunities at MP_0. This creates further unfairness. As discussed above, at the MP_0 hop there is approximately an order of magnitude unfairness between upload and download flows and this can been seen in Fig. 2.3(a).

The setup in Fig. 2.1(b), where download traffic must contend at two hops, is already sufficient to create a level of unfairness whereby download traffic to client stations 1-10 is almost starved of throughput. By introducing contention at further relay hops, the unfairness can evidently be amplified still further. In effect, the potential exists for almost arbitrary levels of unfairness to exist between competing traffic flows in a multi-hop setting. Note that this effect is not associated with interference or other sources of unfairness. Rather it is a direct consequence of the properties of the 802.11 MAC.

2.4   802.11 Unfairness with TCP Traffic

The previous section considers MAC flow unfairness with UDP traffic. However, the majority of traffic on the Internet is based on TCP. The congestion control feedback within
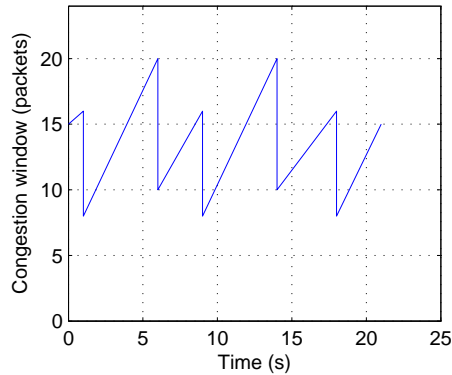
Figure 2.4: Example of TCP congestion window time history.

TCP introduces additional issues that we now consider in this section.

2.4.1 Overview of TCP

TCP is the Internet's transport-layer, connection-oriented, reliable transport protocol, that is implemented on top of the unreliable (IP) network layer. To achieve reliable data transfers, TCP receivers return ACK packets to the data senders confirming the safe arrival of data packets with lost packets being retransmitted.

To achieve congestion control, TCP regulates the number of sent but unacknowledged packets ("in flight" packets) to be no more than the congestion window *cwnd*. TCP also makes use of ACK clocking, whereby data packet transmissions primarily occur on receipt of TCP ACKs. A sender reduces its *cwnd* when a data packet loss event occurs and additively increases its *cwnd* when it perceives that the end-to-end path is congestion free. The process of TCP congestion window adjustment is illustrated in Fig. 2.4.

2.4.2 Prioritising TCP ACKs Using 802.11e

To illustrate the unfairness when using TCP in 802.11 networks, we use the topology depicted in Fig. 2.5 where there are 8 flows altogether, with end-to-end flow 0 traversing three hops, flows 1 and 2 traversing the first hop, and flows 3 – 7 traversing the third hop.

In Fig. 2.6(a) we plot the results with the 802.11 DCF as the MAC layer, i.e., both TCP data and ACK packets are put into the same queue whose parameters are given in Table 2.1. It can be seen that (i) it takes a long time for the TCP flow throughput to become steady (if ever), and (ii) flows sharing the same client station achieve significantly different throughput.

We can understand this behaviour by noting that TCP flows are bidirectional, with data
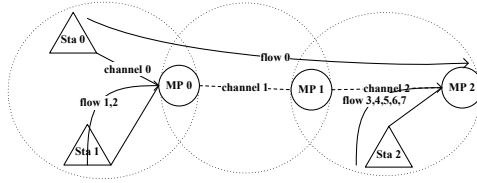
14

Figure 2.5: A linear topology for TCP flows.

packets travelling in one direction and TCP ACK packets in the reverse direction. TCP
ACK packets for all of the 8 flows must be transmitted by MP_2. As discussed in Section. 2.3
for UDP upload/download flows, this can lead to TCP ACKs becoming backlogged at
MP_2 and so result in excessive losses of TCP ACKs due to queue overflow. This potential
queuing and dropping of TCP ACKs can disrupt the TCP ACK clocking mechanism and
so hinder congestion window growth and induce repeated timeouts. This is exacerbated by
the sensitivity of flows with small *cwnd* to packet loss since retransmissions can then only
take place via timeouts rather than by fast retransmit, which tends to create sustained
unfairness between flows with large *cwnd* and those with smaller *cwnd*. To address this
problem, we collect into a single queue the outgoing TCP ACKs and assign high priority
to this queue using a small $CW_{min} = 3$, a small $CW_{max} = 7$ and a small $AIFS = 1$. The
corresponding parameters for TCP data packets use the standard DCF values $CW_{min} = 31$,
$CW_{max} = 1023$ and $AIFS = 2$. This ensures that TCP ACK packets effectively have
unrestricted access to the wireless medium. When the wireless hop is the bottleneck, data
packets will be queued for transmission and packet drops will occur there, while TCP ACKs
will pass freely with minimal queuing, i.e., the standard TCP semantics are recovered.

The effectiveness of this approach is illustrated when we compare Fig. 2.6(a) (no ACK
prioritisation) with Fig. 2.6(b) (with ACK prioritisation). We see that flow throughput is
stabilised and that flows on the same station now have approximately the same throughput.
However, denoting x_i to be the throughput achieved by MP_i, then $x_0 = 0.39$ Mbps,
$x_1 = 0.19 * 2 = 0.38$ Mbps and $x_3 = 0.08 * 5 = 0.40s$ Mbps. However, gross unfairness still
exists when we consider per-flow throughput. For example, the throughput of flow 0 is 0.4
Mbps, whereas that of flows $3 - 7$ is 0.08 Mbps.

For the previous topologies in Figs. 2.1(a) and 2.1(b), we show in Figs. 2.7(a) and 2.7(b)
the corresponding results for TCP flows when TCP ACKs are prioritised. It can be seen
that once again the network exhibits unfairness between upload and download flows and
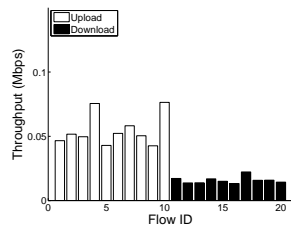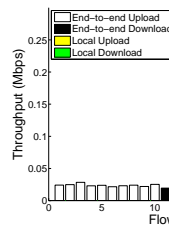between local and end-to-end flows.
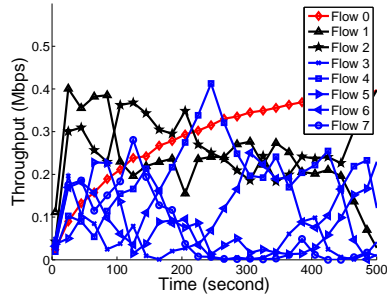
(a) 802.11 DCF
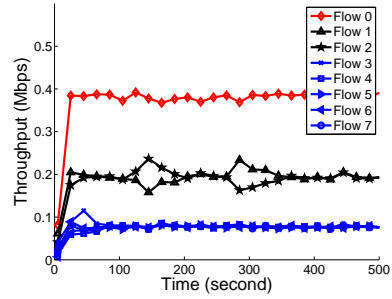
(b) TCP with prioritised ACK

Figure 2.6: Throughput fairness: TCP results in the topology shown in Fig. 2.5. Throughput values are average over 20s.

(a) TCP with prioritised ACK

(b) TCP with prioritised ACK

Figure 2.7: Measured flow throughput with prioritised TCP ACKs in Fig. 2.1(a) and 2.1(b) respectively.

16

Chapter 3

Restoring fairness

3.1 WLAN Throughput Model

We begin by considering modelling the throughput of 802.11 networks, taking TXOP packet bursting into account. Following [5], we divide time into MAC slots, where a slot corresponds to either a successful transmission, a collision (two or more simultaneous transmissions) or an idle slot (corresponding to a PHY idle slot). Consider a WLAN with n stations. Denote the slot transmission attempt probability of station i by τ_i. Let $P_{s,i}$ denote the probability of a successful transmission in a given slot by station i and P_c the corresponding collision probability. We have that:

$$P_{s,i} = \tau_i \prod_{j \neq i} (1 - \tau_j) \tag{3.1}$$

$$P_c = 1 - \sum_{i=1}^{n} P_{s,i} - P_i \tag{3.2}$$

where

$$P_i = \prod_{i=1}^{n} (1 - \tau_i), \tag{3.3}$$

is the probability that a slot is idle.

Let N_i denote the mean TXOP burst size (in packets) at station i. We can express the throughput of station i as follows:

$$s_i(\mathcal{T}, N) = \frac{E[\text{payload information transmitted in a MAC slot}]}{E[\text{length of a MAC slot}]} \qquad (3.4)$$

$$= \frac{P_{s,i} N_i L}{\sigma P_i + T_s \sum_{i=1}^{n} N_i P_{s,i} + T_c P_c} \qquad (3.5)$$

where $\mathcal{T} = [\tau_1 \; \cdots \; \tau_n]^T$, $N = [N_1 \; \cdots \; N_n]^T$, L is the packet size (in bits); T_s is the duration of a successful transmission (the sum of DATA frame duration and MAC ACK duration plus SIFS); T_c is the collision duration (the duration of DATA plus ACK timeout), σ is the duration of a PHY slot.

Letting

$$\mathcal{X} = [x_1 \; \cdots \; x_n]^T,$$

$$x_i = \tau_i / 1 - \tau_i,$$

the throughput expression can be rewritten as:

$$s_i(\mathcal{X}, N) = \frac{x_i N_i}{\sigma/T_c + \sum_{i=1}^{n} N_i \frac{T_s}{T_c} x_i + (\prod_{i=1}^{n}(1 + x_i) - 1 - \sum_{i=1}^{n} x_i)} \frac{L}{T_c} = \frac{x_i N_i}{X} \frac{L}{T_c} \qquad (3.6)$$

where

$$X = (\sigma/T_c - 1) + \sum_i (N_i \frac{T_s}{T_c} - 1) x_i + \prod_i (1 + x_i)$$

Observe that the throughput of each station is proportional to $x_i N_i$ in Equation (3.6), and the denominator X is the same for all stations. When stations are saturated (always have a packet to send) and have the same attempt probability such that $x_i = \bar{x}$, we have that:

$$\frac{s_i}{s_j} = \frac{N_i}{N_j}. \qquad (3.7)$$

For stations which are non-saturated, we assume sufficient buffering that the throughput equals to the offered load, and that the attempt probability is the value which solves Equation (3.6) for the offered load.

Modelling of TCP over wireless is challenging due to the interactions between the TCP congestions control action, interface queue dynamics and the MAC layer channel contention mechanism. However, following [21], we note that TCP ACKs are ready for transmission shortly after the corresponding TCP DATA packet is received, creating a strong time correlation between these events when the propagation delay of the wired component of

18

the TCP path is small. In this case, the TCP throughput model can be modelled using Equation (3.6) but with T_s modified to be the time required to send a TCP DATA packet and receive the TCP ACK. The accuracy of the simplified model is also shown in [21]. The result of the simplified model is very close to the simulation result.

3.2   Per-flow Throughput Fair Allocation

Let the number of flows with packets queued at *station i* be N_i and let T_i denote the TXOP duration. We consider two related approaches. The first approach is to select the TXOP duration

$$T_i = N_i \times T_s, \tag{3.8}$$

where T_s is the time for transmitting a packet. Using this approach, N_i packets are allowed to be transmitted once a transmission opportunity is won, and each packet transmission takes T_s time. Combining this TXOP allocation with the use of a modified queuing discipline that serves one packet per flow at each transmission opportunity, it follows from Equation (3.7) that this scheme ensures that backlogged flows in a WLAN are on average allocated the same throughput.

We illustrate the impact of this scheme under CBR traffic in Fig. 3.2(c) which shows the flow throughputs for the topologies in Fig. 2.1(a) and Fig. 2.1(b). The other parameters are shown in Table. 2.1. For the topology in Fig. 2.1(a), TXOP = 10 at all MPs according the approach in Equation (3.8). It can be seen from Fig. 3.1(a) that all flows are backlogged at the left-hand hop and achieve similar throughput. For the topology in Fig. 2.1(b), TXOP = 10 at $MP_0, \cdots, MP_7$, TXOP = 11 at MP_8, MP_9. All flows are backlogged at the right-hand hop and similarly achieve similar throughput in Fig. 3.1(b).

Then consider the impact of the scheme under TCP traffic. The same TXOPs are set as mentioned above. For the previous topologies in Figs. 2.1(a) and 2.1(b), we show in Figs. 2.7(a) and 2.7(b) the corresponding results when TCP ACKs are prioritised, while in Figs. 3.2(a) and 3.2(b) we show the results when in addition the proposed TXOP mechanism is used. For the topology in Fig. 2.5, flows 0 and $3-7$ are backlogged at the right-hand hop while flows 1,2 are backlogged at the left-hand hop. The capacity at each hop is not the same since the throughput is dependent on the number of contending stations, which differs at each hop. By setting the TXOPs (TXOP=2 at client station 1, TXOP=4 at client station 2), flow 0 and flows $3-7$ achieve almost the same throughput of 0.125 Mbps, while flows 1 and 2 achieve almost the same throughput of 0.33 Mbps. In all cases, with
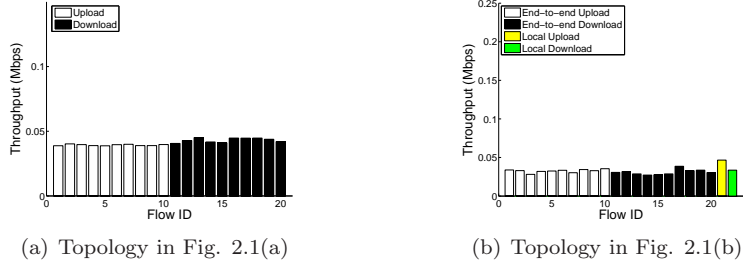
(a) Topology in Fig. 2.1(a)

(b) Topology in Fig. 2.1(b)

Figure 3.1: CBR results

(a) Topology in Fig 2.1(a)

(b) Topology in Fig 2.1(b)

(c) Topology in Fig 2.5
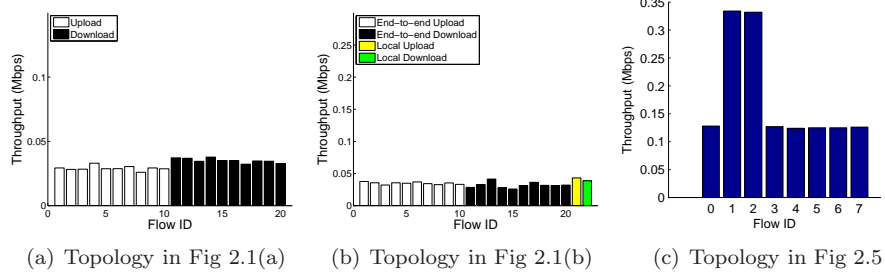
Figure 3.2: TCP results

the proposed TXOP scheme per-flow fairness is restored as expected.

3.3   The MIT Roofnet Topology

We further validate the TXOP scheme proposed in Equation (3.8) in a subset of the MIT Roofnet topology (see Fig. 3.3). In this topology, there is an Internet gateway marked as GW. Orthogonal channels are assigned in neighbouring hops so that transmissions do not interfere. The locations of the client stations and MPs are selected from data derived from the GPS coordinates of the MIT Roofnet network. There are altogether 21 TCP flows and the allocation of flows between client stations is detailed in Table 3.1. Routing for each flow is via the GW and is statically assigned as indicated by the arrows in Fig. 3.3. For example, flow 0 traverses orthogonal channels 0, 1, 3, 4 and 2 from the client station 0 to the gateway GW.

Source station(s)	Number of flows on each station	Flow ID(s)
$0-8$	1	$0-8$
9	3	$9-11$
10, 11	1	12, 20
12	2	16, 17
$13-15$	1	$15-17$
16, 17	1	13, 14

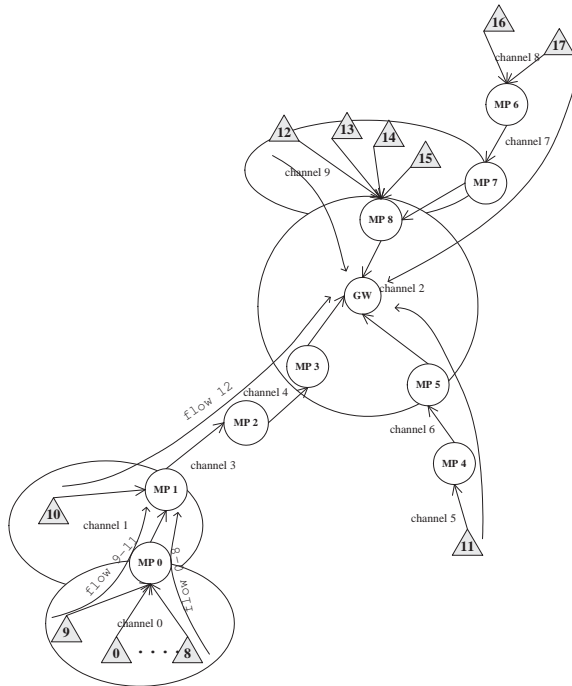Table 3.1: Flows in the roofnet topology in Fig. 3.3.

Figure 3.3: A subset of the MIT Roofnet topology.

In Fig. 3.4(a) we illustrate the resulting throughput when TCP ACKs are prioritised but the proposed TXOP scheme is not used. In this topology, all flows share channel 2 at the gateway GW which forms a shared bottleneck. Specifically, MPs 3, 5 and 8 act as relays between the gateway GW and the three branches of the network. With 802.11, these MPs share channel 2 equally, and receive roughly the same throughput. To see this observe that flow 20 obtains a throughput of 0.252 Mbps, flows $0 - 12$ achieve $0.0188 * 12 + 0.022 = 0.247$ Mbps and flows $13 - 19$ achieve $0.035 * 7 = 0.245$ Mbps. When the TXOP scheme proposed in Equation (3.8) is now used, it can be seen that per-flow fairness is now achieved, see Fig. 3.4(b).

3.4 Time-Based Fair Allocation

When multiple PHY rates are available, stations using slow PHY rates tend to dominate the channel access, resulting in fast and slow stations achieving a similar throughput[17]. For example, we select the PHY rate between client station 0 and MP_2, and that between client station 0 and MP_2 to be 11 Mbps in the topology shown in Fig. 2.5, while keeping other PHY rates at 1 Mbps. If we use the TXOP scheme in Equation (3.8), the resulting flow throughput is shown in Fig. 3.5(a) where we can see that a similar throughput allocation
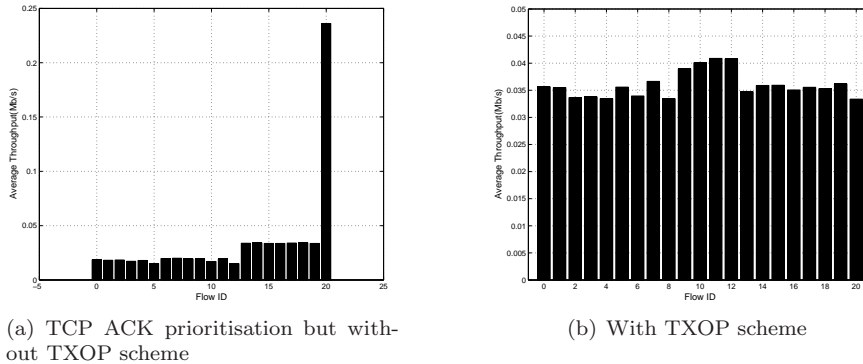
(a) TCP ACK prioritisation but without TXOP scheme

(b) With TXOP scheme

Figure 3.4: TCP flow throughputs for the topology in Fig. 3.4. Note that in Fig. 3.4(a), the throughput of flows $13 - 19$, and that of flows $0 - 12$ are equal, while in Fig. 3.4(b), the throughput of all flows is roughly equal.

to Fig. 3.2(c) is achieved. However, this is not fair for fast stations in terms of channel usage time, as transmitting the same amount of information at higher rates takes a shorter time than at slow rates. The channel usage time of the throughput-fair allocation is plotted in Fig. 3.5(b). In this figure, we plot the time used by each flow over a 100 second interval. It can be seen that the time allocated to flow 0 is much less than that to flows $3 - 7$.

Time-based fairness has been proposed to resolve this issue (see for example [30]). The rationale of time-based fairness is to allocate transmission time fairly amongst contending stations which may have different PHY rates.

Motivated by this, we consider modifying the TXOP scheme in Equation (3.8) to become

$$T_i = N_i \times T_{max} \tag{3.9}$$

where T_{max} is the time for transmitting a packet at the slowest PHY rate[1]. Again, we combine this scheme with a modified queueing discipline. Using this approach, backlogged flows at a link are on average allocated the same *air time* to transmit, regardless of their actual PHY rates. Putting it another way, flows running at a higher PHY rate may send more packets than flows with lower PHY rates. Of course when all stations use the same PHY rate, the two approaches are equivalent. However, when stations have different PHY rates, the fairness properties of the two approaches are different.

Using this scheme for the topology in Fig. 2.5, MP_0 and MP_4 use a TXOP duration $K = 9230 \ \mu s$ where $9230 \ \mu s$ is the time for transmitting a 1000-byte packet at 1 Mbps, and MP_1 and MP_3 use $T_1 = 9230 * 2 \ \mu s$ and $T_3 = 9230 * 5 \ \mu s$. In Figs. 3.6(a) and 3.6(b) we illustrate the corresponding throughput and time results. As we can see, channel usage

---

[1]The slowest rate of 802.11b/g is 1 Mbps, while that of 802.11a is 6 Mbps.
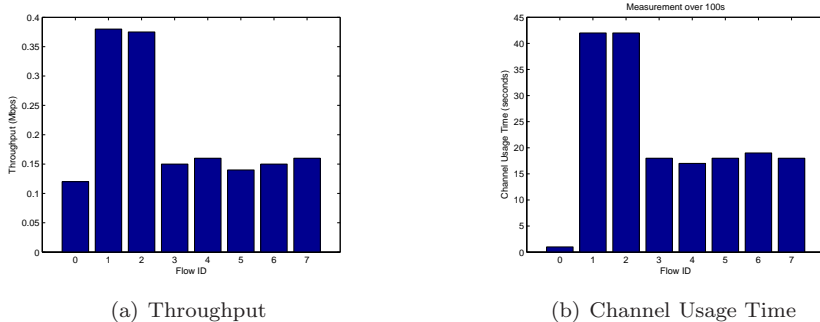
(a) Throughput (b) Channel Usage Time

Figure 3.5: Throughput-fairness allocation when multiple PHY rates are available for the scenario in Fig. 2.5. The link rates between MP_0 and MP_2, and MP_4 and MP_5 are 11 Mbps, while other link rates are still 1 Mbps. Other simulation parameters listed in Table 2.1. Note that in Fig. 3.5(a), the throughput of flow 0 is similar with that of flows $3 - 7$. In Fig. 3.5(b), the channel usage time of flows $3 - 7$ is also similar.
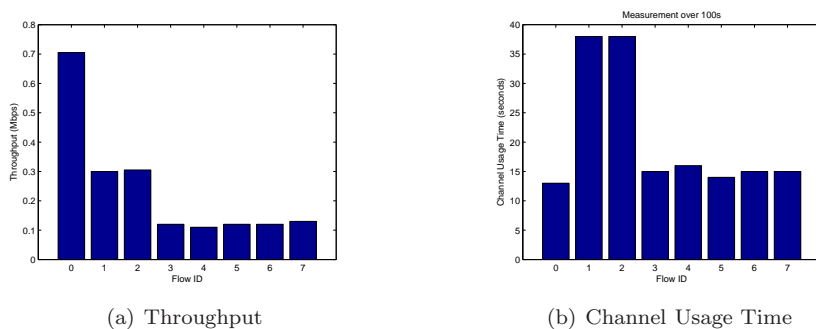

(a) Throughput (b) Channel Usage Time

Figure 3.6: Time-based fairness allocation when multiple PHY rates are available for the scenario in Fig. 2.5. The link rates between MP_0 and MP_2, and MP_4 and MP_5 are 11 Mbps, while other link rates are still 1 Mbps. Other simulation parameters listed in Table 2.1. Note that in Fig. 3.6(a), the throughput of flows $3 - 7$ is roughly equal. In Fig. 3.6(b), the time of flow 0 is almost same with that of flows $3 - 7$.

time is now per-flow fair with the throughput of flow 0 higher than others as expected. The difference between the throughput of flow 0 and that of flows $3 - 7$ is around a factor of 7, which is the same as 9230/1303. In comparison, using the throughput-fair allocation in Equation (3.8), MP_0 and MP_4 use a duration $T = 1303 \, \mu s$ where 1303 μs is the time for transmitting a packet at 11 Mbps, but MP_1 and MP_3 use $T_1 = 9230*2 \, \mu s$ and $T_3 = 9230*5 \, \mu s$.

For the MIT roofnet topology in Fig. 3.3, we illustrate the effectiveness of the time-based allocation Equation (3.9) by increasing all the link rates from station 11 to the gateway GW to be 11 Mbps, while keeping other link rates at 1 Mbps. Other simulation parameters are listed in Table 2.1. The resulting flow throughputs using the throughput-fair allocation Equation (3.8) and the time-based allocation Equation (3.9) are shown in Figs. 3.7(a) and 3.7(c) respectively. In Fig. 3.7(c), we also plot the channel usage time of

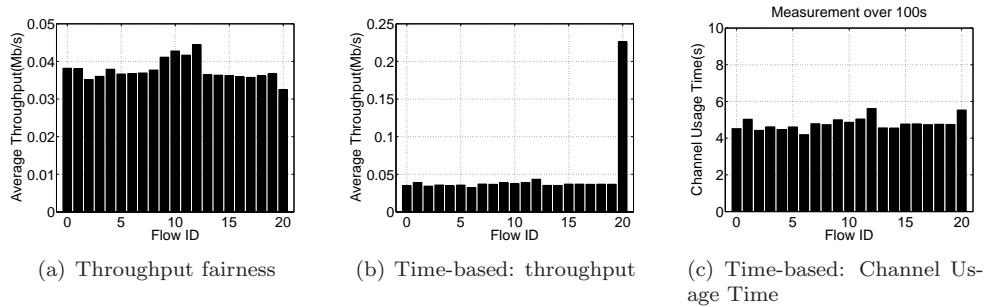| (a) Throughput fairness | (b) Time-based: throughput | (c) Time-based: Channel Usage Time |

Figure 3.7: Time-based fairness allocation when multiple PHY rates are available for the scenario in Fig. 3.3. All the link rates from station 11 to the gateway GW are 11 Mbps, while other link rates are still 1 Mbps. Other simulation parameters listed in Table 2.1.

each flow over 100 seconds. It is evident that per-flow fairness in terms of channel usage time is achieved.

3.5 Max-Min Fairness

While the foregoing simulations consider specific topologies, we comment that it has subsequently been shown[32] that the TXOP scheme in Equation (3.8) achieves max-min fairness in general topologies. When all stations use the same PHY rates, while the TXOP scheme in Equation (3.9) achieves time-based weighted max-min fairness[13].

3.6 Impact of lossy channel

Now we briefly extend consideration to include losses due to channel noise. Consider the two hop network in Fig. 3.8. There are altogether 21 TCP flows in the example: 10 one-hop flows in each hop and 1 end-to-end flow (Flow 10) traversing two hops. In this topology, we expect that each flow should have the same throughput if the resulting allocation is fair, and can be seen from Fig. 3.9(a) that this is indeed this case.
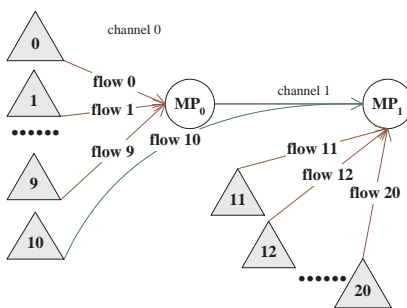

Figure 3.8: Topology with lossy link.

(a) BER=0

(b) BER=10^{-6}
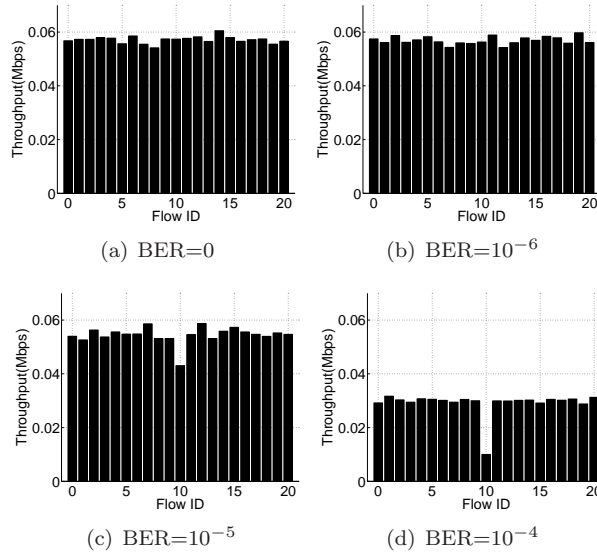
(c) BER=10^{-5}

(d) BER=10^{-4}

Figure 3.9: Throughput allocation vs BER.

Figures 3.9(b), 3.9(c) and 3.9(d) show the goodput allocations with lossy links. Here we think each link has the same loss rate. It can be seen that for a BER of 10^{-6}, the noise related losses have little impact on fairness compared to the noise-free case. For 1000 Byte packets a BER of 10^{-6} corresponds to a packet loss rate of 0.796%, which is even negligible. This is consistent with the recent theoretical results in [34] which establish that the max-min fair allocation varies smoothly as the level of losses is increased. Thus when the loss rate is low it will have only a small impact on fairness.

At higher BERs fairness starts to degrade significantly. From the results at a BER of 10^{-5} ($PER \approx 7.69\%$) which is already higher that is likely to be encounted in production networks, it can be seen that a higher level of channel losses eventually leads to significant unfairness – the two-hop flow now achieves about 20% less throughput than that of the one-hop flows. When each link has a BER of 10^{-4} ($PER \approx 55.1\%$), the throughput of the two-hop flow is only half of that of the one-hop flows. The unfairness arises because the two-hop flow sees a greater end-to-end loss rate than the one-hop flows.

Chapter 4

Experimental measurements

We implemented the topologies shown in Figure ? and Figure ? using a testbed con-
structed from Soekris net4801 stations with Atheros 802.11a/b/g miniPCI cards. The NICs
used support 802.11e EDCF functionality which makes adjustable the MAC parameters
AIFS, CWmin and TXOP. All stations run Linux 2.6.21.1 with a version of the MADWiFi
driver customised to allow the priorisations described in this paper. Otherwise, tests were
performed in infrastructure mode using standard 802.11a parameters and channels with
PHY rate 6Mbps. To implement dual-radio mesh relay points, we joined two net4801s at
100Mbps with a cross-over cable to form a single logical mesh point. This avoided po-
tential interference between network cards sharing the same PCI backplane. Routing in
the network was statically congured. Iperf was used to generate TCP traffic and data was
collected from both iperf and tcpdump. TCP ows used SACK enabled NewReno.

Chapter 5

Improving network capacity

In this chapter we consider improving network capacity while maintaining flow fairness and decentralised operation.

5.1 Maximising WLAN throughput

5.1.1 Homogeneous Stations

We begin by considering a homogeneous network, where all stations have the same attempt probability $\bar{\tau}$ and are saturated, i.e always have a packet to send at every transmission opportunity.

The network throughput model has been given in Section 3.1. Let $\bar{x} = \dfrac{\bar{\tau}}{1 - \bar{\tau}}$. Equation (3.6) can be rewritten as follows:

$$s_i(\bar{x}, N) = \frac{\bar{x} N_i}{X} \frac{L}{T_c} \tag{5.1}$$

where

$$X = (\sigma/T_c - 1) + \bar{x} \sum_i (N_i \frac{T_s}{T_c} - 1) + (1 + \bar{x})^n.$$

Setting the derivative of s_i with respect to $\bar{x}$ to zero leads to:

$$\frac{ds_i}{d\bar{x}} = \frac{N_i L}{T_c} \frac{1}{X^2} (X - \bar{x} \sum_i (N_i \frac{T_s}{T_c} - 1) - n(1 + \bar{x})^{n-1}) = 0 \tag{5.2}$$

Equation (5.2) can be rewritten as:

$$(\sigma/T_c - 1) + (1 + \bar{x})^n (1 - n\bar{x}/(1 + \bar{x})) = 0 \qquad (5.3)$$

That is

$$1 - n\bar{\tau} = \epsilon P_i. \qquad (5.4)$$

where $\epsilon = 1 - \sigma/T_c$, i.e.

$$1 - \sum_{j=1}^{n} \tau_j = \epsilon P_i. \qquad (5.5)$$

Solving for $\bar{x}$ yields the attempt probability that maximizes the stations' (and so WLAN) throughput.

5.1.2   Heterogeneous Stations

Consider now a WLAN with n stations, with station i having offered load s_i. We partition the WLAN into $n - m$ stations with attempt probability $\bar{x}$ such that $s_j \geq \dfrac{\bar{x}}{X}\dfrac{L}{T_c} = \bar{s}, j = m+1, \cdots, n$ and m stations with attempt probability such that $s_j = \dfrac{x_j}{X}\dfrac{L}{T_c}$ and $x_j < \bar{x}, j = 1, 2, \cdots, m$. That is, we partition stations into those which are able to service their offered load and those which cannot. We seek the value of $\bar{x}$ that maximizes the value $\bar{x}/X$. We have that

$$X = \sigma/T_c - 1 + \prod_{j=1}^{m}(1 + x_j)(1 + \bar{x})^{n-m}$$

Differentiating,

$$\frac{d\bar{x}/X}{d\bar{x}} = \frac{d\bar{x}/X}{dX}\frac{dX}{d\bar{x}}$$

Now

$$\frac{dX}{d\bar{x}} = (n - m)\prod_{j=1}^{m}(1 + x_j)(1 + \bar{x})^{n-m-1} > 0$$

Hence, setting the derivative equal to zero observe that

$$\frac{d\bar{x}/X}{d\bar{x}} = 0 \Rightarrow \frac{d\bar{x}/X}{dX} = 0$$

We have that

$$\frac{d\bar{x}/X}{dX} = \frac{\frac{1}{n-m}Y^{\frac{1}{n-m}-1}\frac{dY}{dX}X - \{Y^{\frac{1}{n-m}} - 1\}}{X^2} \qquad (5.6)$$

where

$$Y = \frac{X + \epsilon}{\prod_{j=1}^{m}(1 + Xs_j)}$$

Setting $\dfrac{d\bar{x}/X}{dX}$ equal to 0, we obtain

$$(n - m)\bar{\tau} + \sum_{j=1}^{m} \tau_j + \epsilon P_i = 1 \tag{5.7}$$

i.e.

$$1 - \sum_{j=1}^{n} \tau_j = \epsilon P_i \tag{5.8}$$

which is identical in form to equation (5.4) in the homogeneous case. Solving for $\bar{\tau}$ yields the attempt probability that maximises throughput.

5.2 Decentralized Optimisation

Observe from Equation (5.8) that the optimal $\bar{\tau}$ depends on $\tau_j, j = 1, 2, \cdots, m$, so message-passing between stations is required in order to determine $\bar{\tau}$. In this section we consider decentralised (i.e. no message-passing) suboptimal methods for achieving high throughput. Motivated by the idle sense approach[18] we consider constraining the attempt probabilities such that the probability P_i that a MAC slot is idle satisfies $P_i \geq p$ for some fixed choice of p.

5.2.1 P_i Constraint

We begin by considering the minimum value of p such that the P_i constraint is always active at or before the rate region boundary. This is illustrated by subset P_0 in Fig. 5.1. Along any ray from the origin, the boundary P_0 is reached at or before the boundary of the rate region. In other words, p is the smallest value such that $P_i \leq p$ at the rate region boundary.

From the Arithmetic Mean-Geometric Mean inequality, we have that:

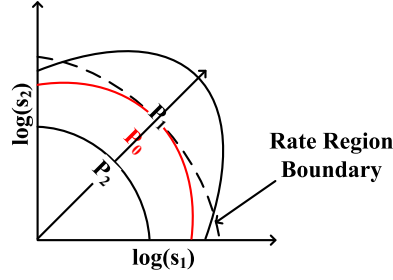$$\sum_i \frac{1}{1 + x_i} \geq n \sqrt[n]{\prod_i (\frac{1}{1 + x_i})}$$

Figure 5.1: P_0, P_1, P_2 illustrate subsets of the rate region corresponding to different values of p in constraint $P_i \geq p$. P_0 corresponding to the smallest value of p such that $P_i \leq p$ at the rate region boundary.

Hence,

$$\sum_j \tau_j + \epsilon P_i = \sum_i \frac{x_i}{1 + x_i} + \frac{\epsilon}{\prod_i (1 + x_i)}$$

$$= \sum_i (1 - \frac{1}{1 + x_i}) + \frac{\epsilon}{\prod_i (1 + x_i)} \qquad (5.9)$$

$$\leq n\{1 - \sqrt[n]{\prod_i (\frac{1}{1 + x_i})}\} + \frac{\epsilon}{\prod_i (1 + x_i)}$$

with equality when $n = 1$ or $x_1 = x_2 = \cdots = x_n, n > 1$, i.e. when moving along the 45° ray marked in Fig. 5.1.

That is,

$$n(1 - \sqrt[n]{P_i}) + \epsilon P_i \geq \sum_i \tau_i + \epsilon P_i. \qquad (5.10)$$

Recall from Equation (5.8) that the rate region boundary is defined by the solution to $\sum_i \tau_i + \epsilon P_i = 1$. Now we consider the monotonicity of the following function,

$$f(n, P_i) = n(1 - \sqrt[n]{P_i}) + \epsilon P_i,$$

which is strictly increasing with respect to n, and strictly decreasing with respect to P_i. Hence,

$$n(1 - \sqrt[n]{P_i}) + \epsilon P_i < \lim_{n \to \infty} n(1 - \sqrt[n]{P_i}) + \epsilon P_i = -\ln P_i + (1 - \eta)P_i$$

where $\epsilon = 1 - \eta$. Let p be the solution to $1 = -\ln p + (1 - \eta)p$. Constraining the idle probability such that $P_i \geq p$ then we have from Equation (5.10) that $\sum_i \tau_i + \epsilon P_i \leq 1$. That is, the constraint on P_i becomes active at or before the rate region boundary constraint
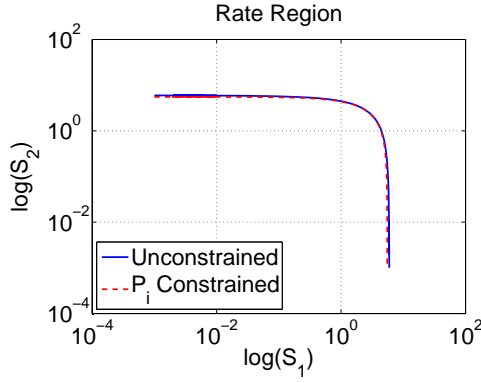
30

Figure 5.2: The boundary of rate region with/without constraint on P_i is shown, for $n = 2$ stations, PHY rate=11Mbps, L=1000bytes, $\sigma/T_c = 67.1$ and $p = 0.8412$.

$\sum_i \tau_i + \epsilon P_i = 1$ as required.

Using a Taylor Series expansion we have that, $-\ln p = (1-p) + \dfrac{(1-p)^2}{2} + \dfrac{(1-p)^3}{3} + \cdots \geq (1-p) + \dfrac{(1-p)^2}{2}$. Solving $(1-p) + \dfrac{(1-p)^2}{2} + (1-\eta)p \leq 1$, we have the lower bound on p

$$p \geq 1 + \eta - \sqrt{2\eta} \tag{5.11}$$

to ensure that the $P_i \geq p$ constraint is always active at or before the rate region boundary is reached.

5.2.2 Throughput Efficiency

By constraining the idle probability such that $P_i \geq p$, there will be loss of throughput compared to the maximum possible throughput. However, it turns out that the loss is so small that it is essentially negligible. This is illustrated in Fig. 5.2 where the unconstrained and constrained rate regions are depicted and from which we can see that the constrained rate region is log-convex[19] and the gap between the two rate regions is small.

To see the difference between the achieved throughput and the maximum possible throughput in more detail, the ratio of achieved to maximum possible throughput is plotted in Fig. 5.3(a). In this figure, data is shown for three different approaches, namely for $P_i = p$ with p chosen according to Equation (5.11), using the attempt probability derived by Bianchi[5] to approximately maximize throughput, and using the Idle Sense approach in [18]. Results are shown for all the available PHY rates in 802.11b (1, 2, 5.5, 11 Mbps). It can be seen that there is at most 8 percent throughput loss, even in the worst case when only one station is transmitting. Additionally, the loss becomes extremely small (less than

1 percent lost) when there are more than 2 stations in the channel. It can be also seen that the higher the PHY rate used, the smaller the throughput loss.
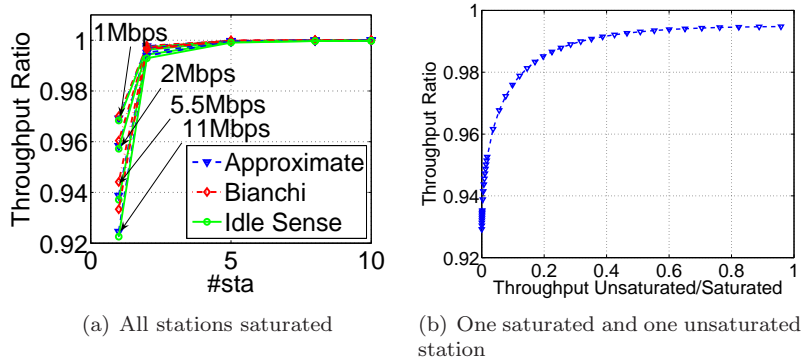

(a) All stations saturated

(b) One saturated and one unsaturated station

Figure 5.3: Illustrating efficiency of approximate throughput maximisation using P_i constraint. L=1000bytes; in Fig. 5.3(b), $\frac{\sigma}{T_c}$=67.2, p=0.8412, PHY rate=11Mbps

Fig. 5.3(b) shows a second example of a network with one saturated stations and one unsaturated station. The throughput efficiency is plotted versus the offered load on the unsaturated station. The throughput ratio is again at most 8 percent, becoming smaller as the offered load increases. In Fig. 5.3(b), when the throughput of the unsaturated station is zero then only one saturated station is active and the throughput efficiency agrees with that in Fig. 5.3(a). When the throughput of unsaturated station equals that of the saturated station, two saturated stations are active and again the righthand paint in Fig. 5.3(b) can be compared with Fig. 5.3(a).

5.2.3 Decentralised implementation

Algorithm

We can see that threshold p in constraint Equation (5.11) provides a design parameter that allows us to control the throughput efficiency in a WLAN. Importantly, since the channel idle probability P_i is observable by all stations in a WLAN the potential exists to enforce constraint (5.11) in a decentralised manner, with no need for message passing.

We can regulate P_i to satisfy constraint Equation (5.11) by adjusting the station CW_{min} values to be sufficiently large. Evidently there are many possible combinations of CW_{min} satisfying the constraint. In particular, this might be achieved in an unfair manner by having some stations with very large CW_{min} and others with small CW_{min}. Such unfair solutions are not of interest in the present context. We borrow from the TCP congestion control literature the use of AIMD to provide a decentralized mechanism for achieving

32

Algorithm 1 Proposed ADMI tuning algorithm.
---
   for Every T seconds **do**
     if $P_i > p$ **then**
       $CW_{min} \leftarrow CW_{min} - \alpha$
     else
       $CW_{min} \leftarrow CW_{min}/\beta$
     end if
   end for
---

fairness. Specifically, we propose Algorithm 1. Providing all stations use the same ADMI (rather than AIMD) decrease rate α and increase factor β, under mild conditions we are guaranteed [33] that stations will, on average, converge to the same value of CW_{min} while respecting constraint $P_i \geq p$ as required.

We briefly comment that the use of channel idle probability has previously been considered in [9] in the context of measuring the link quality in a WLAN. The use of Idle Sense is also considered in [18], [16] in the context of maximizing throughput in a infrastructure mode WLAN. However, this work focusses mainly on optimising throughput in WLANs rather than achieving fairness in mesh networks.

Design Parameters

It can be seen that Algorithm 1 contains four design parameters α, β, p and T. We consider these in turn.

The choice of update interval T is determined by the time required to obtain an accurate estimate of the idle probability P_i. As a rough guideline, if the measurement noise is roughly white and gaussian we expect that the standard deviation of our P_i estimate is proportional to $1/\sqrt{N}$ where N is the number of measurements available. A long update interval T therefore allows more observations and so a more accurate estimate. However, this also slows convergence of Algorithm 1 (see below). In this paper, we use $T = 1$ second as a reasonable compromise. We note that this is also a common upper bound of TCP RTT on the Internet and reflects the time-scale over which network conditions are likely to change [36] (currently, WLAN users are mainly using the Internet and the main traffic type is TCP i.e. 80%–90% of traffic is TCP [31]).

From the AIMD analysis in [33], we have that the mean time between backoff events in Algorithm 1 is proportional to $\alpha/T(1 - \beta)$. The mean time to converge to the stationary distribution is proportional to $1/\log \beta$ backoff events. We have found values of $\beta = 0.75$ (corresponding to a convergence time of roughly 10 backoff events), and $\alpha = 4$ to yield good performance across a wide range of network conditions.
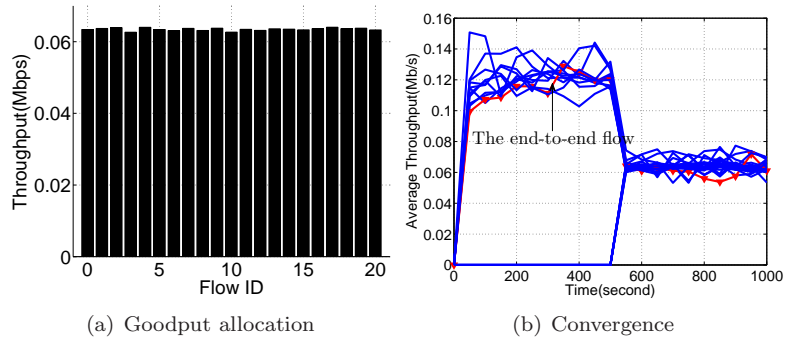
(a) Goodput allocation   (b) Convergence

Figure 5.4: Goodput allocations when the joint TXOP/CW_{min} algorithm is used on the topology in Fig. 3.8.

The choice of p has already been discussed in Section 5.2.1.

Simulation Results

In this section we use simulations to evaluate the fairness performance achieved when the TXOP adaptation approach of Chapter 3 is augmented with the CW_{min} autotuning Algorithm 1. This yields a joint TXOP/CW_{min} tuning algorithm.

We begin by revisiting the example in Section 3.6 with network topology shown in Fig. 3.8. Fig. 5.4(a) shows the throughput allocations obtained with the joint TXOP/CW_{min} tuning algorithm. It can be seen that all flows, i.e. including both one-hop and two-hop flows, converge to essentially the same goodputs as required. And it is obvious that the goodput within the joint TXOP/CW_{min} approach is about 10.9% higher than that with TXOP-only one.

Shown in Fig. 5.4(b) is the convergence of the throughput allocations following a change in the network load. Namely, initially the network starts with 5 flows active on each hop and one 2 hop flow. At time 500s an additional 5 flows start at each hop. It can be seen that when the traffic load is increased the network is able to quickly adapt to the new channel state.

Synchronization of CW_{min} Updates

In the above figures, all stations update their CW_{min} at the same time (i.e. with time synchronization.) This might be achieved in practice by synchronising updates to the access point beacons. Fig. 5.5 illustrates behaviour without time synchronization. It can be seen that throughput fairness is not achieved. This happens because the first station to increase its CW_{min} when $P_i < p$ may lead to P_i subsequently increasing such that $P_i \geq p$, in which case other stations may not increase their CW_{min}s since they do not observe the
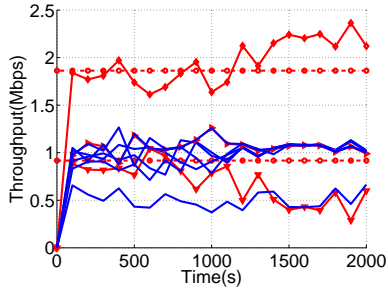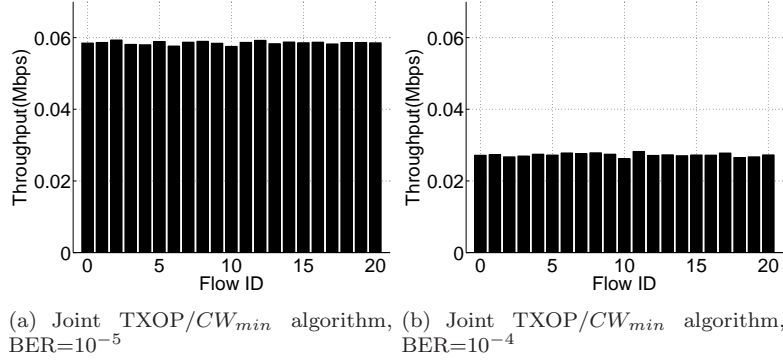
Figure 5.5: Without time synchronization of CW_{min} updates.

(a) Joint TXOP/CW_{min} algorithm, BER=10^{-5}

(b) Joint TXOP/CW_{min} algorithm, BER=10^{-4}

Figure 5.6: Goodput allocation vs BER.

temporary violation of the $P_i \geq p$ constraint, leading to unfairness. Synchronisation of CW_{min} updates prevents such behaviour.

5.3 Channel Noise

In this section we briefly consider the impact of losses due to channel noise. Figures 5.6(a) and 5.6(b) show the goodput allocations for a range of channel BER values for the topology of Figure 3.8 when the joint TXOP/CW_{min} algorithm is used. It can be seen that the flow fairness is still maintained for BERs up to 10^{-4}.

Chapter 6

Conclusions

This thesis presented two TXOP algorithms that provides per-flow fairness instead of per-station fairness in the DCF and a distributed algorithm that maximises the total throughput of the mesh network. It addressed the fairness problem in 802.11 wireless mesh networks and especially the problem under TCP traffic, also discussed the impact of lossy links.

This thesis gave the boundary of the mesh networks and a feasible way to maximise the total throughput in a distributed manner. The algorithm can give about 10% more throughput and work well in lossy channels.

Bibliography

[1] Part 11: wireless LAN medium access control (MAC) and physical layer (PHY) specifications: Medium Access Control (MAC) Quality of Service (QoS) Enhancements, IEEE 802.11e/D8.0, February 2004.

[2] P. Bahl, R. Chandra, and J. Dunagan, "SSCH: Slotted seeded channel hopping for capacity improvement in ieee 802.11 adhoc wireless networks," in Proc. ACM MOBI-COM, Philadelphia, Sep. 2004, pp. 216-230.

[3] B. Li and R. Battiti, "Supporting service differentiation with enhancements of the IEEE 802.11 MAC protocol: models and analysis," University of Trento, technical report, DIT-03-024, May 2003.

[4] D. Bertsekas and R. Gallager, *Data Networks*, Prentice-Hall 1987.

[5] G. Bianchi, "Performance analysis of the IEEE 802.11 distributed coordination function," *IEEE Journal on Selected Areas in Communications*, vol. 18, No. 3, pp. 607-614, Mar. 2000.

[6] M. Bottigliengo, C. Casetti, C. F. Chiasserini, and M. Meo, "Short-term Fairness for TCP Flows in 802.11b WLANs," in *Proc. of IEEE INFOCOM*, Mar. 2004, pp. 1383-1392.

[7] V. Badarla, D. Malone, and D. Leith, "Implementing TCP Flow-Level Fairness Using 802.11e in a Multi-Radio Mesh Testbed," *IEEE Communications Letters*, vol. 12, no. 4, pp. 262–264, Apr. 2008.

[8] D. Leith and P. Clifford, "A Self-Managed Distributed Channel Selection Algorithm for WLANs," in *ACM/IEEE RAWNET*, Apr. 2006.

[9] D. Malone, P. Clifford, and D. leith, "MAC Layer Channel Quality Measurement in 802.11," *IEEE Communications Letters*, vol.11, no.2, pp. 143-145, Feb. 2007.

[10] D. Malone, K. Duffy, and D. Leith, "Modeling the 802.11 Distributed Coordination Function in Nonsaturated Heterogeneous Conditions," *IEEE/ACM Transcactions on Networking*, vol. 15, no. 1, pp. 159-172, Feb. 2007.

[11] D. De Couto, D. Aguayo, J. Bicket, and R. Morris, "A high-throughput path metric for multi-hop wireless routing," in *Proc. of ACM MobiCom*, Sep. 2003, pp. 134-146.

[12] R. Draves, J. Padhye, and B. Zill, "Comparison of routing metrics for static multi-hop wireless networks," in *Proc. of ACM SIGCOMM*, Aug. 2004.

[13] K. Duffy, D. Leith, T. Li, and D. Malone, "Modeling 802.11 Mesh Networks," *IEEE Communication Letters*, vol. 10, no. 8, pp. 635-637, Aug. 2006.

[14] V. Gambiroza, B. Sadeghi, and E. W. Knightly, "End to End Performance and Fairness in Multihop Wireless Backhaul Networks," in *Proc. of ACM MOBICOM*, Sep. 2004.

[15] M. Garetto, T. Salonidis, and E. W. Knightly, "Modeling Per-flow Throughput And Capturing Starvation In CSMA Multi-hop Wireless Networks," in *Proc. of IEEE IN-FOCOM*, Apr. 2006.

[16] Y. Grunenberger, M. Heusse, F. Rousseau, and A. Duba, "Experience with an Implementation of the Idle Sense Wireless Access Method," in *Proc. of ACM CoNext*, 2007.

[17] M. Heusse, F. Rousseau, G. Berger-Sabbatel, and A. Duda, "Performance anomaly of 802.11b," in Proc. IEEE INFOCOM, San Francisco, Mar. 2003, pp. 836-843.

[18] Martin H., Franck R., Romaric G. and Andrzej D. "Idle Sense: An Optimal Access Method for High Throughput and Fairness in Rate Diverse Wireless LANs", IGCOMM'05, August 21-26, 2005 Philadelphia, Pennsylvania, USA

[19] Vijay S. and Douglas L. "Log Convexity of Rate Region in 802.11e WLANs"

[20] Jean-Yves Le Boudec, "Rate Adaptation, Congestion Control and Fairness: A Tutorial," on line.

[21] D. Leith, P. Clifford, D. Malone, and A. Ng, "TCP Fairness in 802.11e WLANs," *IEEE Communications Letters*, vol. 9, no. 11, pp. 964–966, Jun. 2005.

[22] T. Li, D. Leith, D. Malone and V. Badarla, "Achieving End-to-end Fairness in 802.11e Based Wireless Multi-hop Mesh Networks," in Proc. Chinacom, Hangzhou China, Aug. 2008.

[23] R. Maheshwari, H. Gupta, S. R. Das, "Multichannel MAC Protocols for Wireless Networks," in Proc. IEEE SECON. Reston, VA, Sep. 2006, vol. 2, pp. 393-401.

[24] L. Massoulié and J. Roberts, "Bandwidth Sharing: Objectives and Algorithms," *IEEE/ACM Transcactions on Networking*, vol. 10, no. 3, pp. 320-328, Feb. 2002.

[25] J. Mo and J. Walrand, "Fair end-to-end window-based congestion control," IEEE/ACM Transactions on Networking, vol. 8, no. 5, pp. 556-567, Oct. 2000.

[26] K. Ramachandran, E. Belding-Royer, K. Almeroth, and M. Buddhikot, "Interference-Aware Channel Assignment in Multi-Radio Wireless Mesh Networks," in *Proc. of IEEE INFOCOM*, Apr. 2006.

[27] B. Raman, "Channel Allocation in 802.11-based Mesh Networks," in *Proc. of IEEE INFOCOM*, Apr. 2006.

[28] A. Raniwala, P. De, S. Sharma, R. Krishnan, and Tzi-cker Chiueh, "End-to-End Flow Fairness over IEEE 802.11-based Wireless Mesh Networks," in *Proc. of IEEE INFOCOM* Mini-Symposium, May 2007.

[29] R. Stanojevic, and R. Shorten, "Beyond CHOKe: Stateless fair queueing," in *Proc. of EuroFGI NET-COOP* 2007.

[30] G. Tan, and J. Guttag, "Time based Fairness Improves Performance in Multi-rate WLANs," in Proc. USENIX, Boston, MA, Jun. 2004.

[31] D. Tang and M. Baker, "Analysis of A Local-Area Wireless Network," in *Proc. of ACM MobiCom*, Aug. 2000.

[32] I. Tinnirello and S. Choi, "Temporal Fairness Provisioning in Multi-Rate Contention-Based 802.11e WLANs," in *Proc. of IEEE WOWMOM*, Jun. 2005.

[33] R.N. Shorten, F. Wirth, D.J. Leith, "A positive systems model of TCP-like congestion control: Asymptotic results". IEEE/ACM Transactions on Networking, 14(3), pp616-629, June 2006.

[34] V. G. Subramanian, K. R. Duffy and D. J. Leith, "Existence and uniqueness of fair rate allocations in lossy wireless networks", IEEE Transactions on Wireless Communications, in press.

[35] H. Wu, F. Yang, K. Tan, J. Chen, Q. Zhang, and Z. Zhang, "Distributed Channel Assignment and Routing in Multi-radio Multi-channel Multi-hop Wireless Networks," Journal on Selected Areas in Communications, vol.24, pp. 1972-1983, Nov. 2006.

[36] Z. Zhao, S. Darbha, and A. L. N. Reddy, "A Method for Estimating the Proportion of Nonresponsive Traffic At a Router," *IEEE/ACM Transactions on Networking*, vol. 12, no. 4, pp. 708–718, Aug. 2004.

[37] K. Kar, S. Sarkar, and L. Tassiulas, "Achieving Proportional Fairness Using Local Information in Aloha networks," IEEE Transactions on Automatic Control, vol.49, No.10, Oct. 2004.

[38] Douglas J. Leith, Qizhi Cao, Vijay G. Subramanian, "Max-min Fairness in 802.11 Mesh Networks," CoRR abs/1002.1581, 2010.

[39] Qizhi Cao, Tianji Li and Douglas Leith, "Achieving Fairness in Lossy 802.11e Wireless Multi-Hop Mesh Networks," MeshTech2009, co-located with IEEE MASS, October 12, 2009, Macau SAR, P.R. China.

[40] Tianji Li, Douglas Leith, Venkataramana Badarla, David Malone, and Qizhi Cao, "Achieving End-to-end Fairness in 802.11e Based Wireless Multi-Hop Mesh Networks Without Coordination," Mobile Networks and Applications (MONET), Sep. 2009.

[41] A. Kamerman and L. Monteban, "WaveLAN-II: A High-performance wireless LAN for the unlicensed band," AT&T Bell Lab Technical Journal, pages 118–133, Summer 1997.

[42] Lacage Mathieu, Manshaei Mohammad Hossein and Turletti Thierry, "IEEE 802.11 rate adaptation: a practical approach," MSWiM 2004 ACM, pp. 126 – 134, October 2004, Italy

[43] J. Bicket, "Bit-rate selection in wireless networks," Master's thesis, MIT, 2005.

[44] Kelly, "Charging and rate control for elastic traffic," European Transactions on Telecommunications, vol. 8, pp 33-37, 1997

[45] Multiband Atheros Driver for Wifi, http://madwifi.org/

[46] Onoe Rate Control Algorithm, http://madwifi.org/browser/madwifi/trunk/ath_rate/onoe/