

Achieving End-to-end Fairness in 802.11e Based Wireless Multi-Hop Mesh Networks Without Coordination

Tianji Li · Douglas J. Leith ·
Venkataramana Badarla · David Malone · Qizhi Cao

© Springer Science + Business Media, LLC 2009

Abstract To mitigate the damaging impacts caused by interference and hidden terminals, it has been proposed to use orthogonal channels in wireless multi-hop mesh networks. We demonstrate however that even if these issues are completely eliminated with perfectly assigned channels, gross unfairness can still exist amongst competing flows which traverse multiple hops. We propose the use of 802.11e's TXOP mechanism to restore/enforce fairness. The proposed scheme is simple, implementable using off-the-shelf devices and fully decentralised (requires no message passing).

Keywords medium access control (MAC) · transmission control protocol (TCP) · wireless mesh networks · IEEE 802.11 · IEEE 802.11e · max-min fairness

This work is supported by Science Foundation Ireland Grant 03/IN3/I396 and Irish Research Council for Science, Engineering and Technology.

T. Li (✉) · D. J. Leith · V. Badarla · D. Malone · Q. Cao
Hamilton Institute, National University of Ireland
Maynooth, Maynooth, Co. Kildare, Ireland
e-mail: tianji.li@nuim.ie

D. J. Leith
e-mail: doug.leith@nuim.ie

V. Badarla
e-mail: badarla.venkataramana@nuim.ie

D. Malone
e-mail: david.malone@nuim.ie

Q. Cao
e-mail: cao.qizhi@nuim.ie

1 Introduction

We consider ensuring fairness in the backhaul of 802.11 based multi-hop mesh networks. In this type of backhaul networks, if the mesh points (MPs, i.e., wireless routers that compose the backhaul) are equipped with a single radio, it is well known that the network performance can suffer due to the hidden terminal effect. For example, it has been observed that due to this reason end-to-end traffic over more than 3 hops tends to achieve rather limited throughput [13]. With the aim of supporting high performance therefore, neighbouring MPs are normally equipped with multiple radios which operate in orthogonal channels, with 802.11a channels being preferred since these channels are less utilised than the 802.11b/g ones, see e.g., [18, 26, 27] and the references therein.

However, we first demonstrate that even when the channels in the considered backhaul networks are perfectly assigned (i.e., orthogonal channels are always used in neighbouring hops and there is no hidden terminal effect observed), gross unfairness can still exist amongst competing flows. This unfairness is associated with fact that 802.11 DCF ensures that all stations obtain a same transmission opportunity regardless of their bandwidth demands, physical layer rates, etc. To restore/enforce fairness therefore, proper prioritisation mechanisms should be designed.

In this paper, we propose the use of 802.11e's TXOP mechanism to restore/enforce fairness. An immediate advantage of this method is that it is standard compatible and requires no modifications to the current hardware and protocol stacks. In particular, we show in this paper that both i) throughput fairness and ii) time-based fairness can be readily achieved using TXOP.

In general, the choice of the most appropriate fairness requirement is essentially a policy decision for the network operator. Our main point is that TXOP does indeed provide the necessary flexibility to allow fairness to be controlled in a simple and practical manner using off-the-shelf devices and fully decentralised (requires no message passing). We demonstrate the efficacy of this approach with both NS simulations and test-bed implementation. Since TCP currently carries the vast majority of network traffic (in both wired [36] and wireless networks [32]) it is important to investigate the performance of the proposed scheme with TCP, we consider both CBR and TCP traffic in this paper.

2 Related work

In single-hop WLANs, existing approaches normally rely on either contention window size tuning to achieve per-station fairness (e.g., [35]) or rate control at the MAC layer (e.g., [31]) to achieve per-station fairness when stations use different physical layer rates. Neither of them however are standard compatible and thus they require modifications of hardware/protocols/etc. For the former type of methods, arbitrary contention window sizes are potential required, with however that standard sizes being the powers of two [1]. For the latter, specific algorithms should be designed and implemented at the MAC layer. Note also that per-flow rather than per-station fairness is the aim of this paper.

Most previous work in multi-hop networks has focussed on issues such as hidden terminals and interference (e.g., [3, 21, 26, 27, 34]). MAC-related unfairness has been studied in the context of single-hop 802.11 WLANs, e.g., see [7, 17] and references therein. However, fairness in multi-hop networks has received limited attention. In single-channel multi-hop networks, [13] illustrates that unfairness exists in parking lot deployments, and a congestion control algorithm is proposed to mitigate unfairness in [28]. The unfairness issue in [13] and [28] is caused by hidden terminals and interference. There has been even less work regarding the use of the TXOP mechanism. In [33], the authors evaluate the use of TXOP for stations with different physical rates.

To the best of our knowledge, there exists no other prior work on enforcing/restoring per-flow fairness using 802.11e's TXOP in multi-hop networks besides our early work [19, 20] in which early results on throughput fairness are reported. In this paper, we extend significantly [19, 20] with results on simulation results using the MIT roofnet topology (derived from the GPS




Fig. 1 Real slot durations in the 802.11e TXOP mechanism

coordinates file¹), and results on time-based fairness in both simulations and test-bed experiments.

3 Background

3.1 802.11e and TXOP

The 802.11e MAC protocol [2] extends the standard 802.11 DCF (Distributed Coordinated Function) contention mechanism by allowing the adjustment of MAC parameters that were previously fixed. With 802.11, on detecting the wireless medium to be idle for a period *DIFS*, each station initializes a counter to a random number selected uniformly from the interval $[0, CW-1]$ where *CW* stands for contention window. Time is slotted and this counter is decremented each slot that the medium is idle. An important feature is that the countdown halts when the medium becomes busy and only resumes after the medium is idle again for a period *DIFS*. On the counter reaching zero, the station transmits a packet. If a collision occurs (two or more stations transmit simultaneously), *CW* is doubled and the process repeated. On a successful transmission, *CW* is reset to the value CW_{min} and a new countdown starts for the next packet. The 802.11e MAC enables the values of *DIFS* (called *AIFS* in 802.11e) and CW_{min} to be set on a per class basis for each station. Four separate classes at each station is specified in 802.11e. Packets from each class are put into a separate queue. Different parameters including *AIFS*, TXOP, CW_{min} , CW_{max} can be assigned to each class/queue so that differentiation can be realised.

The TXOP mechanism specifies a duration during which a station can keep transmitting without releasing the channel once it wins a transmission opportunity. In order not to release the channel, a SIFS interval

¹<http://www.pdos.lcs.mit.edu/roofnet/roofnet-coords>




Fig. 2 Illustrative wireless multi-hop scenarios. **a** Scenario 1: end-to-end traffic only. **b** Scenario 2: with local traffic

is inserted between each packet-ACK pair (Fig. 1). A successful transmission round consists of multiple packets and ACKs. By adjusting this time, the number of packets that may be transmitted by a station at each transmission opportunity can be controlled. A salient feature of the TXOP operation is that, if a large TXOP is assigned and there are not enough packets to be transmitted, the TXOP period is ended immediately to avoid wasting bandwidth.

4 Unfairness of 802.11

Before proceeding we first describe the notations used (see Fig. 2). 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 both topologies. 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 [10] and [11]. 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. Unless otherwise stated, all CBR traffic has offered load equal to the raw physical layer rate so will lead to saturation of the client stations.

Interestingly, even with such a simple network setup (no interference/hidden terminals, fixed routing, standard 802.11 parameters) significant unfairness can exist between traffic flows. This is illustrated in Fig. 3a which shows the measured throughput achieved by a mix of upload and download CBR flows (with one upload and one download flow per client station; sources for download flows and destinations for upload flows lie in the wired network). 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.




Fig. 3 CBR results for scenario in Fig. 2a. One upload and one download flow per client station. Per flow throughput is shown in Fig. 3a. Per hop aggregate throughput (in Mbps) is plotted

in Fig. 3b. Simulation parameters listed in Table 1. **a** Per flow throughput. **b** Per hop aggregate throughput (in Mbps)

Table 1 MAC/PHY parameters used in simulations

T_{SIFS} (μs)	10
Idle slot duration (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

We can gain some insight into the source of this unfairness by looking at the corresponding per hop measurements shown in Fig. 3b. It can be seen that on the relay hops, the aggregate throughput of the upload flows and of the download flows are approximately equal, as might be expected. However, at the left-hand hop, between the client stations and MP_0 , the situation is very different. We can understand this behaviour by noting that the 802.11 MAC ensures that roughly the same number of transmission opportunities are allocated to every station [15, 17], including the MPs. Thus, if there are n_0 client stations, we expect each of them to obtain roughly a $1/(n_0 + 1)$ share of the bandwidth, and similarly for the MP 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 the MP 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 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., [17]). However, the impact of this unfairness can be far greater in a multi-hop context.

To see this, consider the multi-hop network in Fig. 2b with one local station at MP_8 . End-to-end traffic from the left-hand stations, numbered 1–10 in Fig. 2b, now has to compete with the traffic from 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. 4a. 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. 4a 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 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 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 stations 1–10 obtains only a $1/10(n_8 + 2)$ share. In line with this analysis, Fig. 4a confirms that the upload flow from station 11 obtains roughly an order of magnitude greater throughput than the upload flows from stations 1–10.

The aggregate download traffic to stations 1–11 also obtains a $1/(n_8 + 2)$ share at the MP_8 hop. The download traffic to 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.




Fig. 4 CBR results without TXOP for scenario in Fig. 2b with one client station at MP_8 . Per flow throughput is shown in Fig. 4a. Per hop aggregate throughput (in Mbps) is plotted in Fig. 4b. Simulation parameters listed in Table 1. Note that 0.553 Mbps

in Fig. 4b is the aggregate throughput of local upload originating from station 11 and relay uploads from MP_7 to MP_8 . **a** Per flow throughput. **b** Per hop throughput (in Mbps)

As discussed above, at the MP_0 hop there is approximately an order of magnitude unfairness between upload and download flows and this can be seen in Fig. 4a.

The setup in Fig. 2b, where download traffic must contend at two hops, is already sufficient to create a level of unfairness whereby download traffic to 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.

5 Achieving fairness using TXOP

Since the unfairness behaviour noted above is associated with the MAC operation, it is natural to seek to improve fairness by investigating changes at the MAC layer. In this paper, we propose the use of 802.11e’s TXOP mechanism to restore/enforce fairness.

5.1 The proposed scheme

Let the number of flows with packets queued at MP_i on channel l be n at a transmission opportunity. Let $K_{l,i}$ denote the TXOP duration of the transmission opportunity. We consider two related approaches.

5.1.1 Throughput fair allocation

The first approach is to select the TXOP duration

$$K_{l,i} = n \times T_{now}, \tag{1}$$

where T_{now} is the time for transmitting a packet at the current PHY rate. Using this approach, n packets are allowed to be transmitted once a transmission opportunity is won, and each packet transmission takes T_{now} time. Combining this TXOP allocation with the use of a modified queuing discipline (e.g., [29]) that serves one packet per flow at each transmission opportunity, we will show below that this throughput fair scheme ensures that backlogged flows at a link are on average allocated the same *number* of packet transmissions.

5.1.2 Time-based fair allocation

The second approach is to select

$$K_{l,i} = n \times T_{max} \tag{2}$$

where T_{max} is the time for transmitting a packet at the slowest PHY rate.² 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, in more realistic cases where stations may have different PHY rates, the fairness properties of the two approaches are different. In general, the choice of the most appropriate fairness requirement is essentially a policy decision for the network operator. Our main point here is that TXOP does indeed provide the necessary flexibility to allow fairness to be controlled in a simple and practical manner. In particular, we show in this paper that i) throughput fairness and ii) time-based fairness can be readily achieved using the proposed approaches.

In this section, we provide a theoretical model to analyse the effectiveness of these two approaches. Simulations and test-bed results will be given in the following sections to validate the analysis and to illustrate how the three mentioned fairness requirements are met.

5.2 Modelling TXOP

We design a finite-load model to quantify TXOP’s functionality. We use the approach proposed by Bianchi in [6] and extended in [22] to allow us to calculate the impact of TXOPs.

In multi-hop CSMA/CA based networks, modelling the relay traffic distribution from a previous hop is still an open problem. Following common practice (e.g., [12, 14]) we assume that the offered load at station i is an independent Poisson process with mean rate of λ_i bits/sec.

We therefore consider an intermediate hop between the source and the destination with relaying MP denoted as MP' and $n - 1$ associated MPs/client stations. The quantity of interest is the throughput of station i (recall that by station, we mean both MPs and client stations) i

$$x_i = \frac{P_{i,s} E[L_i]}{E[T]}, \tag{3}$$

where $P_{i,s}$ is the probability that station i has a successful transmission, $E[L_i]$ is the expected number of bits

²The slowest rate of 802.11b/g is 1 Mbps, while that of 802.11a is 6 Mbps.

transmitted in a transmission, and $E[T]$ is the expected slot duration (in seconds).

Let K_i be the TXOP duration in packets. For the first approach (Eq. 1), $K_i = n$. While for the second approach (Eq. 2), $K_i = n \times T_{max}/T_{now}$. Let τ_i be the probability that station i attempts transmission, and p_i be the probability of station i collides with others in a real slot time. Following [22], we assume that for each station i there is a constant probability $1 - q_i$ that the station's queue has no packets awaiting transmission in an expected slot. The probability q_i that one or more packets are available in $E[T]$ time is given by $q_i = 1 - e^{(-\lambda_i/K_i)E[T]}$.

Using a similar coupling technique as in [6], the probability τ_i can be modelled as a function of p_i and q_i using a Markov chain for the contention windows (see Equation (6) in [22]). A second relation relating τ_i and p_i is

$$1 - p_i = \prod_{j \neq i} (1 - \tau_j), \tag{4}$$

i.e., there is no collision for station i when all other stations are not transmitting. With n stations, p_1, \dots, p_n and τ_1, \dots, τ_n can be solved numerically.

Let P_{tr} be the probability that at least one station is transmitting, we then have that

$$P_{tr} = 1 - \prod_{i=1}^n (1 - \tau_i). \tag{5}$$

Let $P_{i,s}$ be the probability that station i successfully wins a transmission opportunity (which may involve transmitting one or multiple packets), then

$$P_{i,s} = \tau_i \prod_{j \neq i} (1 - \tau_j), \tag{6}$$

and combining with Eq. 4, we have that

$$P_{i,s} = \tau_i (1 - p_i). \tag{7}$$

Let P_c be the probability that more than one station starts transmissions at the same time, we have that

$$P_c = P_{tr} - \sum_{i=1}^n P_{i,s}. \tag{8}$$

Now we can represent the expected slot duration as

$$E[T] = (1 - P_{tr})\sigma + \sum_{i=1}^n (P_{i,s}T_{i,s}) + P_cT_c. \tag{9}$$

where σ is the idle slot duration, T_c is the collision duration, and $T_{i,s}$ is the successful duration. In the non-TXOP case, both T_c and $T_{i,s}$ correspond to a

packet transmission and associated overhead, while in the TXOP case multiple packets can be transmitted.

There are two variables ($T_{i,s}$ and $E[L_i]$) in Eq. 3 that are still unknown, with their relationship being that $T_{i,s} = E[L_i]/R + \Delta$ where R (bits/sec) denotes the physical rate, and Δ (in seconds) denotes the overhead including DIFS for 802.11 (AIFS for 802.11e), SIFS and ACKs. For calculating $E[L_i]$, we use an approximation that station i always waits until there are enough packets to transmit in one TXOP (as we will see that analysis with this assumption matches the simulations well), hence $E[L_i] = K_i * L$ where K_i is the TXOP duration in packets and L is the packet size in bits. The aggregate overhead in one TXOP is thus $\Delta_i = DIFS(or AIFS) + K_i(2 * SIFS + T_{ack} + 2 * T_{phy,hdr} + T_{mac,hdr} + T_{other,hdr})$. The model is now complete.

This analysis is verified against simulations. We use a process with mean rate of 64 Kbps as the application traffic. An upload and a download flow are associated with each client station and the MP' . The packet size, physical data rate and physical basic rate used is 80 bytes, 11Mbps and 1Mbps, respectively. The other parameters are listed in Table 1. In Fig. 5a, we illustrate the results for both TXOP enabled 802.11e (labelled as TXOP in the figure³) and 802.11 DCF (labelled as non-TXOP).


It can be seen that (i) as the number of flows increases, in both cases the system throughput increases to a maximum level and remains steady thereafter, (ii) the use of TXOP allows higher throughput to be sustained compared with the 802.11 DCF. In Fig. 5b, the individual throughput achieved by MP' and client stations is depicted. We can see that the throughput achieved by the 802.11 DCF drops rapidly when there are more than 12 pairs of traffic flows. With TXOP, however, MP' can maintain a near constant throughput after the channel becomes saturated. In both cases, user-stations throughput decrease slightly with the number of traffic flows.

For stations which are backlogged, we have that the probability $q_i = 1$. According to Eqs. 4, 6, we know that these saturated stations have the same transmission success probability (represented as P_s^*) in a slot. The throughput ratio between these stations is thus proportional to their TXOPs. i.e.,

$$\frac{x_i}{x_j} = \frac{P_s^* E[L_i]}{P_s^* E[L_j]} = \frac{K_i}{K_j}. \tag{10}$$

³As all links use the same PHY rate, TXOP's selected according to Eqs. 1 and 2 are the same.

Fig. 5 Model vs. simulation for TXOP and 802.11 DCF. We use a process with mean rate of 64 Kbps as the application traffic. An upload and a download flow are associated with each client station and the MP' . The packet size, physical data rate and physical basic rate used is 80 bytes, 11 Mbps and 1 Mbps, respectively. The other parameters are listed in Table 1. **a** System throughput. **b** Individual throughput



Recall that all stations are using the same parameters such as CW_{min} , CW_{max} , AIFS, etc.

In practice, traffic may not be saturated which means that the assigned TXOP durations may not always be utilised sufficiently. To model the fairness in this case, we define the effective TXOP duration K'_i used by station i to be

$$K'_i = \frac{P_{i,s} E[L_i] / L}{P_s^*} \quad (11)$$

where $P_{i,s}$ is the actual successful transmission probability, L is the packet length. Observe that $K'_i = K_i$ for saturated stations, but $K'_i \leq K_i$ for stations which are not persistently saturated. That is, saturated stations can use up to the maximum assigned TXOP, but non-saturated stations can not. The advantage of working in terms of K'_i is that the throughput ratio between any stations can be written as

$$\frac{x_i}{x_j} = \frac{P_{i,s} E[L_i]}{P_{j,s} E[L_j]} = \frac{K'_i}{K'_j}, \quad (12)$$

i.e., this relationship holds for both saturated and non-saturated stations. This equation says that the ratio of throughput achieved by any two stations is equal to the ratio of their TXOPs, which provides the basis for the proposed TXOP approaches.

5.3 Remarks

Note that TXOP may change from transmission opportunity to transmission opportunity as the mix of queued packets varies and so the scheme automatically adapts to changes in the number of flows carried by a station. In practice, this dynamic TXOP allocation scheme can be simplified to select $K_{l,i}$ to equal the average number of flows carried by station i , and by employing FIFO queuing (rather than per-flow fair queueing) with little loss in performance—see the examples below. There is

no message passing required since each station is able to determine the number of flows it carries by inspection of its outgoing packet stream and thus the scheme is fully decentralised, greatly facilitating management and roll-out.

It is important to note that for a station that is assigned a long TXOP length, if during a transmission opportunity it has no packets to send (the network interface queue is empty) then that transmission opportunity is ended automatically. That is, if the offered load at a station is too low to make full use of its allocated TXOP share (or due to burstiness of the traffic, the interface queue is empty from time to time), the excess is not lost but rather becomes available on a best effort basis for use by other stations in the network.

We also comment that with this TXOP approach a station transmits n packets in a single burst. For large n , this can result in the station occupying the channel for a substantial consolidated period of time and this may, for example, negatively impact competing delay-sensitive traffic. We can address this issue in a straightforward manner by using multiple smaller TXOPs instead of a single one. When using smaller packet TXOPs, it is necessary to ensure a corresponding increase in the number of transmission opportunities won by the station. This can be achieved by using a smaller value of CW_{min} for the prioritised traffic class at the station. It is shown in [4] that competing traffic classes gain transmission opportunities approximately in inverse proportion to their values of CW_{min} . Let k denote the ratio of the stations CW_{min} value to the base value used in the network (e.g. 31 in 802.11b/g). Scaling k with the number of transmission opportunities required provides coarse (recall that in 802.11e k is constrained to be a power of two) prioritisation of downstream flows. We then complement this with use of TXOP for fine grained adjustment of the packet burst lengths. For example, when $k = 2$ we halve the

value of CW_{min} and also halve the value of TXOP to $n/2$. Hence, fine grained prioritisation can always be achieved while avoiding unduly large packet bursts.

6 Throughput fairness

In this section, we introduce how to ensure/enforce throughput fairness using the proposed throughput fair approach in Eq. 1.

6.1 CBR traffic

We revisit the earlier multi-hop examples, and illustrate the impact of the proposed TXOP assignment scheme in Eq. 1 with CBR traffic. For the topology in Fig. 2a, Fig. 6a demonstrates the impact of this change—it can be seen that fairness is restored between upload and download flows. For the second topology in Fig. 2b, the corresponding results are shown in Fig. 6b. Again, it is evident that fairness is restored.

6.2 TCP traffic

We now consider fairness behaviour with TCP traffic. This is of key importance since the vast majority of network traffic is TCP (for both wired [36] and wireless networks [32]).

6.2.1 802.11 unfairness for TCP

To achieve reliable data transfers, TCP receivers return ACK packets to the data senders confirming the safe arrival of data packets. In wireless networks however, TCP ACK packets without any prioritisation can be easily queued/dropped [17]. 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.

To clearly illustrate this source of unfairness when supporting TCP flows in 802.11 networks, we use the linear topology depicted in Fig. 7. This kind of topology is often used in studying end-to-end traffic in both wired networks (e.g., [16, 23, 24]) and wireless networks (e.g., [13]). In Fig. 7, there are 8 flows altogether, with end-to-end flow 0 traversing three hops, local flows 1 and 2 in the first hop, and local flows 3, 4, 5, 6 and 7 in the third hop. Local flows in a same hop share the same station.

In Fig. 8a we plot the results with 802.11 as the MAC layer, i.e., both TCP data and ACK packets are put into the same queue whose parameters are given in Table 1. It can be seen that (i) it takes a long time for TCP flows to become steady (if ever), and (ii) flows sharing the same MP achieve significantly different throughput.

6.2.2 Prioritising TCP ACKs using 802.11e

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 are $CW_{min} = 31$, $CW_{max} = 1023$ 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. 8a (no ACK prioritisation) with Fig. 8b (with ACK prioritisation). We see that the resulting MP throughput is approximately the same. That is, 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 $x_3 = 0.08 * 5 = 0.40$ s Mbps. However, gross unfairness




Fig. 6 Throughput fairness: CBR results with TXOP per station fairness scheme. TXOP = 10 at MP_0 , TXOP = 10 at MP_7 , TXOP = 11 at MP_8 . **a** Scenario 1. **b** Scenario 2




Fig. 7 Scenario 3: a linear topology for TCP flows




Fig. 8 Throughput fairness: TCP results in the topology shown in Fig. 7. For prioritising TCP ACKs, TCP ACK packets are put into high priority queue with $CW_{min} = 3$, $CW_{max} = 7$ and $AIFS = 1$, whereas TCP data packets are put into low priority with $CW_{min} =$

31, $CW_{max} = 1023$ and $AIFS = 2$. Other simulation parameters listed in Table 1. Note that in Fig. 8b, the throughput of flows 3–7 overlaps, and in Fig. 8c, the throughput of flows 0, 3–7 overlaps. **a** 802.11. **b** TCP with prioritised ACK. **c** TXOP

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.

6.2.3 Achieving fairness for TCP flows

Using the proposed TXOP approach in Eq. 1, per-flow fairness can be readily restored. In Fig. 8c, the resulting flow throughput for the topology in Fig. 7 is plotted. For this parking lot topology the resulting allocation is close to max-min fairness [5]. According to [23], the vector

$$\left\{ \frac{c_2}{6}, \frac{(c_0 - \frac{c_2}{6})}{2}, \frac{(c_0 - \frac{c_2}{6})}{2}, \frac{c_2}{6}, \frac{c_2}{6}, \frac{c_2}{6}, \frac{c_2}{6}, \frac{c_2}{6} \right\}$$

is the unique max-min allocation where c_i is the current capacity in channel i .

Here, $c_2 = 0.75$ Mbps and $c_0 = 0.785$ Mbps. The capacity at each hop is not the same since 802.11 throughput is dependent on the number of contending stations, which differs at each hop. Flow 0 and flows 3, 4, 5, 6, 7 achieve almost the same throughput of 0.125 Mbps,

while flow 1 and 2 achieve almost the same throughput of 0.33 Mbps.

For the previous topologies in Fig. 2a and b, we show in Figs. 9a and 10a the corresponding results when TCP ACKs are prioritised, while in Figs. 9b and 10b we show the results when in addition the proposed TXOP mechanism is used. In all cases, per-flow fairness is restored as expected.

6.2.4 The impact of channel noise

In the foregoing simulations the channel is error free and packet losses are solely due to buffer overflow and MAC-layer collisions. In practice, of course, we also expect packet losses to be induced by channel noise. However, provided the level of losses is sufficiently small we can expect that these will have only a small impact on network fairness (this is in line with recent analytic results on utility fairness in lossy networks (see [30])). Figure 11a, b and c illustrate the impact of various levels of noise loss in the case when a similar level of noise is present on all links. Results are shown for




Fig. 9 TCP results for topology in Fig. 2a. Simulation parameters listed in Table 1. **a** TCP with prioritised ACK. **b** TXOP




Fig. 10 TCP results for topology in Fig. 2b. Simulation parameters listed in Table 1. **a** TCP with prioritised ACK. **b** TXOP




Fig. 11 The impact of interference on TCP traffic in the topology shown in Fig. 7. Simulation parameters listed in Table 1. Note that in Fig. 11a, b, c and d, the throughput of flows 3–7 overlaps. a

an independent and identically distributed noise model when the bit error rate (BER) is 10^{-6} , 10^{-5} and 10^{-4} for the topology shown in Fig. 7. It can be seen that even though the channel capacity decreases when the BER increases, for BER levels up to and including 10^{-4} the impact on fairness is negligible.

Figures 11a, d and e show corresponding results when the BER level in channel 2 is increased from 10^{-6} to 10^{-4} while other links' BER remains to be 10^{-6} . We can expect that this asymmetry on noise losses will increase the severity of the impact of noise on fairness. Nevertheless, for BERs up to 10^{-5} it can be seen that the impact is negligible, and even at a BER of 10^{-4} the impact remains small.

We note that a BER of 10^{-4} corresponds to a frame error rate of approximately 55% for 1000 byte frames and so represents a high level of noise loss. A rate adaptation algorithm is typically used to adjust the coding and modulations to maintain the BER below a target level, reflecting application and transport layer requirements. For TCP in particular, since TCP congestion control views packet losses as an indicator of congestion, TCP throughput is strongly dependent on

BER= 10^{-6} . **b** BER= 10^{-5} . **c** BER= 10^{-4} . **d** BER= 10^{-6} and 10^{-5} in channels 0 and 2. **e** BER= 10^{-6} and 10^{-4} in channels 0 and 2

the link loss rate (e.g., [8, 9]) and too high a loss rate may then prevent high utilisation of the wireless channel. Normally we would expect that rate adaptation algorithms would be used to select a PHY rate that maintained the loss rate significantly below 10^{-4} .

6.2.5 The impact of hidden terminals

It is well known that hidden terminals can have a considerable impact on the fairness of CSMA/CA based protocols, and this is widely acknowledged in the research community. In severe cases, e.g., when the hidden stations have a high traffic load, they can easily starve other stations in their vicinity. To show this, we add two hidden stations in both channels 0 and 2 in the topology shown Fig. 7, and plot the new topology in Fig. 12a where stations S1/S2 are the sources, and stations R1/R2 are the destinations of a TCP flow in channels 0/2 respectively. We first consider a long-lived FTP transfer as the application traffic from S1 to R1 and from S2 to R2. As can be seen in Fig. 12b when these two hidden flows are active between 100–200 s, all flows on the mesh backhaul are starved.




Fig. 12 The impact of hidden terminal traffic on TCP traffic in the topology shown in Fig. 7. Simulation parameters listed in Table 1. Note that in Fig. 12b and c, the throughput of flows 3–7 overlaps. **a** Topology. **b** FTP traffic. **c** Web traffic

However, in practice we expect that such severe situations occur relatively rarely since the bulk of network traffic is short-lived/bursty, e.g. web traffic, p2p traffic and voice, and so the associated hidden terminal interference is also short-lived/bursty. To explore this further, Fig. 12c shows the behaviour in the same setup as in Fig. 12b but now with the hidden terminals carrying on/off web traffic. During the ON time, a web user visits some web pages, whilst in the OFF time, the user is reading what he/she just downloaded. Following [25], the size of transfer during an ON period follows a Pareto distribution with mean 5KB and shape parameter 1.5, while the length of OFF periods follow an exponential distribution with mean duration of 1 s. It can be seen that in Fig. 12c when these two web flows are active from 100 s on, the impact of the hidden terminals is much reduced, to the extent that the impact on fairness is minor.

Our proposed fairness scheme is therefore reasonably insensitive to hidden terminal effects, and we argue that this is probably the best that can be hoped for. Severe hidden terminals effects, whereby for example affected stations become starved, are essentially a pathology of the CSMA/CA MAC and so we argue these should really be addressed using separate techniques. For example, an appealing approach for eliminating/mitigating hidden terminals is to use non-overlapping channels, which is highly feasible for 802.11a based networks since the number of orthogonal 802.11a channels exceeds 20. We note that 802.11a is commonly advocated for use on mesh backhaul links, with 802.11b/g used on access links. We also note that recently proposed algorithms such as that in [18] can reliably find non-interfering channel allocations in


a decentralised manner (without message passing or packet sniffing) and so are eminently suited to practical implementation.

6.3 Prioritising local traffic

In a large wireless multi-hop network, per-flow fairness can lead to local access traffic being starved of bandwidth at MPs close to the wired gateway. For example, consider the network topology in Fig. 2b with one local station at MP_8 (i.e. $N = 11$). Figure 13a plots the throughput of this local station as the number M of client stations at MP_0 is varied between 0 and 10. The per-flow scheme discussed above is used. The total number of destination stations is $2M + 2$ and, as expected, the bandwidth share of the local station is proportional to $1/(2M + 2)$ and so decreases towards zero as M increases.

Of course, this might be reasonable in some circumstances, e.g. where the network infrastructure is provided by a third party and resources are to be shared equally regardless of proximity to the gateway. However, it can also be undesirable, e.g., where the multi-hop network is formed by individual MP owners allowing shared access, subject to the proviso that each MP owner retains a minimum bandwidth share for his/her own traffic. In general, the choice of the most appropriate fairness requirement here is essentially a policy decision for the network operator. Our main point here is that TXOP does indeed provide the necessary flexibility to allow fairness to be controlled in a simple and practical manner. In particular, we show in this section that protecting local traffic can also be readily achieved using TXOP.

Fig. 13 Protecting local stations for topology in Fig. 2b. One local station at MP_8 ($N = 11$) and number of client stations at MP_0 is varied. Simulation parameters listed in Table 1. **a** Per-flow fairness. **b** Protecting local traffic ($\alpha = 0.25$)



We consider two approaches for protecting local traffic. The first approach only requires adjustment of TXOP but makes use of two 802.11e traffic classes. This approach has the virtue of simplicity and requires only straightforward adjustment of the MAC TXOP parameters. It does, however, require the use of two traffic classes. The 802.11e standard provides a total of only four traffic classes. Hence, in mixed traffic environments with voice, video, data etc we can quickly run out of traffic classes. We therefore also present an alternative approach to protecting local traffic that uses only a single traffic class. This benefit comes at the cost of the need for changes (a software change only) to the queuing discipline employed at the network interface queue.

6.3.1 Protecting local traffic using two traffic classes

As before, let n denote the number of neighbouring wireless stations (client stations and other MPs) on a given radio channel. At the i 'th wireless station we assign local flows to one traffic class which is assigned a TXOP value of $N_{i,local}$ packets. Relay flows are assigned to a second traffic class with TXOP $N_{i,relay}$.

Summing over all of the neighbouring stations, the share of transmission opportunities allocated to local traffic is therefore $\frac{\sum_{i=1}^n N_{i,local}}{\sum_{i=1}^n [N_{i,local} + N_{i,relay}]}$. We require to allocate a minimum proportion α of transmission opportunities to local traffic. That is, we require $\frac{\sum_{i=1}^n N_{i,local}}{\sum_{i=1}^n [N_{i,local} + N_{i,relay}]} = \alpha$. Rearranging yields

$$\sum_{i=1}^n N_{i,local} = \frac{\alpha}{1 - \alpha} \sum_{i=1}^n N_{i,relay}. \tag{13}$$

Any TXOP allocation that satisfies constraint (13) will ensure that local traffic receives an α share of the available transmission opportunities.

Note that freedom remains as to the selection of $N_{i,local}$ and $N_{i,relay}$, and this can be used to further control fairness. In particular, we consider selecting $N_{i,local} = k_l n_{local,i}$ and $N_{i,relay} = k_r n_{relay,i}$ where $n_{local,i}$ is the number of destination stations associated with the outgoing local traffic from node i and $n_{relay,i}$ is the corresponding relay traffic value. The scaling factors k_l




Fig. 14 Scenario 4: a subset of the MIT Roofnet topology

and k_r , are selected to satisfy constraint (13). By similar arguments to those in Section 5.2, this choice for $N_{i,local}$ ensures that the local traffic transmission opportunity share α is allocated among local traffic flows on a per station basis. Similarly, with the relay traffic share $1 - \alpha$.

We illustrate the impact of this strategy in Fig. 13b. In this example we select $\alpha = 0.25$, i.e., local traffic is allocated a minimum of 25% of the available transmission opportunities. It can be seen that the bandwidth share of the the local traffic is now lower bounded as required.

6.3.2 Protecting local traffic using one traffic class

We can also use a single traffic class for both local and relay traffic. Let the TXOP value for this class at the i 'th wireless station be T_i .

We now partition the total transmission time T_i into intervals $T_{i,local}$ and $T_{i,relay}$ with $T_{i,local} + T_{i,relay} = T_i$ and $T_{i,relay} = (1 - \alpha)/\alpha T_{i,local}$. Then $T_{i,local}/T_i = \alpha$ and $T_{i,relay}/T_i = (1 - \alpha)$. By partitioning TXOP in this way at every wireless station, we can protect the bandwidth share of local traffic as required. Specifically, on winning a transmission opportunity at station i , we use $T_{i,local}$ of the available transmit time to send local traffic and $T_{i,relay}$ to transmit relay traffic.

To implement this approach requires a software change that can be implemented in practice in a number of ways—for example, we can modify the wireless card device driver to perform a selective walk of the interface queue on each transmission opportunity. In practice, the interface queue is commonly divided into a device queue and a txqueue, with packets queued by the network stack in the txqueue before transferral to the device queue. While the device queue service disci-

Table 2 Flows in the roofnet topology in Fig. 14


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

pline may be hardware dependent, the txqueue service discipline is generally implemented within the operating system kernel and can be readily modified. The proposed approach can thus be readily implemented by use of a TXOP-sized device queue combined with a selective walk of the txqueue when transferring packets from the txqueue to the device queue.

6.4 The MIT roofnet topology

We further validate the TXOP scheme proposed in Eq. 1 in a subset of the MIT Roofnet topology (see Fig. 14). In this topology, there is an Internet gateway marked as GW in the figure. Orthogonal channels are assigned in neighbouring hops so that transmissions do not interfere. Client stations are marked by shadowed triangles, and MPs by circles. 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 2. Routing for each flow is via the GW and is statically assigned as indicated by the arrows in Fig. 14. For example, flow 0 traverses orthogonal channels 0, 1, 3, 4 and 2 from the client station 0 to the gateway GW.

Fig. 15 Throughput-fairness allocation for the topology in Fig. 15. Note that in Fig. 15a, the throughput of flows 13–19, and that of flows 0–12 overlaps, and in Fig. 15b, the throughput of all flows overlaps. **a** TCP ACK prioritisation. **b** TXOP






Fig. 16 Throughput-fairness allocation when multiple PHY rates are available for the scenario in Fig. 7. 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 1. Note that in Fig. 16a, the throughput of flow 0 overlaps with that of flows 3–7. In Fig. 16b, the channel usage time of flows 3–7 overlaps. **a** Throughput. **b** Channel usage time

In Fig. 15a we illustrate the resulting throughput when TCP ACKs are prioritised but the proposed TXOP scheme is not used. We see 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. 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. These MPs share channel 2 equally, which results in the observed unfairness

amongst flows. When the TXOP scheme proposed in Eq. 1 in Section 5.1 is now used, it can be seen that per-flow fairness now achieved as required, see Fig. 15b.

7 Time-based fairness

When multiple PHY rates are available [15], slow stations tend to dominate the channel access, resulting in fast and slow stations achieving a similar throughput. For example, we increase the link rate between MP_0




Fig. 17 Time-based fairness allocation when multiple PHY rates are available for the scenario in Fig. 7. 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 1. Note that in Fig. 17a, the throughput of flows 3–7 overlaps. In Fig. 17b, the time of flow 0 overlaps with that of flows 3–7. **a** Throughput. **b** Channel usage time




Fig. 18 Time-based fairness allocation when multiple PHY rates are available for the scenario in Fig. 14. 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 1. Note that in Fig. 18c, we also plot the channel usage time of each

flow over 100 s. However, as the actual channel usage time is around 4 s for each flow, we set the maximum of the y-axis to be 30 s to make the figure readable. **a** Throughput fairness. **b** Time-based: throughput. **c** Time-based: channel usage time

and MP_2 , and that between MP_4 and MP_5 to be 11 Mbps in the topology shown in Fig. 7, while keeping other link rates at 1 Mbps. If we use the throughput fairness scheme, the resulting flow throughput is shown in Fig. 16a where we can see that a similar allocation to Fig. 8c is achieved, i.e., it is still max-min fair in terms of flow throughput.

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. 16b. In this figure, we plot the time used by each flow over a 100 s interval. It can be seen that the time allocated to flow 0 is much less than that to flows 3–7 and the time allocation is not max-min fair.

In the literature, time-based fairness has been proposed to resolve this issue (see for example [31]). The rationale of time-based fairness is to allocate transmission time amongst contending stations which may have different PHY rates.

We can readily achieve time-based fairness using our proposed scheme in Eq. 2. Recall that we select

TXOP duration $K_i = n * T_{max}$ where T_{max} is the time for transmitting a packet with the slowest PHY rate. The slowest rate is known to all stations, and it is a fixed value for a specific standard, i.e., $R = 6$ Mbps for 802.11a and $R = 1$ Mbps for 802.11b/g. The time $T_{max} = T_f + T_{SIFS} + T_{ack}$ where T_f , T_{SIFS} and T_{ACK} are respectively the time to transmit a MAC layer data frame with the slowest rate, the duration of a SIFS and the time to transmit a MAC layer ACK frame. Note that MAC layer data frames are still transmitted with higher rates, but we update the TXOP duration based on the lowest rate to ensure that stations with the same number of flows are always able to hold the channel for a same period of air time if they win a transmission opportunity.

Using this scheme for the topology in Fig. 7, 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 $K_1 = 9230 * 2 \mu s$ and $K_3 = 9230 * 5 \mu s$. In Fig. 17a and b we illustrate the corresponding throughput and time results. As we can see, channel usage time is now max-min 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 Eq. 1, MP_0 and MP_4 use a duration $K = 1303 \mu s$ where $1303 \mu s$ is the time for transmitting a packet at 11 Mbps, but MP_1 and MP_3 use $K_1 = 9230 * 2 \mu s$ and $K_3 = 9230 * 5 \mu s$.

For the topology in Fig. 14, we illustrate the effectiveness of the time-based allocation (Eq. 2) by increasing all the link rates from station 11 to the gateway GW

Table 3 MAC and PHY parameters used in the test-bed implementation

T_{SIFS} (μs)	16
Idle slot duration (σ) (μs)	9
T_{DIFS} (μs)	34
CW_{min}	15
CW_{max}	1023
Retry limit	11
Packet size (bytes)	1500




Fig. 19 Throughput fairness: test-bed results for the topology in Fig. 7. Parameters used is listed in Table 3. Note that in Fig. 19c, the throughput of flow 0 overlaps with that of flows 3–7. **a** 802.11. **b** TCP ACK prioritisation. **c** Throughput fairness

to be 11 Mbps, while keeping other link rates at 1 Mbps. Other simulation parameters are listed in Table 1. The resulting throughput using the throughput-fair allocation (Eq. 1) and the time-based allocation (Eq. 2) is shown in Fig. 18a and c respectively. In Fig. 18c, we also plot the channel usage time of each flow over 100 s. It is evident that max–min fairness in terms of channel usage time is achieved. Note that as the actual channel usage time is too short (around 4 s for each flow), we set the maximum of the y-axis to be 30 s to make the figure readable.

8 Testbed implementations

We have implemented the topology shown in Fig. 7 using a test-bed constructed from Soekris net4801⁴ stations with Atheros 802.11a/b/g miniPCI cards. All stations run the Linux 2.6.21.1 kernel with a version of the MADWiFi⁵ wireless driver which is customised to allow the prioritisation described in this paper. In order to ensure a non-interfering channel allocation at each MP and to avoid interference with neighboring WLANs, all of these tests are performed with 802.11a channels. We use channels 40, 48 and 56 of 802.11a for channels 0, 1 and 2, respectively. The channel rate is fixed at 6 Mbps. To implement dual-radio MPs, we join two net4801 stations at 100 Mbps with a crossover cable to form a single logical MP. Routing in the network is statically configured. We use iperf⁶ to generate TCP traffic and data is collected from both iperf and tcpdump. All the control operations such as

initializing flows, collecting statistics etc., are carried out using the wired Ethernet of net4801 stations. SACK enabled TCP NewReno with a large receiver buffers (16 MBytes) is used. The TCP data packet size is 1500 bytes. Default TCP parameters of Linux Kernel 2.6.21.1 are used. To prioritise TCP ACK packets, we put ACK packets into the highest priority queue (Queue 3) which is assigned with $CW_{min} = 3$, $CW_{max} = 7$ and $AIFS = 2$. TCP data packets are collected into lower priority queue (Queue 2) which is assigned with $CW_{min} = 31$, $CW_{max} = 1023$ and $AIFS = 6$.

Using the proposed throughput fair approach in Eq. 1, the resulting allocation is shown in Fig. 19c where we use 0, 4200 and 11000 μs (which correspond to duration of transmitting 1, 2 and 5 packets at 6 Mbps) as TXOPs for MP_0 , MP_1 and MP_3 . It can be seen that an approximate max-min fair allocation is achieved. Here, $c_2 = 4.5$ Mbps and $c_0 = 4.75$ Mbps where c_i is the current channel capacity. In comparison, we also plot the results when 802.11 and TCP ACK prioritisation are used in Figs. 19a and 19b.

To validate the effectiveness of the time-based approach (Eq. 2) in test-bed experiments, we increase the PHY rate between MP_0 and MP_2 , and that between MP_4 and MP_5 to 12 Mbps while keep the PHY rates unchanged on other links. If we use the throughput-fair approach (Eq. 1), we can see in Fig. 20a that a similar max-min fair allocation to that in Fig. 19c is achieved. That is, while the PHY rates for flow 0 has been increased, its end-to-end performance has not improved.

Using the time-based approach, the corresponding channel usage time is shown in Fig. 20b. As can be seen, a max–min fair allocation in the time domain is now achieved. Note that here we plot the channel usage time of each flow over 10 s. With regard to throughput, it can be seen in Fig. 20c that the aggregate throughput is

⁴<http://www.soekris.com/net4801.htm>

⁵<http://sourceforge.net/projects/madwifi/>

⁶<http://dast.nlanr.net/Projects/Iperf/>




Fig. 20 Test-bed results when the PHY rate between MP_0 and MP_2 and that between MP_4 and MP_5 are increased to be 12 Mbps for the topology in Fig. 7. While PHY rates on all other links are still 6 Mbps. Other parameters used is listed in Table 3. Note that in Fig. 20a, the throughput of flow 0 overlaps with that

of flows 3–7. In Fig. 20c, the throughput of flow 0 overlaps with that of flows 1 and 2. In Fig. 20b, the channel usage time of flow 0 overlaps with that of flows 3–7. **a** Throughput fairness. **b** Time-based: channel usage time. **c** Time-based: throughput

improved to around 12 Mbps. In comparison, 10 Mbps is achieved when we use the throughput fairness approach.

9 Conclusions and future work

We have shown in this paper that using 802.11 at the MAC layer, gross unfairness can exist in multi-hop networks even if orthogonal channels are assigned in neighbouring hops. We have demonstrated that the TXOP mechanism of 802.11e can be used to ensure/restore fair allocation of resources to flows. The proposed TXOP based scheme is implementable on standard hardware in a simple and fully decentralised way without the need for message passing.

The network setups considered in the present paper are 802.11e multi-radio multi-hop networks, where the level of MAC layer contention is not excessive (10 stations sharing a link, but not 100 stations), and there is negligible channel noise and interference due to hidden/exposed terminals. When these factors are present however, tuning TXOP alone may not be sufficient. It is well known that adaptation of the contention window sizes can be used to mitigate excessive MAC layer contention. Further, if there are significant levels of packet loss due to channel noise and/or hidden/exposed terminals, then joint tuning of TXOP, contention window sizes and other parameters together may be necessary to ensure fairness.

In this paper, we restrict ourselves to provide fairness using only standard parameters so that the proposed algorithms can be implemented using off-the-shelf 802.11e hardware and, importantly, are compatible with legacy 802.11 equipment. We have designed

an algorithm to tune related parameters to achieve fairness in the more general cases mentioned above, and have obtained promising results in simulations. We are currently implementing the new algorithm in our experimental test-bed and will report on these results in the future.

References

1. IEEE 802.11 WG (1999) Part 11: wireless LAN medium access control (MAC) and physical layer (PHY) specifications. IEEE Std. 802.11
2. IEEE 802.11 WG (2004) 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
3. Bahl P, Chandra R, Dunagan J (2004) SSCH: slotted seeded channel hopping for capacity improvement in IEEE 802.11 adhoc wireless networks. In: Proc. ACM MOBICOM, Philadelphia, pp 216–230
4. Li B, Battiti R (2003) Supporting service differentiation with enhancements of the IEEE 802.11 MAC protocol: models and analysis. University of Trento, technical report, DIT-03-024
5. Bertsekas D, Gallager R (1987) Data networks. Prentice-Hall, Englewood Cliffs
6. Bianchi G (2000) Performance analysis of the IEEE 802.11 distributed coordination function. IEEE J Sel Areas Commun 18(3):607–614
7. Bottigliengo M, Casetti C, Chiasserini CF, Meo M (2004) Short-term fairness for TCP flows in 802.11b WLANs. In: Proc. of IEEE INFOCOM, pp 1383–1392
8. Celandroni N (2006) Comparison of FEC types with regard to the efficiency of TCP connections over AWGN satellite channels. IEEE Trans Wirel Commun 5(7):1735–1745
9. Celandroni N, Davoli F, Ferro E, Gotta A (2006) Long-lived TCP connections via satellite: cross-layer bandwidth allocation, pricing, and adaptive control. IEEE/ACM Trans Netw 14(5):1019–1030

10. De Couto D, Aguayo D, Bicket J, Morris R (2003) A high-throughput path metric for multi-hop wireless routing. In: Proc. of ACM MobiCom, pp 134–146
11. Draves R, Padhye J, Zill B (2004) Comparison of routing metrics for static multi-hop wireless networks. In: Proc. of ACM SIGCOMM
12. Duffy K, Leith D, Li T, Malone D (2006) Modeling 802.11 Mesh Networks. *IEEE Commun Lett* 10(8):635–637
13. Gambiroza V, Sadeghi B, Knightly EW (2004) End to end performance and fairness in multihop wireless backhaul networks. In: Proc. of ACM MOBICOM
14. Garetto M, Salonidis T, Knightly EW (2006) Modeling per-flow throughput and capturing starvation in csma multi-hop wireless networks. In: Proc. of IEEE INFOCOM
15. Heusse M, Rousseau F, Berger-Sabbatel G, Duda A (2003) Performance anomaly of 802.11b. In: Proc. IEEE INFOCOM, San Francisco, pp 836–843
16. Le Boudec J-Y (2008) Rate adaptation, congestion control and fairness: a tutorial. http://ica1www.epfl.ch/PS_files/LEB3132.pdf
17. Leith D, Clifford P, Malone D, Ng A (2005) TCP fairness in 802.11e WLANs. *IEEE Commun Lett* 9(11):964–966
18. Leith D, Clifford P (2006) A self-managed distributed channel selection algorithm for WLANs. In: ACM/IEEE RAWNET
19. Li T, Leith D, Malone D, Badarla V (2008) Achieving end-to-end fairness in 802.11e based wireless multi-hop mesh networks. In: Proc. Chinacom, Hangzhou China
20. Badarla V, Malone D, Leith D (2008) Implementing TCP flow-level fairness using 802.11e in a multi-radio mesh testbed. *IEEE Commun Lett* 12(4):262–264
21. Maheshwari R, Gupta H, Das SR (2006) Multichannel MAC protocols for wireless networks. In: Proc. IEEE SECON. Reston, VA, vol 2, pp 393–401
22. Malone D, Duffy K, Leith D (2007) Modeling the 802.11 distributed coordination function in nonsaturated heterogeneous conditions. *IEEE/ACM Trans Netw* 15(1):159–172
23. Massoulié L, Roberts J (2002) Bandwidth sharing: objectives and algorithms. *IEEE/ACM Trans Netw* 10(3):320–328
24. Mo J, Walrand J (2000) Fair end-to-end window-based congestion control. *IEEE/ACM Trans Netw* 8(5):556–567
25. Prasad RS, Dovrolis C, Thottan M (2007) Router buffer sizing revisited: the role of the output/input capacity ratio. CoNEXT
26. Ramachandran K, Belding-Royer E, Almeroth K, Buddhikot M (2006) Interference-aware channel assignment in multi-radio wireless mesh networks. In: Proc. of IEEE INFOCOM
27. Raman B (2006) Channel allocation in 802.11-based mesh networks. In: Proc. of IEEE INFOCOM
28. Raniwala A, De P, Sharma S, Krishnan R, Chiueh T (2007) End-to-end flow fairness over IEEE 802.11-based wireless mesh networks. In: Proc. of IEEE INFOCOM, Mini-Symposium
29. Stanojevic R, Shorten R (2007) Beyond CHOKe: stateless fair queueing. In: Proc. of EuroFGI NET-COOP
30. Subramanian VG, Duffy KR, Leith DJ (2009) Existence and uniqueness of fair rate allocations in lossy wireless networks. *IEEE Trans Wirel Commun* 8:3401–3406
31. Tan G, Gutttag J (2004) Time based fairness improves performance in multi-rate WLANs. In: Proc. USENIX, Boston
32. Tang D, Baker M (2000) Analysis of A Local-Area Wireless Network. In: Proc. of ACM MobiCom
33. Tinnirello I, Choi S (2005) Temporal fairness provisioning in multi-rate contention-based 802.11e WLANs. In: Proc. of IEEE WOWMOM
34. Wu H, Yang F, Tan K, Chen J, Zhang Q, Zhang Z (2006) Distributed channel assignment and routing in multi-radio multi-channel multi-hop wireless networks. *J Sel Areas Commun* 24:1972–1983
35. Yang Y, Wang J, Kravets R (2005) Distributed optimal contention window control for elastic traffic in wireless LANs. In: Proc. of IEEE INFOCOM, pp 35–46
36. Zhao Z, Darbha S, Reddy ALN (2004) A method for estimating the proportion of nonresponsive traffic at a router. *IEEE/ACM Trans Netw* 12(4):708–718