

Max-min Fairness in 802.11 Mesh Networks

Douglas J. Leith, Qizhi Cao, Vijay G. Subramanian
Hamilton Institute, NUI Maynooth

Abstract—In this paper we build upon the recent observation that the 802.11 rate region is log-convex and, for the first time, characterise max-min fair rate allocations for a large class of 802.11 wireless mesh networks.

I. INTRODUCTION

In this paper we build upon the recent proof in [15] that the 802.11 rate region is log-convex and, for the first time, characterise max-min fair rate allocations for a large class of 802.11 mesh networks. By exploiting features of the 802.11e/n MAC, in particular TXOP packet bursting, we are able to use this characterisation to establish a straightforward, practically implementable approach for achieving max-min throughput fairness. We demonstrate that this approach can be readily extended to encompass time-based fairness in multi-rate 802.11 mesh networks.

Fairness in 802.11 networks has been the subject of a considerable body of literature. A large part of this literature is concerned with *unfairness* behaviour in 802.11 networks due to hidden terminals, exposed terminals, capture, upload/download unfairness *etc.*, see for example [1], [17], [5], [10], [13] and references therein. Proportional fairness over a single 802.11 hop is considered by [24], but this work makes the simplifying assumption that every wireless station in a WLAN is always saturated, which cannot be expected to hold in general and is an unreasonable hypothesis for multi-hop networks. An extensive literature relates to utility fairness in wired networks, but the CSMA/CA scheduling used in 802.11 differs fundamentally from wired networks due to carrier sense deferral of the contention window countdown and the occurrence of colliding transmissions – both of which act to couple together the scheduling of transmissions by stations in a WLAN and lead to the rate region being non-convex. Utility fairness has been considered in random access wireless networks, but this work has been confined to the Aloha MAC, see [9], [6], [11], [27] and references therein. The Aloha framework assumes that idle and transmission slots are of the same duration and so does not encompass standard 802.11 frame structure where (i) it is common for transmissions to be more than an order of magnitude longer than the idle slot duration in order to improve throughput efficiency and (ii) the mean transmission duration is not identical at all stations but instead depends on the packet size and PHY rate selected. While it has been known for some time that Aloha networks have a log-convex rate region [9], [27], it has only recently been established that the 802.11 rate region is also log-convex [15]; it is this fundamental result that underpins the max-min utility fair analysis in the present paper.

This material is based upon works supported by the Science Foundation Ireland under Grant No. 07/IN.1/901.

II. NETWORK MODEL

A. Network Architecture

We consider a mesh network formed from a set of inter-connected WLANs and assume that the WLANs are non-interfering *i.e.* that they either transmit on orthogonal channels or are physically separated so that transmissions on the same channel do not interfere. Traffic is routed between WLANs via mesh points equipped with multiple radios. Communication between mesh points is peer-to-peer so that sending a packet from WLAN i to WLAN j involves a single transmission (rather than routing via a central access point). We assume that stations within a WLAN are within sensing distance of one another *i.e.* there are no hidden terminals; we comment later on incorporating hidden terminals. Such a mesh network is illustrated, for example, in Figure 1. In this example the network is formed from five inter-connected WLANs such that three orthogonal channels are sufficient to achieve a non-interfering allocation.

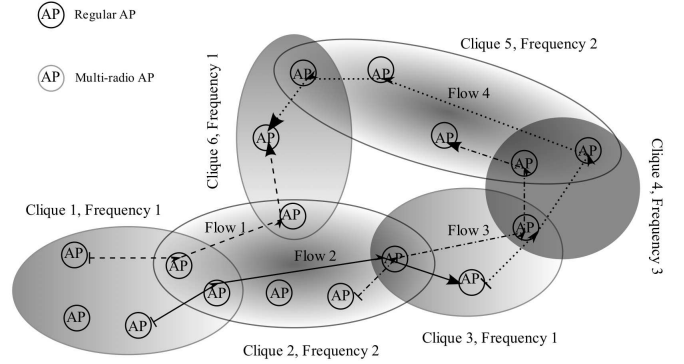


Fig. 1. Illustrating class of mesh networks considered.

B. Station throughput

Consider one of the WLANs in the mesh network and let n denote the number of stations in the WLAN. Following [18], we divide time into MAC slots where each MAC slot may consist either of a PHY idle slot, a successful transmission or a colliding transmission (where more than one station attempts to transmit simultaneously). Let τ_i denote the probability that station i attempts a transmission in a slot. The mean throughput of station i is then (*e.g.* see [18])

$$s_i(\mathcal{T}) = \frac{\tau_i \prod_{k=1, k \neq i}^n (1 - \tau_k) D_i}{\sigma P_{idle} + T_s P_{succ} + T_c (1 - P_{idle} - P_{succ})}$$

where $P_{idle} = \prod_{k=1}^n (1 - \tau_k)$ is the probability that a slot is a PHY idle slot, $P_{succ} = \sum_{i \in N} \tau_i \prod_{k=1, k \neq i}^n (1 - \tau_k)$ is the probability that a slot is a successful transmission, $\mathcal{T} = [\tau_1 \dots \tau_n]^T$ is the vector of attempt probabilities, D_i is the mean number of bits sent by station i in a successful transmission, σ is the PHY idle slot duration, T_s is the mean duration of a successful transmission (including time to transmit each data frame, receive the MAC ACK and wait for DIFS) and T_c the mean duration of a collision.

C. Incorporating TXOP

Later, we will make use of the TXOP packet bursting in 802.11e/n to facilitate achieving max-min fairness. With TXOP, the length of time during which a station can keep transmitting without releasing the channel once it wins a transmission opportunity is specified as a control parameter. In order not to release the channel, a SIFS interval is inserted between each packet-ACK pair and a successful transmission round then consists of multiple packets and ACKs. By adjusting the TXOP time the number of packets that may be transmitted by a station at each transmission opportunity can be controlled. We can readily generalise the above throughput expression to support TXOP packet bursting as follows. Firstly observe that when TXOP packet bursting is used colliding transmissions end after sending the first packet in a burst and so T_c is unchanged. However, the duration T_s of a successful transmission now depends on the size of the TXOP packet burst. To encompass situations where stations may transmit different sized bursts on winning a transmission opportunity we let $T_{s,i}$ denote the mean duration of a successful transmission by station i . The throughput of station i is then

$$s_i(\mathcal{T}) = \frac{\frac{\tau_i}{1-\tau_i} P_{idle} D_i}{\sigma P_{idle} + \sum_{i=1}^n T_{s,i} \frac{\tau_i}{1-\tau_i} P_{idle} + T_c (1 - P_{idle} - P_{succ})}$$

It will prove useful to work in terms of the quantity $x_i = \tau_i / (1 - \tau_i)$ rather than τ_i . With this transformation we have that $P_{idle} = 1 / \prod_{k=1}^n (1 + x_k)$ and $P_{succ} = \sum_{i=1}^n x_i / \prod_{k=1}^n (1 + x_k)$ and so

$$s_i(x, N) = \frac{N_i x_i}{X(x, N)} \frac{L_i}{T_c} \quad (1)$$

where $N_i = T_{s,i} / T_c$, $L_i = D_i / N_i$ and

$$X(x, N) = a + \sum_{k=1}^n (N_k - 1) x_k + \prod_{k=1}^n (1 + x_k) - 1 \quad (2)$$

with $a = \sigma / T_c$, $x = [x_1, \dots, x_n]^T$ and $N = [N_1, \dots, N_n]^T$. We also have that the mean fraction of time spent by station i on successful transmissions is

$$t_i(x, N) = \frac{N_i x_i}{X(x, N)} \quad (3)$$

which is simply a rescaling of the station throughput expression (1).

In the foregoing we have implicitly assumed that packet losses only occur due to colliding transmissions *i.e.*

Assumption 1: Packet losses from sources other than collisions can be neglected.

We discuss relaxing this assumption and including channel noise losses in Section VIII below. In addition, we will generally make the following assumption,

Assumption 2: Frame transmissions are of duration T_c .

A TXOP burst therefore consists of a sequence of frame transmissions each of duration T_c . This assumption yields the useful technical benefit that the collision duration T_c is invariant with the attempt rates x_i used in a WLAN – if stations used frames of different duration then the duration of a collision would depend on the specific set of stations involved in a collision and so on the attempt rates x_i . More importantly, however, it is also a natural assumption in the context of 802.11e where TXOP bursts are specified in terms of their duration in seconds (which, in turn, is motivated by consideration of time-based fairness when stations use different PHY rates). With this assumption, N_i can be interpreted as the mean number of transmissions in a burst and L_i as the mean size, in bits, of the payload of each frame.

D. Constraining burst size

Before proceeding, it is important to note that it is necessary to suitably constrain the size N_i of allowed TXOP packet bursts. To see this, say we let $N_i = \lambda n_i$ with $\lambda > 0$, $n_i > 0$ and look at the behaviour as $\lambda \rightarrow \infty$. It can be verified that $ds_i/d\lambda$ equals

$$\frac{n_i x_i}{X} \left(1 - \frac{\lambda \sum_{j=1}^n n_j x_j}{\sum_{j=1}^n (\lambda n_j - 1) x_j + a + \prod_{j=1}^n (1 + x_j) - 1} \right) \frac{L_i}{T_c}$$

which can be seen to be strictly positive. That is, increasing λ (and so burst size) *always* increases throughput. In the limit, $s_i \rightarrow \frac{n_i x_i}{\sum_{j=1}^n n_j x_j} \frac{L_i}{T_c}$ as $\lambda \rightarrow \infty$. Observe that the idle time and collision time terms (which remain of finite duration) are washed out in the denominator and so the efficiency of the network is maximised subject to the fixed per packet overhead embodied by L_i / T_c . In effect, this says that any point strictly in the interior of the simplex $\{s : \sum_{i=1}^n \frac{s_i T_c}{L_i} \leq 1\}$ is achievable by an appropriate choice of N_i s. This high efficiency comes at the price of unbounded delays and so is not of practical interest. Instead, to maintain bounded delay it is necessary to constrain the burst size and we let \bar{N}_i denote the maximum burst size admissible at station i .

E. Finite-load

It is useful to distinguish between the attempt probability τ_i and the attempt probability design parameter $\bar{\tau}_i$. $\bar{\tau}_i$ is the probability that station i considers making a transmission in a slot, but a transmission will not actually take place unless at least one packet is available to send. It is the attempt probability τ_i which is relevant for the foregoing throughput expressions.

When $\tau_i = \bar{\tau}_i$ a station is said to be *saturated* and sends a packet at every transmission opportunity, otherwise it is *unsaturated*. For unsaturated stations the attempt probability τ_i depends jointly on the offered load and $\bar{\tau}_i$. We will assume that when a station is unsaturated the throughput is equal to the offered load *i.e.* that stations have sufficient buffering

that queue overflow losses can be neglected when a station is unsaturated¹. We also assume that the corresponding attempt probability τ_i is the just value that makes throughput expression (1) equal the offered load, *i.e.*

Assumption 3: Let \mathcal{U} denote the set of unsaturated stations in a WLAN and \mathcal{S} the set of saturated stations, with $\mathcal{U} \cup \mathcal{S} = \{1, 2, \dots, n\}$ and $\mathcal{U} \cap \mathcal{S} = \emptyset$. Let $y_i = \bar{\tau}_i / (1 - \bar{\tau}_i)$. The attempt rate at a saturated station $j \in \mathcal{S}$ is then y_j and the set of admissible station attempt rates is $\mathcal{X} = [\underline{x}_1, y_1] \times \dots \times [\underline{x}_n, y_n]$ with $\underline{x}_i = y_i$ for $i \in \mathcal{S}$ and 0 otherwise. Let s_i denote the offered load at unsaturated station $i \in \mathcal{U}$. If a solution $(x, N) \in \mathcal{X} \times \mathcal{N}$, where $\mathcal{N} = [1, \bar{N}_1] \times \dots \times [1, \bar{N}_n]$, exists to the throughput balance equations

$$s_i = \frac{N_i x_i}{X(x, N)} \frac{L_i}{T_c} \quad \forall i \in \mathcal{U} \quad (4)$$

then the offered load of s_i can be serviced by unsaturated station $i \in \mathcal{U}$ with the attempt rate x_i and burst-size N_i solving the balance equations.

Note that for solutions to (4) to exist it is necessary and sufficient that the set

$$C' = \left\{ (x, N) \in \mathcal{X} \times \mathcal{N} : \forall i \in \mathcal{U} \quad s_i \leq \frac{N_i x_i}{X(x, N)} \frac{L_i}{T_c} \right\}$$

be non-empty²

F. Realisation in 802.11e/n

Following the approach taken in Bianchi-like throughput models (*e.g.* see [18] and references therein), transmissions by an 802.11 station can be modelled as a renewal process, with renewals occurring after a successful transmission or discard. The attempt probability can then be directly related to the 802.11 MAC parameters CW_{min} , CW_{max} , *etc.* For simplicity, we will hereafter assume that the attempt probability design parameter $\bar{\tau}_i$ can be freely selected. However, this is not a fundamental requirement of our analysis and can be readily relaxed provided any constraints imposed on $\bar{\tau}_i$ continue to yield a log-convex rate region; in particular, Theorem 3 below carries over in the obvious way. As an example of admissible constraints on $\bar{\tau}_i$, consider an 802.11 WLAN where we select $CW_{min} = CW_{max} = CW$, where CW is an appropriate constant *e.g.* 32. Then $\bar{\tau}_i$ is constrained to take the single value $2/(CW - 1)$ and the attempt probability τ_i can take values in $[0, 2/(CW - 1)]$ as the offered load on stations is varied. By Theorem 1 in Section III, the corresponding WLAN rate region is log-convex. Indeed, we can constrain $\bar{\tau}_i$ to take any finite set of values (*e.g.* corresponding to CW taking powers of 2) since the resulting rate region is the intersection of the log-convex rate regions corresponding to each of the individual constraints on $\bar{\tau}_i$ and is therefore log-convex.

¹Conservation of packets then means that the mean throughput must equal the mean arrival rate. Observe also that, by Loynes theorem [16], for sufficiently large buffering we have the intuitive property that a station will be unsaturated whenever the mean packet inter-arrival time is less than the mean service time.

²This is because solutions to (4) are a superset of solutions to the optimization problem $\min_{(x, N) \in C'} \prod_{i=1}^n x_i N_i$. From the proofs of Theorems 1 and 3 (see later) this optimisation can be transformed into a convex problem such that all the constraints are satisfied with equality at the optimal solution(s), if the problem is feasible.

G. Additional notation

We represent the connectivity between WLANs via graph \mathcal{G} with vertices \mathcal{V} and edges \mathcal{E} . Each vertex in \mathcal{V} corresponds to a WLAN and an edge exists between WLANs that can communicate. Edges are labelled by the radio channel used. Let $\mathcal{N}_i(c)$ denote the set of neighbours of WLAN i on channel c *i.e.* a set of peering mesh points. We will assume that each such set uses a channel with a unique label, but this is just a notational assumption and does not require that the physical channels are all different (in practice physical channels would be reused to exploit spatial multiplexing). Since there are no hidden terminals, peering mesh points form a clique *i.e.* $\mathcal{N}_j(c) \cup \{j\} = \mathcal{N}_i(c) \cup \{i\} \quad \forall j \in \mathcal{N}_i(c)$ and we let $n(c) = |\mathcal{N}_j(c) \cup \{j\}|$ denote the number of peers on channel c . Let \mathcal{P} denote the set of network flows. Associated with each flow p is a source client station and route $r(p)$ (assumed loop-free) consisting of edges in \mathcal{G} (*i.e.* triples (i, j, c) , $j \in \mathcal{N}_i(c)$) traversed by the flow. For notational simplicity we assume that flows do not start/finish at mesh points. Let $\mathcal{P}_{i,j}(c)$ denote the set of flows $\{p : (i, j, c) \in r(p), p \in \mathcal{P}\}$ relayed from WLAN i to WLAN j on channel c , $\mathcal{P}_i(c) = \cup_{j \in \mathcal{N}_i(c)} \mathcal{P}_{i,j}(c)$ denote the set of all flows relayed by WLAN i and $\mathcal{P}(c) = \cup_{i \in \mathcal{V}} \mathcal{P}_i(c)$ denote the set of all flows relayed by peers on channel c .

III. LOG-CONVEXITY OF RATE REGION

We begin by extending the log-convexity analysis in [15] to include TXOP packet bursting, and then use this to establish log-convexity of the mesh network rate region. We present a new method of proof that makes use of theory of posynomials and geometric programming [2], [3].

Definition 1: WLAN Rate Region. The rate region of a WLAN is the set \mathcal{R} of achievable throughput vectors $S(x, N) = [s_1 \dots s_n]^T$, with i 'th element given by (1), as the vector x ranges over $\mathcal{X} = [0, \bar{x}_1] \times \dots \times [0, \bar{x}_n]$ and the vector N ranges over $\mathcal{N} = [1, \bar{N}_1] \times \dots \times [1, \bar{N}_n]$.

Definition 2: Log-convexity. A set $C \in \mathbb{R}^n$ is convex if for any $s^1, s^2 \in C$ and $0 \leq \alpha \leq 1$, there exists an $s^* \in C$ such that $s^* = \alpha s^1 + (1 - \alpha) s^2$. A set C is log-convex if the set $\log C := \{\log s : s \in C\}$ is convex.

Theorem 1: The WLAN rate-region \mathcal{R} is log-convex.

Proof: The throughput of station i is given by

$$s_i = \frac{N_i x_i}{X(x, N)}$$

where

$$\begin{aligned} X(x, N) &= a + \sum_{j \in \mathcal{M}} N_j x_j + \prod_{j \in \mathcal{M}} (1 + x_j) - 1 - \sum_{j \in \mathcal{M}} x_j \\ &= a + \sum_{j \in \mathcal{M}} N_j x_j + \sum_{k=2}^n \sum_{A \subseteq \mathcal{M}, |A|=k} \prod_{j \in A} x_j. \end{aligned}$$

and $\mathcal{M} = \{1, 2, \dots, n\}$ denotes the set of stations in the WLAN. Making a change of variables to $y_j = \log(x_j)$ and

$\eta_j = \log(N_j)$, we have

$$\log(s_i) = y_i + \eta_i - \log \left(a + \sum_{j \in \mathcal{M}} \exp(\eta_j + y_j) + \sum_{k=2}^n \sum_{A \subseteq \mathcal{M}, |A|=k} \exp \left(\sum_{j \in A} y_j \right) \right).$$

with $y_j \in Y_j := (-\infty, \log(\bar{x}_j)]$, $\eta_j \in E_j := (0, \log(\bar{N}_j)]$. Note that the right-hand-side is a concave function of (y, η) since it is the transformed version of the reciprocal of a posynomial [2]. Then the definition of a concave function implies that

$$C_i = \left\{ (\mu, y, \eta) \in \mathbb{R}^n \times \prod_{j \in \mathcal{M}} Y_j \times \prod_{j \in \mathcal{M}} E_j : \mu_i \leq y_i + \eta_i - \log \left(a + \sum_{j \in \mathcal{M}} \exp(\eta_j + y_j) + \sum_{k=2}^n \sum_{A \subseteq \mathcal{M}, |A|=k} \exp \left(\sum_{j \in A} y_j \right) \right) \right\}$$

is a convex set. Therefore, $C = \cap_{i \in \mathcal{M}} C_i$ is also a convex set. The log rate-region is then the image of C under the (linear) projection map that takes (μ, y, η) to μ . Thus, the log rate-region is convex. ■

We also have the following corollary that will prove useful later. Let $\mathcal{R}'(\bar{p})$ denote the set of achievable throughput vectors $S(x, N) = [s_1 \dots s_n]^T$ as the vector x ranges over $\mathcal{X} \cap \{x : \prod_{i=1}^n (1 + x_i) \leq \bar{p}\}$ and the vector N ranges over \mathcal{N} .

Corollary 1: The constrained WLAN rate region $\mathcal{R}'(\bar{p})$ is log-convex for any $\bar{p} \geq 1$.

Proof: We require $\bar{p} \geq 1$ for \mathcal{R}' to be non-empty. Now using the same transformation as in the proof of Theorem 1, the constraint that $\prod_{i=1}^n (1 + x_i) \leq \bar{p}$ translates to restricting attention to the following set

$$D = \left\{ (\mu, y, \eta) \in \mathbb{R}^n \times \prod_{j \in \mathcal{M}} Y_j \times \prod_{j \in \mathcal{M}} E_j : \sum_{j \in \mathcal{M}} \log(1 + \exp(y_j)) \leq \log(\bar{p}) \right\}$$

which is a convex set as a consequence of $\log(1 + \exp(\cdot))$ being a convex function. The log rate-region is then $C \cap D$ which is convex, thus establishing the corollary. ■

We note that the proof above can be readily extended to show that other constraints on τ (or x) and N vectors also yield a convex set under our chosen transformation³. Since the station transmission time (3) is simply a rescaling of the station throughput expression (1) we also have the following result.

³For example, consider a constraint of the form $\sum_{i=1}^n x_i^2 \leq 1$. Since the left-hand-side becomes $\log(\sum_{i=1}^n \exp(2y_i)) \leq 0$, log-convexity continues to hold. Similarly, the constraint $\sum_{i=1}^n \tau_i^2 \leq 1$ can be transformed to $\frac{\tau_i}{1-\tau_i} \leq x_i$ for all i and $\sum_{i=1}^n \tau_i^2 \leq 1$ with x_i replacing $\frac{\tau_i}{1-\tau_i}$ in all the throughput formulae. Since the first set of constraints can be transformed to $\tau_i + \frac{\tau_i}{x_i} \leq 1$ for all i , the constraints are now posynomial constraints in (τ, x) and log-convexity continues to hold.

Corollary 2: The sets of feasible transmission times corresponding to rate regions \mathcal{R} and \mathcal{R}' are also log-convex.

A mesh network carries flows which traverse the component WLANs. Let $\mathcal{R}(\mathcal{G})$ denote the network rate region *i.e.* the set of feasible flow throughputs. Since the throughput of unsaturated stations equals their offered load (see Assumption 3 and related discussion regarding buffering requirements), the network rate region is obtained by the appropriate intersection of the individual WLAN rate regions. It follows immediately from the log-convexity of the component WLAN rate regions that the mesh network rate region is log-convex, *i.e.* we have the following corollary.

Corollary 3: The mesh network rate-region $\mathcal{R}(\mathcal{G})$ is log-convex.

Proof: We will once again use the property that convexity is preserved when taking intersections. Thus, it suffices to outline the key steps of the proof. Consider mesh point i on channel c . Let \mathcal{P}_i be the set of flows relayed by this mesh point. Using the transformation from the proof of Theorem 1 the additional constraint that we need to satisfy is the flow-balance constraint at every mesh point *i.e.*

$$\log \left(\sum_{p \in \mathcal{P}_i} \exp(\tilde{s}(p)) \right) \leq y_i + \eta_i - \log \left(a + \sum_{j \in \mathcal{N}_i(c)} \exp(\eta_j + y_j) + \sum_{k=2}^{n_l(c)} \sum_{A \subseteq \mathcal{N}_i(c), |A|=k} \exp \left(\sum_{j \in A} y_j \right) \right)$$

where $\tilde{s}(p) = \log(s(p))$, which is again a convex constraint. Such constraints have to be satisfied for all the mesh points, and hence we get log-convexity for the entire rate-region. ■

IV. MAX-MIN THROUGHPUT FAIRNESS

In this section we establish our main result, characterising max-min fair throughput allocations in the class of 802.11 mesh networks considered.

A. Assumptions

Before proceeding we make the following assumptions. We will relax all of these assumptions later, but they are useful for gaining initial insight into the nature of the max-min fair throughput allocation.

Assumption 4: PHY rate. All stations in the WLAN on channel c use the same PHY rate for transmissions.

It follows from Assumption 2 that stations use the same frame size *i.e.* $L_i = L$.

Assumption 5: Maximum burst-size. A station can transmit a maximum of one frame per flow at each successful transmission. It follows that $N_i \leq \bar{N}_i = |\mathcal{P}_i|$, where $|\mathcal{P}_i|$ is the number of flows carried by station i , and we have an additional constraint for each flow, namely $s(p) \leq \frac{x_i}{X(x, N)} \frac{L_i}{T_c}$ for flows $p \in \mathcal{P}_i$ carried by station i .

Note that the additional constraint introduced here can once again be transformed to a log-convex constraint and therefore

Corollary 3 still holds and the network rate-region is still log-convex.

Assumption 6: Attempt probability. All stations in the WLAN on channel c use the same attempt probability design parameter $y(c) = \bar{\tau}(c)/(1 - \bar{\tau}(c))$.

Recall that $\bar{\tau} \in [0, 1)$ is the transmission attempt probability when a station is saturated (always has a packet to send), but the actual attempt probability will be lower when a station is unsaturated. Note that $y(c)$ need not be the same for every WLAN, but stations within a WLAN are assumed to use the same value of attempt parameter.

The channel idle probability $P_{idle}(c)$ in the WLAN on channel c is $1/\prod_k(1 + x_k)$.

Assumption 7: Idle probability. $\prod_k(1 + x_k) \leq \bar{p}(c)$.

This assumption involves no loss of generality as by selecting $\bar{p}(c)$ sufficiently large we can always ensure that the constraint is inactive. Nevertheless, including this assumption allows us to also consider smaller values of $\bar{p}(c)$ as we will see later. By Corollary 1, the rate region is log-convex for any value of $\bar{p} \geq 1$.

B. Water-filling & Bottleneck links

Assumptions 4-7 do not change the log-convexity of the network rate region and so we immediately have that a unique max-min rate allocation exists. The network rate region also has the free disposal property [22] (same as coordinate-convexity) since each co-ordinate of the throughput vector is lower bounded by 0 and any non-zero feasible vector can always be decreased – by scaling the attempt rate vector x – while staying within the rate region. By [22, Theorem 3] the max-min solution can therefore be found by water-filling.

Recall the water-filling algorithm in [22]:

- 1: Let $\mathcal{P}^0 = \mathcal{P}$, $\mathcal{R}^0 = \mathcal{R}(\mathcal{G})$, $n = 0$
- 2: **do**
- 3: Find $\max T^n$ s.t. $S_p = T^n \forall p \in \mathcal{P}^n$, $S \in \mathcal{R}^n$
- 4: $\mathcal{R}^{n+1} = \{S \in \mathcal{R}^n : S_p \geq T^n \forall p \in \mathcal{P}^n\}$
- 5: $\mathcal{P}^{n+1} = \{p \in \mathcal{P}^n : \forall S \in \mathcal{R}^{n+1}, S_p > T^n\}$
- 6: $n = n + 1$
- 7: **until** $\mathcal{P}^n = \emptyset$

where \mathcal{P} is the set of network flows, $\mathcal{R}(\mathcal{G})$ denotes the network rate region (*i.e.* the set of feasible flow throughputs), S denotes the vector of flow throughputs and S_p is the throughput of flow p (element p of vector S). On termination of this water-filling algorithm, the remaining point in \mathcal{R}^n is the max-min fair allocation of flow throughputs.

Step 3 is the key step in the algorithm. It finds the maximum throughput T^n that the flows in set \mathcal{P}^n may collectively use while remaining within the network rate region. The flows whose throughput cannot be increased above T^n are then removed from set \mathcal{P}^n , and step 3 repeated. We can express step 3 more explicitly in our wireless mesh network context

as:

$$\max_{x, T^n} T^n \quad (5)$$

$$\text{s.t. } s(p) = T^n \forall p \in \mathcal{P}^n \quad (6)$$

$$s(p) \leq \frac{x_k(c)}{X(c)} \frac{L}{T_c(c)}, \forall p \in \mathcal{P}, (k, \bullet, c) \in r(p) \quad (7)$$

$$\sum_{p: (k, \bullet, c) \in r(p)} s(p) = \frac{N_k x_k(c)}{X(c)} \frac{L}{T_c(c)} \quad (8)$$

$$x \geq 0 \quad (9)$$

$$\prod_k (1 + x_k) \leq \bar{p}(c) \quad (10)$$

Constraints (7)-(10) ensure that the vector of flow rates lies within the network rate region.

For all flows there exists an iteration n such that the flow is eventually removed from set \mathcal{P}^n because its throughput cannot be increased above T^n . When a flow is removed the constraint (7) is necessarily tight (*i.e.* it cannot be loosened by any choice of x while respecting the other constraints) for some WLAN c . We say that flow p is *bottlenecked* at this WLAN. Our interest in bottlenecks stems from the following property, which follows immediately from these observations,

Theorem 2: A throughput allocation is max-min fair if and only if every flow has a bottleneck.

Observe also that all of the flows bottlenecked at the same WLAN c have the same throughput (owing to constraint (6)), and this is strictly greater than the throughput of the other flows which traverse this WLAN but are not bottlenecked there. We have therefore established that the well-known bottleneck property of max-min throughput allocations in wired networks also carries over to 802.11 mesh networks.

C. Main result

Surprisingly, despite the complex nature of the mesh network rate region (where flow rates are strongly coupled at each WLAN), we can obtain an almost complete characterisation of the max-min allocation of station attempt probabilities and burst sizes within each WLAN. This makes use of the characterisation of the max-min allocation in terms of waterfilling and bottlenecks.

Recall that we say that a flow is *saturated* if it has a packet available to send at every transmission attempt by the station, and is otherwise *unsaturated*.

Theorem 3: Under Assumptions 1-7, the max-min fair throughput allocation within each WLAN possesses the following properties:

- 1) The attempt rate design parameter $y(c) \geq \bar{x}(c)$ in each WLAN where $\bar{x}(c)$ is the attempt rate that maximises the throughput of saturated flows.
- 2) Flows bottlenecked at the WLAN send one frame at every successful transmission made by the station. When $y(c) = \bar{x}(c)$, all bottlenecked flows are saturated. When $y(c) > \bar{x}(c)$ they are unsaturated.
- 3) Non-bottlenecked flows are always unsaturated.

Proof: See Appendix A. ■

The importance of Theorem 3 is that it goes a long way to telling us how we might realise a max-min fair allocation in wireless mesh networks. Specifically, consider a mesh network where each WLAN is configured as follows:

- 1) Stations in a WLAN all use the same attempt rate parameter $y(c)$ (e.g. in 802.11 terminology, all stations in a WLAN use the same value of $CW_{min} = CW_{max}$).
- 2) Stations use per flow queueing and at each transmission opportunity send one frame of data from each non-empty queue.
- 3) Parameter $y(c)$ is selected to maximise the throughput of saturated flows in WLAN c .

The network then satisfies Assumptions 4-7. Observe that the per flow queueing discipline trivially ensures that $y(c) = \bar{x}(c)$ (saturated flows will transmit a packet at every transmission opportunity). By Theorem 3 we then have an equivalence between bottlenecked flows and saturated flows. This equivalence is of fundamental importance. Specifically, suppose each flow uses ideal congestion control i.e. adjusts the flow rate to ensure that the flow is saturated at one or more WLANs without incurring queue overflow losses. Then congestion control will ensure that every flow is bottlenecked and so, without further effort, by Theorem 3 the network throughput allocation will be max-min fair. That is, we have the following important corollary of Theorem 3.

Corollary 4: Suppose each flow uses ideal congestion control and each WLAN in a mesh network is configured as stated above. Then the resulting flow throughput allocation is max-min throughput fair.

Of course, in practice we must work with real rather than ideal congestion control. Nevertheless, under suitable continuity conditions, we can expect that any congestion control algorithm that approximates ideal behaviour sufficiently closely will, by Corollary 4, yield a throughput allocation that is close to max-min fair and this is indeed confirmed in simulations, see Section VI.

The network configuration in Corollary 4 also requires that attempt probability parameter $y(c)$ is selected to maximise the throughput of saturated flows in a WLAN. This is considered in detail in the next section. However, we note briefly here that the reason for introducing Assumption 7 is that by appropriately selecting $\bar{p}(c)$ then it turns out that $y(c)$ can be found in a completely decentralised manner (i.e. no message-passing or packet-sniffing) using an approach similar to the idle-sense strategy for maximising WLAN throughput studied in [7]. Assumption 7 could alternatively be replaced by another constraint that simplifies selection of $y(c)$ so long as we retain log-convexity of the rate region. For example, as noted earlier we could simply impose the constraint that $y(c) = y$ for an appropriate fixed value y , in which case no adaptation is required (this corresponds to trivially selecting $CW_{min} = CW_{max} = CW$ where CW is some fixed value), although this appealing simplicity comes at the cost of a reduction in network capacity.




Fig. 2. Illustrating unconstrained rate region and rate region with P_{idle} constraint.

V. MAXIMISING THROUGHPUT

A. Rate region boundary

We begin by studying the boundary of the rate region of WLAN c . For this we will take a vector y , normalised such that $\sum_i y_i = 1$, and set $x_i(c) = \lambda y_i / N_i$, $\lambda \geq 0$. The vector of station throughputs is then $s = \lambda \frac{y}{X(c)} \frac{L}{T_c}$. Since λ , $X(c)$, L and $T_c(c)$ are all scalars it can be seen that varying λ adjusts the position of the throughput vector on the ray in direction y passing through the origin. To determine the rate region boundary we need to find the values of λ and N_i that solve the optimisation

$$\begin{aligned} & \max_{\lambda, N_i} \frac{\lambda}{X} \\ & \text{s.t. } \lambda \geq 0, \quad 0 \leq N_i \leq |\mathcal{P}_i|, i \in \{1, \dots, n(c)\} \\ & \prod_{i=1}^{n(c)} (1 + \lambda y_i / N_i) \leq \bar{p}(c) \end{aligned} \quad (11)$$

Since the objective is strictly increasing in N_i (as already noted) and constraint (11) becomes looser as N_i increases, at the maximum N_i will lie on the constraint $|\mathcal{P}_i|$. It can be verified by inspection of the second derivative that λ/X is a concave function of λ and so has a unique turning point. To find the maximising value of λ , we observe that this will be determined either by constraint (11) becoming active or by the turning point of λ/X , whichever occurs first. This is illustrated in Figure 2 – the dashed line marks the unconstrained rate region (i.e. without constraint (11)) and the solid curves mark the rate region boundary for different values of \bar{p} . For a sufficiently small value of \bar{p} it is the constraint (11) that determines the boundary of the rate region, see curve marked P_2 in the figure.

To determine the turning point of λ/X , and so the unconstrained rate region boundary (marked by the dashed line in Figure 2), differentiating λ/X with respect to λ yields

$$\frac{1}{X^2} \left(X - \lambda \left(\sum_{i=1}^n y_i + \sum_{i=1}^n \frac{y_i}{N_i} \prod_{j \neq i} (1 + \frac{\lambda y_j}{N_j}) - \sum_{i=1}^n \frac{y_i}{N_i} \right) \right)$$

and setting this derivative equal to zero we have that the λ^* corresponding to the turning point solves

$$\sum_{i=1}^n \frac{\lambda^* y_i}{N_i} \prod_{j \neq i} (1 + \frac{\lambda^* y_j}{N_j}) + 1 - a = \prod_{i=1}^n (1 + \frac{\lambda^* y_i}{N_i})$$

Substituting, we therefore have that the turning point (*i.e.*, boundary of the rate-region) satisfies

$$\sum_{i=1}^n \frac{x_i^*}{1+x_i^*} \prod_{j=1}^n (1+x_j^*) + 1-a = \prod_{j=1}^n (1+x_j^*)$$

This can be rewritten as

$$\sum_{i=1}^n \tau_i^* + (1-a)P_{idle}^* = 1$$

where $P_{idle}^* = \prod_{i=1}^n (1-\tau_i^*)$. Note that this is a generalization of the result from [19], [21] to the scenario with different slot lengths (*i.e.*, $a < 1$) and TxOP.

Using the Arithmetic Mean-Geometric Mean inequality, we have

$$\sum_{i=1}^n \frac{1}{1+x_i} \geq n \sqrt[n]{\prod_{i=1}^n \left(\frac{1}{1+x_i} \right)}$$

i.e.

$$1 - \frac{1}{\sqrt[n]{\prod_{i=1}^n (1+x_i)}} \geq \frac{1}{n} \sum_{i=1}^n \frac{x_i}{1+x_i}$$

After some algebra, it follows that selecting $\bar{p}(c) \leq 1/(1+a-\sqrt{2a})$ ensures that constraint (11) is guaranteed to become tight either before or at the turning point of λ/X . Note that $1+a-\sqrt{2a} \leq 1$ when $a \leq 1$ and this bound on $\bar{p}(c)$ is tight (with equality along the ray $\lambda \mathbf{1}$, where $\mathbf{1}$ denotes the all 1's vector, as $n \rightarrow \infty$). This is illustrated by the middle curve marked P_0 in Figure 2, which touches the unconstrained rate region along the 45 degree ray. With this choice of $\bar{p}(c)$ constraint (11) is active at the solution to the above optimisation and so it is this constraint that determines the maximum value of λ , and thereby the maximum throughput of saturated flows.

B. Decentralised optimisation

Recall that our task is to select attempt rate parameter $y(c)$ to maximise the throughput of saturated flows. Selecting $\bar{p} = 1/(1+a-\sqrt{2a})$ so as to maximise the constrained rate region, it follows from the discussion in the preceding section that the throughput of saturated flows is maximised when $P_{idle}(c) = 1/\bar{p}(c)$. That is, we need to select $y(c)$ such that $P_{idle}(c) = 1/\bar{p}(c)$. This can be achieved in an entirely decentralised manner since (i) the idle probability $P_{idle}(c)$ can be directly observed by all stations in a WLAN (via carrier-sense, see for example [17]) and (ii) algorithms such as AIMD can be used to ensure stations converge to using the same parameter $y(c)$, see for example [7].

C. Degree of sub-optimality

Using any non-zero value of $1/\bar{p}$ necessarily comes at the cost of a reduction in throughput. To see this note that when only a single station is active in a WLAN, and so no collisions are possible, then we ought to select the attempt probability equal to 1 (*i.e.* $y(c) \rightarrow \infty$) in order to maximise the throughput, in which case any value of $1/\bar{p}$ greater than zero must reduce throughput below its maximum value. Nevertheless, the throughput loss is generally small. For example, Figure 3 illustrates the throughput cost of selecting \bar{p} to ensure operation




Fig. 3. Cost of operating on $P_{idle} = 1 + a - \sqrt{2a}$ constraint vs number of stations in WLAN. $a = 1/100$, each station carries a single saturated flow.

on the $P_{idle}(c) = 1/\bar{p}(c)$ constraint. The figure plots the ratio of the throughput when $P_{idle}(c) = 1/\bar{p}(c) = 1 + a - \sqrt{2a}$ to the maximum possible throughput when there is no P_{idle} constraint. It can be seen that the throughput efficiency is remarkably high, with a throughput reduction of less than 0.5% (compared to the maximum possible throughput) even when only a single station is active. This is similar to the observation made in [7]. In return for this small cost we gain the advantage of a fully decentralised implementation with no message-passing. The final choice of whether the additional network capacity to be gained by message-passing warrants the additional complexity lies with the network designer.

VI. SIMULATION RESULTS

We illustrate the foregoing analysis via *ns2* packet-level simulations. We begin by considering a mesh network with the topology shown in Figure 4(a). Mesh points (MP) are marked by circles and client stations by triangles. Each WLAN operates on an orthogonal channel and MP0, MP1 are equipped with two radios to allow relaying of traffic between WLANs. Flows 0-2 travel one hop to MP0, flow 3 travels two hops to MP3, flows 4-7 travel one hop to MP1, flow 7 travels two hops to MP3. Flow 8 travels one hop from station 8 to MP2.

In the simulations all flows are long-lived TCP traffic and so are bidirectional (*i.e.* consisting of TCP data and TCP ACK packets). Following [13], TCP ACKs are prioritised so that their loss rate is negligible (link asymmetry leading to excessive loss of TCP ACKs is well known to induce unfairness due to disruption of ACK clocking and repeated TCP timeouts). The TCP ACK transmit time (including MAC ACK *etc*) is lumped in with the TCP DATA transmit time to obtain the T_s value for throughput formula (1). See [12] for a more detailed discussion of the accuracy of this approximation, but we note here the good agreement in Figure 4(b) between the theory values derived using this assumption and the simulation measured throughputs.

The stations in each WLAN measure the idle probability P_{idle} using their carrier-sense functionality (*e.g.* see [17]) and run a local AIMD algorithm to adjust their CW_{min} to satisfy the constraint $P_{idle} \geq 1/\bar{p} = 1 + a - \sqrt{2a}$, see Algorithm

VI for details. Due to the use of the AIMD algorithm the station CW_{min} 's vary over time in a sawtooth pattern and do not settle on a constant value, see Figure 5(a). Moreover, CW_{min} is restricted to take integer values thereby introducing further granularity. By adjusting the AIMD β parameter the amplitude of the CW_{min} sawtooth can be changed. Decreasing β reduces the size of the CW_{min} fluctuations, but this comes at the cost of slower convergence to steady-state operation, *e.g.* see [23] for a detailed analysis of AIMD dynamics. We choose $\beta = 0.25$ as a compromise between fast convergence and reasonably small fluctuations in CW_{min} . Due to these implementation issues, as can be seen from Figure 5(b), the WLANs do not operate exactly on the $P_{idle} = 1/\bar{p}$ constraint as assumed in the calculation of the theoretical throughput values shown in Figure 4(b). Nevertheless, as can be seen from Figure 3 the throughput efficiency is relatively insensitive to P_{idle} fluctuations around the optimum value and this is reflected in the good agreement between the theory and simulation throughputs in Figure 4(b).

Other simulation parameters used are detailed in Table I. Figure 4(b) compares the theoretical max-min fair throughput allocation with the measured simulation throughputs. It can be seen that they agree remarkably well. We can investigate the structure of the throughput allocation in the simulations in more detail. By inspection of the topology in Figure 4(a) we expect that the max-min throughput allocation has flows 0-3 bottlenecked at the left-hand WLAN, flows 4-7 at the right-hand WLAN and flow 8 at the centre WLAN. Figure 6 plots the flow throughputs in each WLAN, from which it can be seen that flows 0-3 are indeed the maximal throughput flows in the left-hand WLAN and similarly for flows 4-7 and flow 8 in the right-hand and centre WLANs respectively. By inspection of the station queue occupancies (not plotted here), we can also confirm that flows 0-3 are saturated in the left-hand WLAN, and similarly for flows 4-7 and flow 8 at their respective bottlenecks, in accordance with Theorem 3.

Figure 7 shows simulations results for a second topology. An additional WLAN has been added containing station 8 and MP0 now carries two flows, namely flow 3 and flow 8. Flow 8 is bottlenecked at the link between MP0 and MP3 while flow 3 is not, and simulations confirm that flow 8 is saturated at MP3 while flow 3 is not as per Theorem 3. Also note that in this modified topology the one-hop flow 8 is allocated a slightly higher throughput than in Fig 4 because there are now fewer collisions in the centre channel which is the bottleneck for this flow – MP0 and MP1 are transmitting data packets and MP3 transmitting TCP ACK packets, while in Fig 4 we additionally have traffic between station 8 and MP2 in this channel. Once again, observe that the simulation measurements agree extremely well with the theoretical max-min throughput allocation.

VII. TIME-BASED MAX-MIN FAIRNESS

We can readily extend the foregoing analysis to encompass weighted max-min fairness, *i.e.* where rather than max-min fairness of the flow throughputs $s(p)$, $p \in \mathcal{P}$ we require max-min fairness of the weighted flow throughputs $s(p)/w(p)$, $p \in$

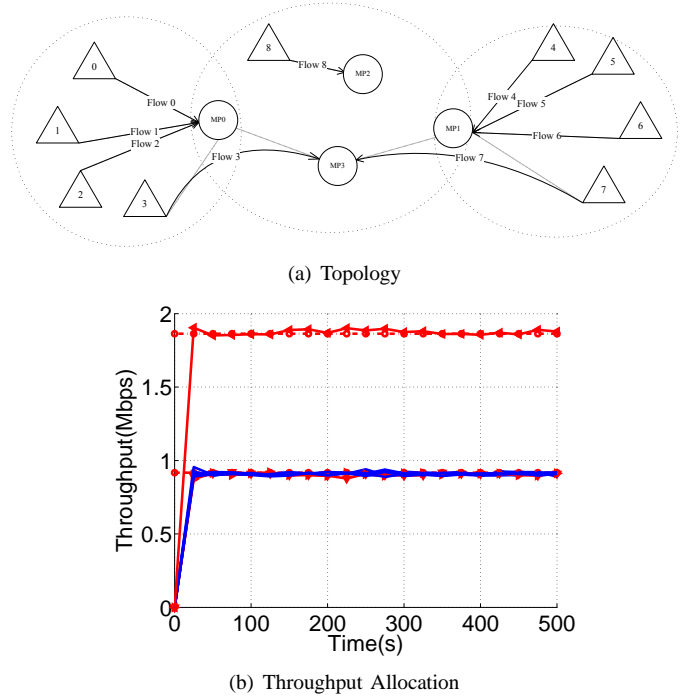


Fig. 4. Example 1. The measured simulation throughputs are compared against the theoretical max-min fair throughput values in the lower plot. The lower valued lines are the throughputs for flows 0-7, while the upper valued lines are for flow 8. The measurement points plotted are averages over 50s time windows. The theory values for the bottlenecks are 0.92127Mbps and 1.8479Mbps respectively (indicated by the dashed red lines and also marked on y-axis by circles). It can be seen that the simulation values are in good agreement with theory.

Algorithm 1 AIMD algorithm used at each station to adjust its CW_{min} value.

```

1: for Every T seconds do
2:   Check the measured idle probability  $P_{idle}$ 
3:   if  $P_{idle} > 1/\bar{p}$  then
4:      $CW_{min} \leftarrow CW_{min} + \alpha$ 
5:   else
6:      $CW_{min} \leftarrow CW_{min} \times (1 - \beta)$ 
7:   end if
8: end for

```

\mathcal{P} for specified weights $w(p) > 0$. This is of particular interest when we relax Assumption 4 that stations within a WLAN use the same PHY rate. When flows can use different PHY rates, max-min throughput fairness leads to flows with a low PHY rate grabbing bandwidth from higher PHY rate flows, potentially leading to a large reduction in network capacity.

PHY rate (Mbps)	11
NIC Buffer (Packets)	50
Packet Length (Bytes)	1000
α	4
β	0.25
T(s)	1
$1/\bar{p}$	0.8412

TABLE I
SIMULATION PARAMETERS



(a) CW_{min} (b) P_{idle}

Fig. 5. CW_{min} and WLAN P_{idle} time histories for the source station of flow 0 in the topology of Figure 4(a). These are representative of the time histories for other stations and illustrate the AIMD adjustment of CW_{min} . The dashed line in the lower plot indicates the ideal P_{idle} constraint value.

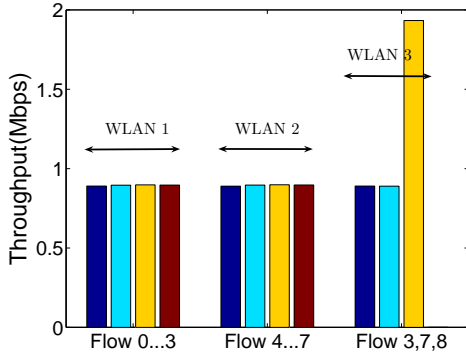
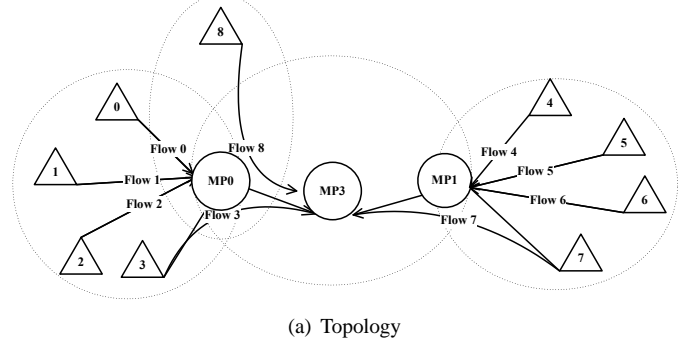



Fig. 6. Histogram of flow throughputs in each WLAN in topology Figure 4(a). WLAN 1 refers to the left-hand WLAN, WLAN 2 to the right-hand WLAN and WLAN 3 to the centre WLAN in Figure 4(a).

Time-based fairness is therefore typically of greater interest than throughput fairness in multi-rate networks, *e.g.* see [7], [26], [8] and references therein. Let $R(p)$ denote the PHY rate used by flow p , which for simplicity we assume is the same at every hop along the flow route $r(p)$. The airtime used by flow p is then given by $t(p) = s(p)/R(p)$ and so time-based fairness corresponds to weighted max-min fairness with weights $w(p) = R(p)$.

Since the airtime is just a rescaling of the throughput it follows that the feasible set of times is log-convex and a unique max-min time allocation exists. Retaining Assumptions 5-7 (for the moment), step 3 of the water-filling algorithm becomes



(a) Topology



(b) Throughput Allocation

Fig. 7. Example 2. As before, the measured simulation throughputs and theoretical max-min fair throughput values (indicated by dashed red lines and also marked on y-axis by circles) are compared in the lower plot.

$$\begin{aligned}
 & \max_{x, T^n} T^n \\
 & s.t. \quad t(p) = T^n \quad \forall p \in \mathcal{P}^n \\
 & \quad t(p) \leq \frac{x_k(c)}{X(c)}, \quad \forall p \in \mathcal{P}, \quad (k, \bullet, c) \in r(p) \\
 & \quad \sum_{p: (k, \bullet, c) \in r(p)} t(p) = \frac{N_k x_k(c)}{X(c)} \\
 & \quad x \geq 0 \\
 & \quad \prod_k (1 + x_k) \leq \bar{p}(c)
 \end{aligned}$$

An identical argument to that used in the proof of Theorem 3 can be applied (since $L/T_c(c)$ is just a constant scaling in the expressions used in the proof) to obtain

Theorem 4: Under Assumptions 1-3,5-7, the max-min fair time allocation within each WLAN possesses the following properties:

- 1) The attempt rate design parameter $y(c) \geq \bar{x}(c)$ in each WLAN where $\bar{x}(c)$ is the attempt rate that maximises the throughput of saturated flows.
- 2) Flows bottlenecked at the WLAN send one frame at every successful transmission made by the station. When $y(c) = \bar{x}(c)$, all bottlenecked flows are saturated. When $y(c) > \bar{x}(c)$ they are unsaturated.
- 3) Non-bottlenecked flows are always unsaturated.

It can be seen that the properties of the max-min time allocation are *identical* to those of the max-min fair throughput

allocation with a single PHY rate and so the same network configuration (together with ideal congestion control) can be used to realise the max-min time allocation *i.e.*

- 1) Stations in a WLAN all use the same attempt rate parameter $y(c)$.
- 2) Stations use per flow queueing and at each transmission opportunity send one frame from the head of each non-empty queue (recall that by Assumption 2 that all frames are of equal duration, regardless of the PHY rate used).
- 3) Parameter $y(c)$ is selected to maximise the throughput of saturated flows in WLAN c .

A. Simulation results

We revisit the previous simulation example in Figure 4, but now extend consideration to a multi-rate situation where flow 0 in the left-hand WLAN uses a PHY rate of 5.5 Mbps while all other flows in the mesh network use a PHY rate of 11 Mbps. Figure 8 compares simulation measurements with theoretical values for a max-min fair time allocation. It can be seen from Figure 8(a) that flow 0 (the flow with lower PHY rate) is now allocated a lower throughput than the other flows in the left-hand WLAN. This ensures that all flows in the left-hand WLAN are allocated the same air-time for transmitting their payloads, see Figure 8(b). Observe that the flows in the right-hand WLAN achieve slightly higher throughput and air-time than those in the left-hand WLAN due to the difference in frame overheads at different PHY rates.

VIII. ASSUMPTIONS

In this section we review the assumptions used in our analysis, and in particular try to identify those assumptions that can be readily relaxed and those that cannot. Assumption 1 (non-collision losses negligible) can be removed, but see the detailed discussion below. Assumption 2 (homogeneous frame transmission duration) can be readily relaxed to the requirement that stations have the same *mean* frame duration. Removing this assumption altogether should be possible but requires modifying the denominator (2) of the throughput formula to take account of the fact that the duration of a collision now depends on the specific set of stations involved in a collision and so on the attempt rates x_i . Assumption 3 (throughput model) is the fundamental assumption used in our analysis. This assumption might be weakened in various ways, but is not straightforward to remove. As discussed in Section VII, it is trivial to remove Assumption 4 (homogeneous PHY rates) and so accommodate multi-rate operation and time-based fairness. Assumptions 5 and 6 can be removed, but similarly to Assumption 1 this is at the cost of a considerable increase in the practical difficulty of realising a max-min allocation. See the following sections for a detailed discussion, but we note here that perhaps the most notable casualty is that by relaxing these assumptions we lose the equivalence between bottlenecked flows and saturated flows. This means that standard flow congestion control algorithms (which work by developing a queue backlog) can no longer be relied upon to guarantee flows are bottlenecked. As already commented upon already, Assumption 7 can be replaced by a variety of

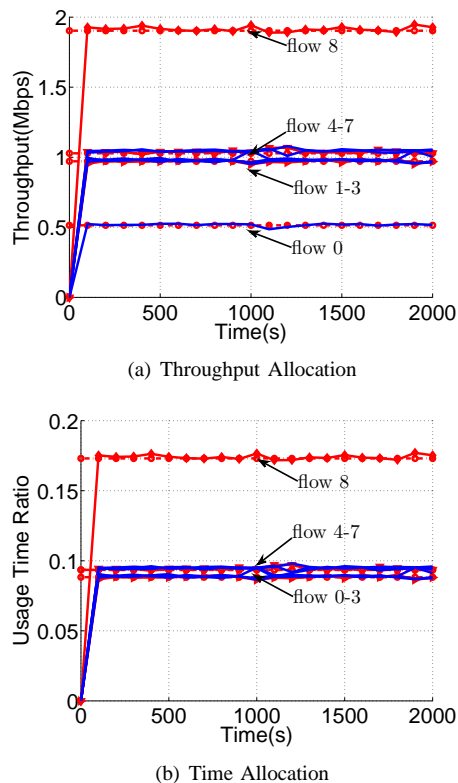


Fig. 8. Multi-rate variant of Example 1. Flow 0 uses PHY rate of 5.5Mbps, other flows a rate of 11Mbps. Plots compare simulation measurements and theoretical values (indicated by dashed red lines and also marked on y-axis by circles) of a max-min fair time allocation. The measurement points plotted are averages over 50s time windows.

alternative constraints provided we retain log-convexity of the network rate region.

Lastly, we note that while we have assumed that stations have sufficient arriving traffic to be able to make full use of the max-min fair throughput allocation, our analysis carries over essentially unchanged to situations where the rate of traffic arrivals at stations is itself constrained. The upper bound on throughput created by the finite traffic load introduces an additional convex constraint, and this constraint becomes the bottleneck when it is less than the max-min fair allocation in the absence of the finite-load constraint.

A. Relaxing Assumption 1: non-collision losses negligible

In this section we consider in more detail what is involved in relaxing Assumption 1. The main non-collision sources of loss are channel noise losses, packet discards after too many retries and queue overflow losses. We begin by noting that excessive channel noise losses can be avoided by appropriate choice of modulation/coding rate, discard losses by use of an appropriate retry limit (the standard value of 11 retries requires a combined channel-noise/collision loss rate exceeding 65% for the discard probability to exceed 1%) and queue overflow losses by provisioning links with sufficient buffering. That is, Assumption 1 can often be satisfied by appropriate network design. When such losses cannot be neglected, more effort is required. Assume use of a block ACK so that TXOP burst

transmissions do not terminate early on detecting a corrupted packet (as they would with per packet ACKing). This ensures that the duration of TXOP burst transmissions is independent of the specific packet loss pattern experienced by each burst – the analysis could be extended to include such dependence, but at the cost of a considerable increase in complexity. TXOP transmissions may consist of multiple blocks destined to different receivers which undergo losses dependent on the receiver. Under such a model we can use the formulation from [25]. Let $s(p)$ now denote the goodput of flow p , *i.e.*, the rate received correctly at the destination. Let $A_{i,p}s(p)$ denote the rate at which station i has to send packets from flow p in order to ensure $s(p)$ goodput is received at the destination after undergoing losses at intermediate hops along the route $r(p)$ to the destination. The scaling term $A_{i,p} \geq 1$ is equal to 1 if and only if there are no losses along the route from station i to the destination of flow p . Log-convexity of the goodput rate region still holds and in equations (7) and (8) we now need to replace $s(p)$ with $A_{i,p}s(p)$ to obtain a revised water-filling algorithm that includes the effect of noise losses.

To maintain equal throughput for flows bottlenecked at the same WLAN the station attempt rates x_k have to be adjusted taking into account the term $A_{i,p}$. When $A_{i,p}$ is not the same for all stations then with per flow burst constraints those bottlenecked flows with smaller values of $A_{i,p}$ will be unsaturated *i.e.* we will lose the equivalence between bottlenecked and saturated flows. We illustrate this with an example. Consider the network in Figure 9 where the capacities and the loss rates on the links are chosen such that all of the flows are bottlenecked in WLAN A. With the restriction that every flow has a maximum burst-size of 1 (Assumption 5), it is easy to see that at the max-min fair solution flow 2 is bottlenecked in clique A but is unsaturated. This is despite the fact that all three flows get the same goodput. As noted above, fortunately such difficulties can be avoided by the simple expedient of selecting a modulation/PHY rate and retry limit such that losses can be neglected.

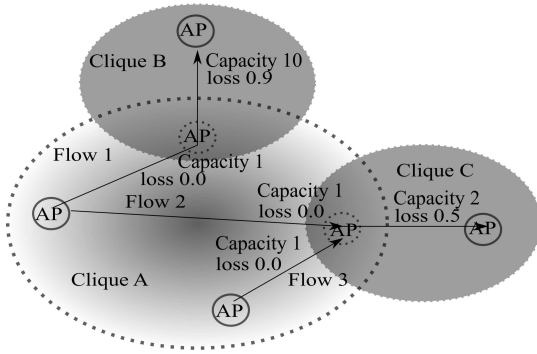


Fig. 9. Example network with losses to illustrate nature of max-min fair solution.

B. Relaxing Assumption 5: per flow burst-size constraint

We now consider in more detail removing Assumption 5. This removes constraint (7) from the water-filling algorithm and the relaxed optimisation in the proof of Theorem 3

becomes

$$\begin{aligned}
 & \max_{x_k, \bar{x}, N_k, n_k} \frac{\bar{x}(c)}{X(c)} \frac{L}{T_c(c)} \\
 & \text{s.t. } \bar{x}(c) \leq n_k x_k(c) \quad \forall k \in \mathcal{V}_B(c) \\
 & \sum_{q \notin \mathcal{B}(c): i(q,c)=k} s(q) \\
 & \leq \frac{N_k x_k(c) - |\mathcal{B}(c) \cap \mathcal{P}_k(c)| \bar{x}(c)}{X(c)} \frac{L}{T_c(c)} \quad \forall k \\
 & 1 \leq N_k(c) \leq \bar{N}_k \quad \forall k \\
 & x_k(c) \leq y(c) \quad \forall k \\
 & \prod_k (1 + x_k) \leq \bar{p}(c), \quad x_k(c) \geq 0, \quad \bar{x} \geq 0 \quad \forall k
 \end{aligned}$$

where n_k is the burst size used by bottlenecked flows at station k (which must be the same for all bottlenecked flows carried by station k since these flows have the same throughput $(n_k x_k / X(c)) (L/T_c)$). Using similar arguments as those in the proof of Theorem 3, the first three constraints will be tight at the optimum. That is, the burst size n_k will be such that $N_k = \bar{N}_k$ (*i.e.* the maximum admissible value) and the station attempt rate is correspondingly adjusted to maintain $\bar{x} = n_k x_k$. In general, the burst size n_k and attempt rate x_k will therefore now be different for every station carrying bottlenecked flows (depending on both the number of bottlenecked flows carried by a station and the load imposed by non-bottlenecked flows). The WLAN attempt rate parameter $y(c) \geq \max_k \bar{x}/n_k$. Due to the maximisation over k needed here, we may have $x_k < y(c)$ for some stations carrying bottlenecked flows *i.e.* there can exist bottlenecked flows which are unsaturated for all admissible values of $y(c)$ and we lose the equivalence between bottlenecked and saturated flows. Moreover, it seems clear that stations will generally need to communicate in order to agree the value of \bar{x} and enforce constraint $\bar{x} = n_k x_k$ (equality of bottleneck flow throughputs). In particular, the selection of \bar{x} is no longer amenable to the decentralised P_{idle} optimisation approach used previously.

C. Relaxing Assumption 6: homogeneous station attempt rate parameters

Removing Assumption 6 removes constraint (17) from the relaxed optimisation in the proof of Theorem 3. For stations carrying bottlenecked flows this change has little effect – all such stations must still use the same attempt rate \bar{x} . In contrast, for stations which carry no bottlenecked flows the attempt rate design parameter can now be selected equal to x_k in which case some of the non-bottlenecked flows will be saturated. That is, once again we lose the equivalence between bottlenecked and saturated flows.

IX. THE HIDDEN TERMINAL ISSUE

Perhaps the most significant omission from our analysis is hidden terminals. The basic difficulty here is that we currently lack simple, accurate, generally applicable throughput models when hidden terminals are present, and so we lack the basic tool needed for any max-min fairness analysis. The modelling difficulty arises from the fact that hidden terminals can start

transmitting even when a transmission by another station has already been in progress for some time. The class of slotted-time models pioneered by Bianchi for 802.11 is therefore no longer valid, since these require all transmissions to occur on well-defined MAC slot boundaries, and indeed this suggests that a fundamental change in modelling paradigm is required. The development of throughput models in the presence of hidden terminals continues to be the subject of an active research effort, and so in this paper we consider it prudent to leave consideration of utility fairness with hidden terminals to future work.

It is perhaps also worth noting here that the prevalence of severe hidden terminals in real network deployments presently remains unclear. While it is relatively easy to construct hidden terminal configurations in the lab that exhibit gross unfairness, it may well be that such configurations are uncommon in practical deployments. For example, recent measurement studies report that severe hidden terminal effects typically affect only a relatively small subset of stations in the WLAN deployments considered, *e.g.* see [4], [20]. In mesh network deployments it additionally seems likely that network designers will proactively seek to avoid (or at least minimise) creating hidden terminals thereby further reducing their impact. In addition to appropriate placement of mesh points, hidden terminals can be avoided/mitigated by judicious radio channel assignment and power control (*e.g.* see [14] and references therein). Looking to the future, the latter solutions are facilitated by the trend in next generation networks towards multi-radio architectures and the use of the 5GHz band for mesh backhaul (with its greater number of orthogonal channels compared to the 2.4GHz band).

Setting the hidden terminal issue to one side for the moment therefore, we stress that the class of mesh networks considered here is a substantial step beyond Aloha, previously the state of the art in wireless utility-fair analysis. In contrast to Aloha, this class is indeed sufficiently powerful and general to encompass at least some real 802.11 mesh network implementations. As support for this we comment that we have already implemented one of the max-min fair approaches derived here in an experimental 802.11 testbed using standard hardware and we will report our experimental measurements in due course.

X. CONCLUSIONS

In this paper we characterise, for the first-time, max-min fair rate allocations for a large class of 802.11 mesh networks. To our knowledge, this is also the first work to extend max-min fair mesh network analysis beyond Aloha networks. The class of 802.11 mesh networks considered is large enough to cover realistic network architectures and, by exploiting the features of the 802.11e/n MAC (in particular TXOP packet bursting), we are able to use this characterisation to establish a simple class of network configurations for achieving max-min throughput fairness. We demonstrate the efficacy of this approach using detailed packet-level simulations and establish that the approach can be readily extended to encompass time-based fairness in multi-rate 802.11 mesh networks.

XI. ACKNOWLEDGEMENTS

The authors would like to thank colleague Ken Duffy for his numerous insightful comments and helpful discussions relating to this paper.

APPENDIX A

APPENDIX – PROOF OF THEOREM 3

We proceed by analysing the optimisation (5)-(10) at step 3 of the water-filling algorithm. Let c denote a WLAN which becomes a bottleneck at iteration n of the algorithm. When considering bottlenecked flows at WLAN c we can ignore the constraints at other WLANs since these constraints must be either loose (or else that WLAN would be the flow bottleneck) or equivalent to the constraints at WLAN c (in the case of a flow having multiple bottlenecks). Flows which are not bottlenecked at WLAN c must be bottlenecked at other WLANs and the constraints at these WLANs determine the throughput of these flows. Let $\mathcal{B}(c)$ denote the set of flows bottlenecked at WLAN c and $\mathcal{V}_{\mathcal{B}}(c) = \{j \in \mathcal{V} : p \in \mathcal{B}(c), p \in \mathcal{P}_j(c)\}$ denote the set of stations carrying one or more bottlenecked flows. For bottlenecked flows we have that

$$s(p) = s(q) = T^n = \frac{\bar{x}(c)}{X(c)} \frac{L}{T_c(c)} \quad \forall p, q \in \mathcal{B}(c)$$

for some $\bar{x} \leq x_i \quad \forall i \in \mathcal{V}_{\mathcal{B}}(c)$. This bottleneck flow throughput is strictly greater than the throughputs of non-bottlenecked flows traversing the WLAN. By Assumptions 2-5 all flows q satisfy $s(q) \leq \frac{x_k(c)}{X(c)} \frac{L}{T_c(c)}$. Let us relax, for the moment, equality in (4) and replace it by the RHS upper bounding the LHS. By Assumption 6 all stations use the same attempt probability design parameter $y(c)$ and $x_k(c) \leq y(c)$ for every station k . Combining these observations, leads us to consider the following relaxed optimisation problem,

$$\max_{x_k, \bar{x}, N_k} \frac{\bar{x}(c)}{X(c)} \frac{L}{T_c(c)} \quad (12)$$

$$s.t. \quad \bar{x}(c) \leq x_k(c) \quad \forall k \in \mathcal{V}_{\mathcal{B}}(c) \quad (13)$$

$$s(q) \leq \frac{x_k(c)}{X(c)} \frac{L}{T_c(c)} \quad \forall q \in \{p \notin \mathcal{B}(c) : i(p, c) = k\}, \forall k \quad (14)$$

$$\begin{aligned} & \sum_{q \notin \mathcal{B}(c) : i(q, c) = k} s(q) \\ & \leq \frac{N_k x_k(c) - |\mathcal{B}(c) \cap \mathcal{P}_k(c)| \bar{x}(c)}{X(c)} \frac{L}{T_c(c)} \quad \forall k \end{aligned} \quad (15)$$

$$1 \leq N_k(c) \leq |\mathcal{P}_k(c)| \quad \forall k \quad (16)$$

$$x_k(c) \leq y(c) \quad \forall k \quad (17)$$

$$\prod_k (1 + x_k) \leq \bar{p}(c), \quad x_k(c) \geq 0, \quad \bar{x} \geq 0 \quad \forall k \quad (18)$$

where $i(p, c)$ denotes the access point relaying flow p on channel c . It can be verified that this relaxed optimisation can be transformed into a convex problem and so has a unique solution⁴.

⁴Change variables to $\log x_k$, $\log N_k$ and $\log \bar{x}$. X is a posynomial and so when expressed in terms of these transformed variables $\log X$ is the log sum of exponentials and convex.

Consider the following constraints on station $k \in \mathcal{V}_B(c)$ carrying at least one bottlenecked flow,

$$\bar{x}(c) \leq x_k(c) \quad (19)$$

$$s(q) \leq \frac{x_k(c)}{X(c)} \frac{L}{T_c(c)} \quad \forall q \in \{p \notin \mathcal{B}(c) : i(p, c) = k\} \quad (20)$$

$$x_k(c) \leq y(c) \quad (21)$$

The last constraint is satisfied provided $y(c) \geq \bar{x}$ – we return to the choice of $y(c)$ shortly. It can be verified (e.g. by inspecting derivatives with respect to x_k) that x_k and x_k/X are strictly increasing in x_k , while $1/X$ is strictly decreasing in x_k . Hence, if inequalities (19) and (20) are both loose then decreasing $x_k(c)$ decreases the RHS while improving the cost function and making the other inequalities looser, which leads to a contradiction. Hence we may must have equality in either/both (19) and (20) (for at least one q). Recalling that for non-bottlenecked flows $s(q) < \frac{\bar{x}}{X(c)} \frac{L}{T_c(c)} \leq \frac{x_k}{X(c)} \frac{L}{T_c(c)}$, it can be seen that constraint (19) will always become tight before constraint (20). Hence, we must have equality in (19). That is, $x_k = \bar{x}$ for all stations carrying a bottlenecked flow and for any bottlenecked flow the burst-size used is exactly one frame per successful transmission by the station. For non-bottlenecked flows the average burst size per successful transmission by the station must be strictly less than one frame, which implies that these flows are unsaturated.

Turning now to station $k \notin \mathcal{V}_B(c)$ that carries no bottlenecked flows, constraint (19) no longer applies but (21) and (20) are still in force. Since all flows on the station are, by assumption, non-bottlenecked they have throughput strictly less than $\frac{\bar{x}}{X(c)} \frac{L}{T_c(c)}$. Hence, if we have equality in (20) for one or more flows then $x_k(c) < \bar{x}$. But from (21) and the discussion in the foregoing paragraph $y(c) \geq \bar{x}$ and so $x_k(c) < \bar{x} \leq y$. Since $x_k < y$ the station is unsaturated and therefore also every flow is unsaturated. If we have inequality in (20) for all flows then the average flow burst size must be strictly less than one frame per successful transmission by the station which implies that, once again, every flow is unsaturated.

To gain insight into the burst size N_k , we need to consider constraint (15). Since $N_k x_k/X$ is increasing in N_k and $1/X$ is decreasing, using a similar contradiction argument as previously we must have equality in (15) for all stations.

Consider now the value of $y(c)$. It can be seen that $x_k(c)$ is invariant in $y(c) \geq \bar{x}$. Hence any $y(c) \geq \bar{x}$ is an admissible solution and yields the same allocation of x_k 's and N_k 's. Since we have equality in (15), these solutions to the relaxed optimisation are also feasible for the true/unrelaxed constraints. Observe, however, that when $y > \bar{x}$ no flow is saturated (for stations $k \in \mathcal{V}_B(c)$, $x_k = r_k(y) = \bar{x} < y$ and so the stations are unsaturated and thus every flow must be unsaturated, for stations $k \notin \mathcal{V}_B(c)$ we already have that every flow is unsaturated). When $y = \bar{x}$ we have that all bottlenecked flows are saturated and all non-bottlenecked flows are unsaturated (for stations $k \in \mathcal{V}_B(c)$, $x_k = r_k(y) = \bar{x} = y$ and so the station is saturated plus bottlenecked flows send one packet at every successful transmission by a station and so are also saturated since a flow cannot know in advance which

transmissions will be successful, for all stations we already have that non-bottlenecked flows are unsaturated). Observe also that while we have some freedom in the choice of y , since the max-min allocation for the original problem is unique the values of the x_k 's and N_k 's (which are invariant in $y \geq \bar{x}$) are unique.

REFERENCES

- [1] V. Bharghavan, A. Demers, S. Shenker, L. Zhang, "MACAW: A Media Access Protocol for Wireless LAN's," *Proc. ACM SIGCOMM*, pp. 212–225, 1994
- [2] S. Boyd, L. Vandenberghe, "Convex Optimisation," Cambridge University Press, 2008
- [3] S. Boyd, S.-J. Kim, L. Vandenberghe, and A. Hassibi, "A Tutorial on Geometric Programming," *Optimization and Engineering*, 8(1), pp. 67–127, 2007.
- [4] Y.-C. Cheng, et. al., "Jigsaw: Solving the Puzzle of Enterprise 802.11 Analysis," *Proc. ACM SIGCOMM*, 2006.
- [5] L. B. Jiang, S. C. Liew, "Improving Throughput and Fairness by Reducing Exposed and Hidden Nodes in 802.11 Networks," *IEEE Transactions on Mobile Computing*, 7(1), 34–49, 2008
- [6] P. Gupta, A. L. Stolyar, "Optimal Throughput Allocation in General Random-Access Networks," *Proc. CISS*, 2006.
- [7] M. Heusse, F. Rousseau, R. Guillier, A. Duda, "Idle sense: an optimal access method for high throughput and fairness in rate diverse wireless LANs," *Proc. ACM SIGCOMM*, 2005.
- [8] L. B. Jiang, S. C. Liew, "Proportional Fairness in Wireless LANs and Ad Hoc Networks," *Proc. IEEE WCNC*, pp.1551– 1556, 2005.
- [9] K. Kar, S. Sarkar, L. Tassiulas, "Achieving Proportional Fairness Using Local Information in Aloha Networks," *IEEE Trans. Auto. Control*, 49(10), pp. 1858–1862, 2004.
- [10] A. Kochut, et al., "Sniffing out the correct physical layer capture model in 802.11b," *Proc. IEEE ICNP*, 2004.
- [11] J. W. Lee, M. Chiang, A. R. Calderbank, "Jointly Optimal Congestion and Contention Control Based on Network Utility Maximization," *IEEE Communications Letters*, 10(3), pp. 216–218, 2006.
- [12] D. J. Leith, P. Clifford, "Modelling TCP Dynamics in Wireless Networks," *Proc. Wirelesscom*, 2005.
- [13] D. J. Leith, P. Clifford, A. Ng, "TCP Fairness in 802.11e WLANs," *IEEE Communications Letters*, 9 (11), 964– 966, 2005.
- [14] D. J. Leith, and P. Clifford, "A Self-Managed Distributed Channel Selection Algorithm for WLANs," *Proc. ACM/IEEE RAWNET*, 2006.
- [15] D. J. Leith, V. G. Subramanian and K. R. Duffy, "Log-convexity of rate region in 802.11e WLANs," *IEEE Communications Letters*, 14 (1), pp. 57–59, 2010.
- [16] R.M. Loynes, "The stability of a queue with non-independent inter-arrival and service times", *Proc. Camb. philos. Soc*, 58, 497–520, 1962
- [17] D. Malone, P. Clifford, D. J. Leith, "MAC layer channel quality measurement in 802.11," *IEEE Communications Letters*, 11(2), pp. 143–145, 2007.
- [18] D. Malone, K. Duffy, and D. Leith, "Modeling the 802.11 Distributed Coordination Function in Nonsaturated Heterogeneous Conditions," *IEEE/ACM Trans. Networking*, 15(1), pp. 159–172, 2007.
- [19] J. Massey and P. Mathys, "The collision channel without feedback," *IEEE Trans. Inform. Theory*, 31(2), pp. 192–204, 1985.
- [20] D. Niculescu, "Interference map for 802.11 networks," *Proc. 7th ACM SIGCOMM conference on Internet measurement*, 2007
- [21] K. A. Post, "Convexity of the Nonachievable Rate Region for the collision channel without feedback," *IEEE Trans. Inform. Theory*, 31(2), pp. 205–206, 1985.
- [22] B. Radunovic, J.-Y. Le Boudec, "A unified framework for max-min and min-max fairness with applications," *IEEE/ACM Trans. Networking*, 15(5), pp. 1073–1083, 2007.
- [23] R. N. Shorten, D. J. Leith and F. Wirth, "Products of random matrices and the internet: Asymptotic results," *IEEE/ACM Trans on Networking*, 14(6), pp. 616–629, 2006.
- [24] V. A. Siris, G. Stamatakis, "Optimal CWmin Selection for Achieving Proportional Fairness in Multi-Rate 802.11e WLANs," *Proc. WinTECH*, 2006.
- [25] 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*, 8(7), pp. 3401–3406, 2009.
- [26] G. Tan, J. Gutttag, "Time-based fairness improves performance in multi-rate WLANs," *Proc. USENIX Annual Technical Conference*, 2004.

- [27] X. Wang, K. Kar, “Distributed Algorithms for Max-Min Fair Rate Allocation in ALOHA Networks,” *Proc. Allerton Conference*, 2004.

PLACE
PHOTO
HERE

Doug Leith graduated from the University of Glasgow in 1986 and was awarded his PhD, also from the University of Glasgow, in 1989. In 2001, Prof. Leith moved to the National University of Ireland, Maynooth to assume the position of SFI Principal Investigator and to establish the Hamilton Institute (www.hamilton.ie) of which he is Director. His current research interests include the analysis and design of network congestion control and resource allocation in wireless networks.

PLACE
PHOTO
HERE

Qizhi Cao

PLACE
PHOTO
HERE

Vijay Subramanian received the B.Tech. degree from IIT Madras, in 1993, the M.S. degree from the Indian Institute of Science, Bangalore, in 1995, and the Ph.D. degree from the University of Illinois at Urbana-Champaign, Urbana, in 1999. From 1999 to 2006, he was with the Networks Business, Motorola, Arlington Heights, IL where he worked on developing wireless scheduling algorithms deployed in many of Motorola’s wireless data products. Since May 2006 he is a Research Fellow at the Hamilton Institute, NUIM, Ireland. His research interests include

information theory, communication networks, queueing theory, mathematical immunology and applied probability.