Existence and Uniqueness of Fair Rate Allocations in Lossy Wireless Networks

Vijay G. Subramanian, Ken R. Duffy, and Douglas J. Leith

Abstract—To extend established concepts of fair resource allocation in wired networks to wireless networks, wired model assumptions must be adapted to be relevant for wireless networks as for example, in wireless networks losses due to environmental conditions may occur even in the absence of queueing congestion. Thus fundamental questions of the existence and uniqueness of fair rate allocations must be reconsidered. We treat wireless networks characterized by lossy channels, spatial channel reuse, multiple routes and multiple frequencies. We establish the existence and uniqueness of utility fair and max-min fair solutions and that, as loss rates decrease, fair allocations converge to the loss-less ones.

Index Terms—Lossy networks, utility fairness, max-min fairness, location of bottlenecks, convergence of fair solutions.

I. INTRODUCTION

THERE is considerable interest in achieving fair resource allocation in multi-hop wireless networks. Most contemporary work focuses on identifying algorithms that enable the discovery of max-min fair solutions, e.g. [1]–[6]. Some of this work treats wireless-network specific features such as frequency reuse at non-interfering distances. However, with the notable exception of [9], the only works that we are aware of that explicitly treat the lossy nature of wireless networks, such as [7], [8], do so by assuming that losses are sufficiently small that they can be ignored at each hop and tallied at the receiver when calculating utility. This is equivalent to assuming losses only occur at the receiver and we refer to it as the *last-link loss approximation*. Here we consider the case where losses can occur at any (and every) link and this is reflected in the output of each link. We show that this can yield significantly different fairness solutions to those given by the last-link loss approximation.

There are two ways one could extend notions of the utility of a flow to wireless networks. One is to make each flow's utility be a function of the bandwidth it receives at each link. The other, which we adopt, is to define the utility of a flow to be a function of its goodput at the receiver, so that a flow gets no utility for data lost in transit. This ties in naturally with max-min fairness, where it is clear that fair solutions should be determined by their receive-rate. Starting from this viewpoint, the mathematical framework we introduce enables us to establish the fundamental questions of existence and uniqueness of utility fair and max-min fair solutions for lossy networks that have multiple routes, resource reuse at non-interfering distances and multiple transmit/receive antennas. We prove that as loss rates converge to zero, the utility fair and max-min fair solutions converge to their loss-less equivalents. A by-product of our formulation enables us to deduce that with the last-link loss approximation utility fair and max-min fair solutions also converge to their loss-less equivalents. This helps to justify that approximation in the presence of low levels of loss, but we also demonstrate by example that it can lead to inaccurate solutions.

We introduce a generalized notion of bottleneck links and prove that max-min fair solutions can be characterized by (and characterize) each flow's bottlenecked links, but introduce an example that demonstrates bottlenecked links do not necessarily converge as loss rates tend to zero even though the corresponding max-min fair solutions do. Through examples, we show how the framework can be used to study the nature of fair solutions in WiFi mesh networks with lossy links and illustrate new phenomena in the fair solutions that occur in wireless networks.

II. MODEL ASSUMPTIONS

Concepts of fair allocation of bandwidth in wired networks date back to at least the early 1980s. A summary of early developments can be found in [10], where the focus was originally on max-min fairness. The introduction of the notion of proportional fairness [11], which is equivalent to maximizing the sum of a concave utility function of the goodput of each flow, was a major development. While max-min fairness cannot be placed directly in that framework, it arises as the limit of a sequence of the widely adopted (w, α) fair solutions [12]. Fundamental assumptions in wired network frameworks include: (1) flows are modeled as fluid; (2) each link has a fixed capacity that can be allocated arbitrarily between flows; (3) the output of a link is unchanged before becoming the input to another link or departing the network; and (4) the medium is point-to-point with no interaction between distinct links. Some of these are appropriate for wireless networks, but others need to be reconsidered. We adopt (1) and (2), but - roughly speaking - adapt (3) and (4) to: (3') links can be lossy (for each link a proportion of every flow passing through it is lost); and (4') links may have joint constraints. (3') corresponds to losses due to environmental conditions and (4') includes multiple-access channels and primary interference constraints. We define the goodput of a flow to be the rate received at its destination and the goodput of the network to be the sum of the goodputs of all flows.

III. EXISTENCE, UNIQUENESS AND CONVERGENCE OF FAIR SOLUTIONS

Utility fairness. Using the notation in [13], we represent a network by a directed multigraph $\vec{\mathcal{G}} = (\mathcal{N}, \vec{\mathcal{E}})$ with nodes $\mathcal{N}$ and edges $\vec{\mathcal{E}}$, and with a set $\mathcal{P} = \{1, 2, \cdots, P\}$ representing data flows. Nodes represent stations and an edge exists from node a to node b if a can send data to b. Let N, E and P denote the cardinality of the sets $\mathcal{N}$, $\vec{\mathcal{E}}$ and $\mathcal{P}$. Associated with each flow p is a source node $s(p)$, a destination node $d(p)$, and a single fixed route consisting of edges $r(p)$ of $\vec{\mathcal{E}}$ connecting $s(p)$ to $d(p)$ without a cycle. Multiple routes are more likely to occur in wireless networks than in wired networks due to the possibility of many non-overlapping frequencies/channels being used in a single physical space. There is no technical difficulty in having multiple flows taking distinct routes between a source and destination pair and this can be incorporated by defining the goodput of a single super-flow to be the sum of the goodputs over sub-flows. For every edge e along route $r' : \mathcal{P} \mapsto 2^{\vec{\mathcal{E}}} \setminus \emptyset$, where $2^{\vec{\mathcal{E}}}$ is the power-set of $\vec{\mathcal{E}}$ and $\emptyset$ is the empty set, define the proximity $g(r', e)$ of the edge to the destination by the number of edges (including itself) to the end node of route r'. For edges that do not belong to route r' define $g(r', e) := \infty$. Define a route r'-based order $\leq_{r'}$ on the set of edges that make up route r': for two links e and f in r' if $g(r', e) \leq g(r', f)$, define $e \leq_{r'} f$.

Let the maximum link-rate of each edge $e \in \vec{\mathcal{E}}$ be $c_e > 0$ and define C to be the $E \times E$ matrix with diagonal entries $C_{e,e} = c_e$ and $C_{e,e'} = 0$ if $e \neq e'$. To treat (3'), each edge represents a (possibly) lossy link that drops a certain fraction of the traffic being transmitted by each flow that traverses it. For each edge e define $q_e \in (0, 1]$ to be the network-layer throughput, i.e. the fraction of traffic that is not dropped at edge e, and define $\mathbf{q}$ to be the corresponding $E \times 1$ vector. The rate successfully received from edge e is at most $q_e c_e$. Define the $E \times P$ connectivity/routing matrix $A(\mathbf{q})$ whose $(e, p)^{\text{th}}$ element is $A_{e,p} = 1/\prod_{f \in \vec{\mathcal{E}}: f \leq_{r(p)} e} q_f$ if $e \in r(p)$ and 0 if $e \notin r(p)$. If flow p has goodput 1, then its input to edge e is rate $A_{e,p}$. The non-zero elements of $A(\mathbf{q})$ are at least 1 because of the lossy nature of the links, whereas for loss-less networks the elements of A take values in $\{0, 1\}$ (e.g. [10], [12], [14]). The last-link loss approximation is modeled by replacing $A(\mathbf{q})$ with a matrix $A'(\mathbf{q})$ whose $(e, p)^{\text{th}}$ element is $A'(\mathbf{q})_{e,p} = \max_{e' \in r(p)} A_{e',p}$. The maximum over edges e' of $A_{e',p}$ identifies the greatest loss on the route of flow p in the lossy network. The routing matrix A', based on the last-link loss approximation, has the effect that all losses occur at each flow's last link and nowhere else on its route.

As we have started with a directed multigraph, the routing matrix corresponds to all links operating independently (e.g. full duplex). To encompass the (4') assumption and model shared wireless links between multiple nodes, we introduce a conflict matrix $B = [I^T, J^T]^T$ where T denotes transpose, I is the $E \times E$ identity and J is a $\zeta \times E$ matrix with $\{0, 1\}$ entries where $\zeta \in \{0, \ldots, 2^E - E - 1\}$, so that B is a $(E + \zeta) \times E$ matrix. I is the individual conflict matrix and gives each link its individual capacity constraint. J is the joint conflict matrix: if the links $e_{i_1}, \ldots, e_{i_K}$ do not operate independently, then we insert a row in J with entries 1 at

each e_{i_k} and 0 at all other positions. Due to the insertion of this row, these links will experience a joint constraint such as a shared wireless resource. There are at most $2^E - E - 1$ rows in the joint conflict matrix as it includes all subsets of the links apart from the E individual link constraints and the empty-set. We say that a flow p is *involved in a conflict* $i \in \{1, \ldots, E + \zeta\}$ and write $p \in i$ if for at least one $e \in r(p_i)$ we have that $B_{i,e} = 1$. As we will see in Section IV, the joint conflict matrix enables us to treat situations such as in WiFi networks where a single wireless resource is shared by two or more distinct nodes, while retaining an individual loss rate and distinct link rate for each pair of stations. We also introduce an $(E + \zeta) \times 1$ degrees of freedom vector $\mathbf{D}$ with non-negative entries that will represent either MIMO gains or the availability of multiple radio resources at each conflict.

Denote by $x_p \geq 0$ the goodput of flow p, i.e. the received rate at the flow's destination $d(p)$, and $\mathbf{x}$ the corresponding $P \times 1$ vector. With these quantities defined, the network places restrictions on possible goodputs through the following constraint:

$$BC^{-1}A(\mathbf{q})\mathbf{x} \leq \mathbf{D} \qquad \text{(goodput constraint)}. \qquad (1)$$

With I denoting the $E \times E$ identity matrix and $\mathbf{1}$ denoting the $E \times 1$ vector with all entries equal to 1, for a wired network $B = I$ and $\mathbf{D} = \mathbf{1}$ so that each link is directional and there are no joint conflicts. Moreover $\mathbf{q} = \mathbf{1}$, so that links are loss-less. The key observation for existence and uniqueness of utility fair solutions is that even if $B \neq I$ and $\mathbf{D} \neq \mathbf{1}$ and $\mathbf{q} < \mathbf{1}$ (entry-wise), equation (1) is still a linear constraint set. The set of goodput rate vectors $\mathbf{x}$ that satisfy equation (1) is called the *rate region* and is denoted by $\mathcal{X}(\mathbf{C}, \mathbf{q}) \subset \Re^P$. It is clear that $A(\mathbf{q}) \geq A(\mathbf{1})$ and that $\mathcal{X}(\mathbf{C}, \mathbf{q}) \subseteq \mathcal{X}(\mathbf{C}, \mathbf{1})$. Note that $A(\mathbf{q})_{e,p} \leq A'(\mathbf{q})_{e,p}$ for each link $e \in \vec{\mathcal{E}}$ and flow $p \in \mathcal{P}$. The consequence of this is that the rate region for the last-link loss approximation is necessarily smaller than that for the real lossy network. Defining a utility function $U : \Re^P \mapsto [-\infty, \infty)$ of the goodput $\mathbf{x}$, the fair allocation is an optimizer in the solution of the following optimization problem:

$$\sup_{\mathbf{x} \in \mathcal{X}(\mathbf{C}, \mathbf{q})} U(\mathbf{x}). \qquad (2)$$

The following proposition follows from the linearity of the constraints in equation (1) [15].

Proposition 1 (Existence and uniqueness of Utility Fair solutions)**.** *If U is a strictly concave function, then the optimization in* (2) *has a unique optimizer.*

We relate the arguments in the solution of the optimization (2) for $\mathbf{q} \leq \mathbf{1}$ to the loss-less case by proving a continuity property of the optimal solutions to (2) as $\mathbf{q}$ approaches $\mathbf{1}$. This is a consequence of showing that a stronger property holds: a type of set convergence [15] for the regions $\mathcal{X}(\mathbf{C}, \mathbf{q})$. Define the Pompeiu-Hausdorff distance [15] between two non-empty closed sets $D, E \subset \Re^P$ as $d_\infty(D, E) := \sup_{\mathbf{x} \in \Re^P} |d_D(\mathbf{x}) - d_E(\mathbf{x})|$, where $d_D(\mathbf{x}) := \inf_{\mathbf{y} \in D} d(\mathbf{y}, \mathbf{x})$ and $d(\cdot, \cdot)$ is the usual Euclidean metric on $\Re^P$. This metric is a well established measure of distance between closed sets and is widely used in the consideration of the convergence of optimization problems. The following theorem establishes our convergence

result for utility fair solutions. Its proof can be found in the Appendix.

Theorem 1 (Convergence of utility fair solutions). *Consider a sequence of link loss rates $\{1 - q_e^{(k)}\}_{k=1}^{\infty}$ such that $\lim_{k \to \infty} q_e^{(k)} = 1$ for each link $e \in \bar{\mathcal{E}}$. Then the rate regions in the lossy networks converge to the rate region in the corresponding loss-less network, $\lim_{k \to \infty} d_{\infty}(\mathcal{X}(\mathbf{C}, \mathbf{q}^{(k)}), \mathcal{X}(\mathbf{C}, \mathbf{1})) = 0$ and utility fair solutions converge to the corresponding loss-less utility fair solutions.*

This theorem also holds with the last-link loss approximation in force. However, in Section IV we present an example where loss rates are not small that will illustrate the failure of the last-link loss approximation.

Max-min fairness. A vector $\bar{\mathbf{x}}(\mathbf{C}, \mathbf{q})$ is *max-min fair on the set* $\mathcal{X}(\mathbf{C}, \mathbf{q})$ if and only if for all $\mathbf{x} \in \mathcal{X}(\mathbf{C}, \mathbf{q})$ there exists $p \in \mathcal{P}$ such that $x_p > \bar{x}_p \Longrightarrow \exists o \in \mathcal{P} \setminus \{p\}$ such that $x_o < \bar{x}_o \leq \bar{x}_p$ [10]. Max-min fairness does not correspond to the solution of (2) for any utility function, but arises as the limit of particular utility fair solutions. The (w, α) fair solution [12] uses the family of utility functions given for $w > 0$, $\alpha \geq 0$ and $x > 0$ by

$$U_{w,\alpha}(x) = \begin{cases} wx^{1-\alpha}/(1-\alpha) & \text{if } \alpha \neq 1; \\ w \log(x) & \text{if } \alpha = 1 \end{cases} \quad (3)$$

where we define $U_{w,\alpha}(0) := 0$ if $\alpha \in [0, 1)$ and $U_{w,\alpha}(0) := -\infty$ if $\alpha \geq 1$. For this family of strictly increasing utility functions we denote the unique maximizer of equation (2) as $\mathbf{x}^*(\mathbf{w}, \boldsymbol{\alpha}, \mathbf{C}, \mathbf{q})$, where $\mathbf{w} = (w_1, \ldots, w_P)^T$ and $\boldsymbol{\alpha} = (\alpha_1, \ldots, \alpha_P)^T$. Lemma 3 in [12] proves that max-min fair solutions arise as the limiting solution as $\alpha \to \infty$. As our network goodput constraints (1) are still linear, we can apply that lemma to see that the solutions $\mathbf{x}^*(\mathbf{1}, \alpha\mathbf{1}, \mathbf{C}, \mathbf{q})$ converge to $\bar{\mathbf{x}}(\mathbf{C}, \mathbf{q})$ as $\alpha \to \infty$.

Proposition 2 (Existence and uniqueness of Max-Min Fair solutions). *With goodput constraints given by equation (1), there exists a unique max-min fair solution.*

This max-min fair solution is unique by Theorem 1 of [14]. The following establishes that max-min fair solutions have the same convergence property as error rates tend to zero as utility fairness.

Corollary 1 (Convergence of max-min fair solutions). *As loss-rates tend to zero, max-min fair solutions converge to the loss-less max-min fair solution.*

As with Proposition 1 and Theorem 1, Proposition 2 and Corollary 1 continue to hold with the last-link approximation in lieu of the full lossy links formulation. This corollary is surprising as the original definition of max-min fairness [10] is in terms of bottlenecked links. For a sequence of networks with decreasing loss rates, in the next section we give examples where the location of these bottlenecked links do not converge, even though the max-min fair solutions do. First, however, we identify a suitable generalized definition of bottlenecked links that is appropriate in the present framework.

In the loss-less case, max-min fairness can, equivalently, be defined in terms of bottlenecked links [10]: given a feasible goodput rate vector $\mathbf{x} \in \mathcal{X}(\mathbf{C}, \mathbf{q})$, a link e is a *bottlenecked link* with respect to $\mathbf{x}$ for a flow p with link e along its route if $\sum_{p \in \mathcal{P}: e \in r(p)} x_p = c_e$ and $x_p \geq x_{p'}$ for all flows p' with link e along their routes. A feasible goodput rate vector $\mathbf{x} \in \mathcal{X}(\mathbf{C}, \mathbf{q})$ is then max-min fair if and only if each flow has a bottlenecked link with respect to $\mathbf{x}$ (e.g. [10] pg. 527). This definition yields a procedure called the water-filling algorithm to identify the max-min fair solution. The algorithm operates as follows: starting from the all zero goodput rate vector every flow's rate is increased until a first set of link constraints become active, i.e., bottlenecked for the flows that pass through these links. Only the rates of flows not passing through the bottlenecked links are increased further until another set of link constraints become active/bottlenecked. This procedure is repeated until all the flows pass through at least one bottlenecked link. The loss-less water-filling algorithm always identifies the max-min fair solution due to the coordinate convexity of the goodput rate region [14].

With lossy links we need to generalize the definition of a bottleneck link given above since each flow has a (potentially) different rate at the ingress and egress of each link along the routes that it traverses. Given a feasible goodput rate vector $\mathbf{x} \in \mathcal{X}(\mathbf{C}, \mathbf{q})$, we say that a conflict $i \in \{1, \ldots, E + \zeta\}$ is a *bottlenecked conflict* with respect to $\mathbf{x}$ for a flow p with at least one link in that conflict if $(BC^{-1}A(\mathbf{q})\mathbf{x})_i = \mathbf{D}_i$ and $x_p \geq x_{p'}$ for all flows p' at least one link in conflict i. That is, a conflict is a bottleneck conflict for a flow if the capacity constraint is met at that conflict and if no other flow involved in that conflict has higher goodput. The following Theorem, whose proof can be found in the Appendix, proves that bottlenecked conflicts provide a characterization of max-min fair solutions for networks with lossy links.

Theorem 2 (Bottlenecked conflict representation of lossy max-min fair solutions). *For a network with or without lossy links, a feasible goodput rate vector $\mathbf{x} \in \mathcal{X}(\mathbf{C}, \mathbf{q})$ is max-min fair if and only if each flow has a bottlenecked conflict with respect to $\mathbf{x}$.*

Due to Theorem 2 and the coordinate convexity of the goodput rate region, the obvious generalization of the water-filling algorithm necessarily identifies the unique lossy max-min fair solution.

IV. EXAMPLES

In the examples we assume that all flows have the same $(w, \alpha) = (1, 1)$ utility function (3), commonly called proportional fairness. Max-min fair solutions were obtained by the generalized water-filling procedure.

Example 1, Bottlenecked Conflicts and Continuity of Max-Min Fair Solutions:

Consider the three flow, two link network in Figure 1, where we assume that $c_2 = 1$. For this network, the matrices in equation (1) are $B = I$ and $\mathbf{D} = \mathbf{1}$, and

$$A(\mathbf{q}) = \begin{pmatrix} 1/(q_1 q_2) & 1/q_1 & 0 \\ 1/q_2 & 0 & 1/q_2 \end{pmatrix}.$$

Since Flows 2 and 3 pass through only one link each, they will be bottlenecked only on those conflicts. However, the bottlenecked conflicts for Flow 1 can change depending on

Fig. 1. Three flow, two link bottlenecked network example. Used to illustrate lack of convergence of bottlenecked conflicts despite the convergence of max-min fair solutions proved in Corollary 1. Also demonstrates the asymmetry in max-min fair solutions that is induced by lossy links.

Fig. 3. Two flow, two link network example illustrates a failing of the last-link loss approximation.


Fig. 2. Network in Figure 1. Continuity of max-min fair solutions.

the loss rates of links 1 and 2 and the capacity of link 1. For a max-min fair solution, define $\beta \in \{\{1\}, \{2\}, \{1,2\}\}$ to be the bottlenecked conflicts for Flow 1. Then we have that

$$\beta = \begin{cases} \{1\} & \text{if } \frac{2q_1 c_1}{1+q_2} < 1; \\ \{2\} & \text{if } \frac{2q_1 c_1}{1+q_2} > 1; \\ \{1,2\} & \text{otherwise.} \end{cases}$$

For the loss-less max-min fair solution with $q_1 = 1$, $q_2 = 1$ and $c_1 = 1$, the bottlenecked conflicts for Flow 1 are $\beta = \{1,2\}$. Consider a sequence of loss rates converging to 0, $\{q_1^{(n)}, q_2^{(n)}\}$ such that $q_1^{(n)} = q_2^{(n)} = 1 - 1/n$, then $\beta^{(n)}$ is always equal to $\{1\}$. Even though the max-min fair solutions are converging by Corollary 1, the bottlenecked conflicts for Flow 1 are not. Similarly, if the sequence of loss rates are $q_1^{(n)} = 1 - 1/n$ and $q_2^{(n)} = 1 - 3/n$ then $\beta^{(n)} = \{2\}$ for all n and if for (even) $n = 2m$ we have $q_1^{(n)} = q_2^{(n)} = 1 - 1/m$ and for (odd) $n = 2m + 1$ we have $q_1^{(n)} = 1 - 1/m$ and $q_2^{(n)} = 1 - 3/m$, then sequence of Flow 1's bottlenecked conflicts $\beta^{(n)}$ oscillates between $\{1\}$ and $\{2\}$. Despite the convergence of max-min fair solutions, the location of bottlenecked conflicts need not converge.

This example also demonstrates that even though the loss-less max-min solution is symmetric in its goodputs, the lossy max-min fair solutions need not be. With $c_1 = 1$, assume the link error rates are equal on all links ($\mathbf{q} = q\mathbf{1}$) and examine the behavior of the max-min fair solutions as $q \to 1$. From the solutions in Figure 2 it is clear that the max-min fair solutions are converging to the loss-less max-min fair solution as $q \uparrow 1$. The lossy solutions exhibit an asymmetry not seen in the loss-less case, whereby flow 3 is favored and gets more goodput.

It can be understood as follows: for the max-min fair solution the input rates are chosen so that flows 1 and 2 achieve the same goodput, but as Flow 1 experiences loss before sharing a link with flow 3, flow 3 can take up the additional left-over capacity and so gets higher goodput. This is a consequence of the location of the bottlenecked conflict for Flow 1.

Example 2, Failure of the last-link loss approximation:

With the topology in Figure 3, set $c_1 = c_2 = 1$. We consider the situation where losses occur on link 1, so that $q_1 = q$ and $q_2 = 1$. The last-link loss approximation effectively assumes that for Flow 1 $q_1 = 1$ and $q_2 = q$, while for Flow 2 $q_2 = 1$. For this network, the matrices in equation (1) are $B = I$ and $\mathbf{D} = \mathbf{1}$, the identity. The routing matrix for the real network, $A(\mathbf{q})$, and with last-link loss approximation, $A'(\mathbf{q})$ are defined by

$$A(\mathbf{q}) = \begin{pmatrix} 1/q & 0 \\ 1 & 1 \end{pmatrix} \text{ and } A'(\mathbf{q}) = \begin{pmatrix} 1/q & 0 \\ 1/q & 1 \end{pmatrix}.$$

Both the max-min fair and proportionally fair solutions can be determined explicitly for this network and they coincide. For the real network, with $A(\mathbf{q})$ the max-min fair and proportionally fair solution is $x_1 = \min(1/2, q)$ and $x_2 = \max(1/2, 1-q)$, while for the last-link loss approximation, with $A'(\mathbf{q})$, $x_1 = x_2 = q/(1+q)$. These solutions are plotted in Figure 4. It can be seen that the solutions converge to the loss-less ones when $q \to 1$, as anticipated by Theorem 1 and Corollary 1. However, the solutions diverge for $q < 1$. Indeed for $q < 1$, the rate region with the last-link loss approximation in force is smaller than the real rate region, leading to a smaller network goodput and highly divergent solutions when $q < 1/2$. This indicates a typical, simple situation in which the last-link loss approximation is inappropriate in the presence of non-zero losses.

Example 3, a WiFi network: Here we show how the model framework can be used to model a network where links have joint constraints, as conveyed by (4'). We consider a WiFi network where links are coupled through a shared wireless resource and discuss the well reported performance anomaly of wireless networks employing the IEEE 802.11 Distributed Co-ordination Function (DCF) [16]. When one station has a low link rate to the Access Point (AP), say, 1 Mbps, and others have a higher link rate to the AP, say, 11 Mbps, then the bandwidth allocation attained by the operation of the 802.11 DCF is such that all flows get less than 1 Mbps. We show that this anomaly arises as a consequence of that protocol enforcing max-min fairness and that MIMO gains can overcome it.

Consider the network depicted in Figure 5. Although each flow has its own link with its own loss rate, the links are
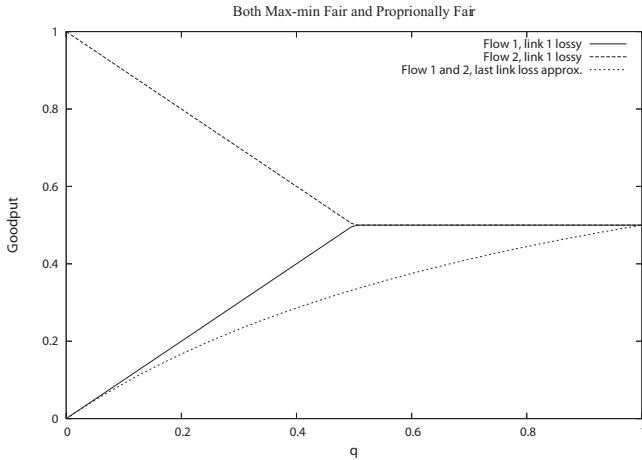
Fig. 4. Network in Figure 3. Difference in max-min fair and proportionally solutions based on loss at first link and last-link loss approximation.
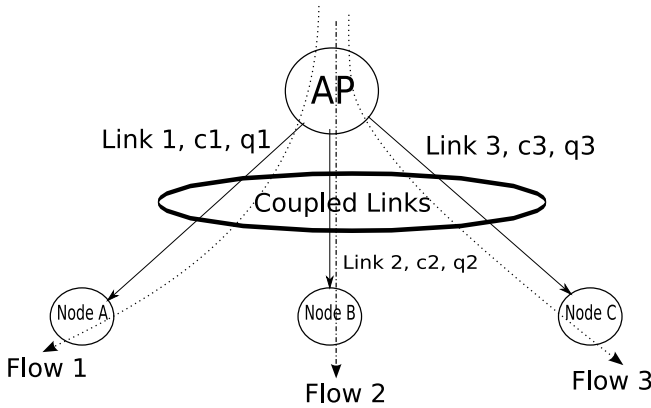

Fig. 5. WiFi network with three flows on three coupled links. Illustrates construction of conflict matrix and the enforcement of max-min fairness by IEEE 802.11

coupled by the IEEE 802.11 DCF [17]. As links share the wireless resource the Joint Conflict Matrix J has a single row with every entry equal to 1. The conflict and routing matrices are:

$$B = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \\ 1 & 1 & 1 \end{pmatrix} \text{ and } A(\mathbf{q}) = \begin{pmatrix} 1/q_1 & 0 & 0 \\ 0 & 1/q_2 & 0 \\ 0 & 0 & 1/q_3 \end{pmatrix}.$$

Consider the following degrees of freedom vector that corresponds to each station having a single receive antenna, but the access point having $d \in \{1, 2, \dots\}$ transmit antennas: $D^T = (1\,1\,1\,d)$.

With $q_1 = q_2 = q_3 = 1$ so that links are not lossy, $d = 1$ so that all stations only have one transmitter and receiver, and $C_1 = C_2 = 11$ Mbps and $C_3 = 1$ Mbps, the goodput rate region must satisfy $x_1/11 + x_2/11 + x_3 \le 1$. The max-min fair solution is readily identified to be $x_1 = x_2 = x_3 = 11/13$ Mbps. While the detailed operation of the IEEE 802.11 DCF is quite complex for unsaturated stations [18], at a high level it gives each station an approximately equal chance of winning the medium for a packet transmission so that it too would give rise to the same solution and to the reported 802.11 performance anomaly.

One solution to this anomaly is to use the TXOP functionality of the 802.11e protocol. Alternatively, it is sufficient to have $d = 2$, corresponding to a two-transmitter MIMO gain at the access point. The new constraints give $x_1/11 + x_2/11 + x_3 \le 2$, so that by the water-filling algorithm, $x_3 = 1$ Mbps and $x_1 = x_2 = 11/2$ Mbps. Thus the station with the low rate link does not throttle the goodput of the stations with the high rate links in the presence of MIMO gains.

V. CONCLUSIONS

We present a natural extension of the notions of utility fairness and max-min fairness from wired networks to their wireless counterpart. We prove the existence and uniqueness of solutions. We show that as loss rates converge to zero, fair solutions in systems with loss converge to the corresponding solution in the loss-less network. We extend the definition of bottlenecked links to bottlenecked conflicts and prove that max-min fairness can be defined in terms of these bottlenecked conflicts. Through examples, we demonstrate that even though max-min fair solutions converge as loss rates tend to zero, the location of bottlenecked conflicts may not.

APPENDIX

Proof of Theorem 1: Let $\{\mathbf{q}^{(k)}\}_{k=1}^{\infty}$ be such that $\lim_{k \to \infty} \mathbf{q}^{(k)} = \mathbf{1}$, which implies $\lim_{k \to \infty} \min_{e \in \vec{\mathcal{E}}} q_e^{(k)} = 1$. For every $0 < \epsilon < 1$ there exists a K_ϵ such that $1 \ge \min_{e \in \vec{\mathcal{E}}} q_e^{(k)} \ge (1 - \epsilon)^{1/|\vec{\mathcal{E}}|}$ for all $k \ge K_\epsilon$, which ensures that $1 \ge \min_{\mathcal{A} \in 2^{\vec{\mathcal{E}}} \setminus \emptyset} \prod_{e \in \mathcal{A} \subseteq \vec{\mathcal{E}}} q_e^{(k)} \ge 1 - \epsilon$ for all $k \ge K_\epsilon$. This implies that $\mathcal{X}(\mathbf{C}(1 - \epsilon), \mathbf{1}) \subseteq \mathcal{X}(\mathbf{C}, \mathbf{q}^{(k)}) \subseteq \mathcal{X}(\mathbf{C}, \mathbf{1})$ for all $k \ge K_\epsilon$.

The Pompeiu-Hausdorff distance between $\mathcal{X}(\mathbf{C}(1 - \epsilon), \mathbf{1})$ and $\mathcal{X}(\mathbf{C}, \mathbf{1})$ is (pg. 117 [15])

$$d_\infty(\mathcal{X}(\mathbf{C}(1 - \epsilon), \mathbf{1}), \mathcal{X}(\mathbf{C}, \mathbf{1}))$$
$$= \sup_{\mathbf{x} \in \mathcal{X}(\mathbf{C}, \mathbf{1}) \setminus \mathcal{X}(\mathbf{C}(1-\epsilon), \mathbf{1})} d_\infty(\mathcal{X}(\mathbf{C}(1 - \epsilon), \mathbf{1}), \{\mathbf{x}\})$$
$$\le \sup_{\mathbf{x} \in \mathcal{X}(\mathbf{C}, \mathbf{1})} d((1 - \epsilon)\mathbf{x}, \mathbf{x}) = \epsilon \sup_{\mathbf{x} \in \mathcal{X}(\mathbf{C}, \mathbf{1})} \|\mathbf{x}\|.$$

Thus for all $k \ge K_\epsilon$ we have

$$d_\infty(\mathcal{X}(\mathbf{C}, \mathbf{q}^{(k)}), \mathcal{X}(\mathbf{C}, \mathbf{1}))$$
$$\le d_\infty(\mathcal{X}(\mathbf{C}(1 - \epsilon), \mathbf{1}), \mathcal{X}(\mathbf{C}, \mathbf{1}))$$
$$\le \epsilon \sup_{\mathbf{x} \in \mathcal{X}(\mathbf{C}, \mathbf{1})} \|\mathbf{x}\|.$$

Therefore $\lim_{k \to \infty} d_\infty(\mathcal{X}(\mathbf{C}, \mathbf{q}^{(k)}), \mathcal{X}(\mathbf{C}, \mathbf{1})) = 0$ proving the first part of the theorem.

Define the indicator function of a convex set $D \in \Re^P$ to be $\delta(\mathbf{x}|D) = 0$ if $x \in D$ and $+\infty$ otherwise. We have shown that $\delta(\cdot | \mathcal{X}(\mathbf{C}, \mathbf{q}^{(k)}))$ epi-converges [15, Section 7.B] to $\delta(\cdot | \mathcal{X}(\mathbf{C}, \mathbf{1}))$. Let $\gamma(\mathbf{x})$ be a continuous, convex function that is level-bounded (i.e. $\{\gamma(\mathbf{x}) \le \eta\}$ is a bounded set for all $\eta \in \Re$). Then $\gamma(\mathbf{x}) + \delta(\mathbf{x}|\mathcal{X}(\mathbf{C}, \mathbf{q}^{(k)}))$ epi-converges to $\gamma(\mathbf{x}) + \delta(\mathbf{x}|\mathcal{X}(\mathbf{C}, \mathbf{1}))$. Since each of these functions are level-bounded from [15, Theorem 7.33] $\inf_{\mathbf{x}} \gamma(\mathbf{x}) + \delta(\mathbf{x}|\mathcal{X}(\mathbf{C}, \mathbf{q}^{(k)}))$ converges to $\inf_{\mathbf{x}} \gamma(\mathbf{x}) + \delta(\mathbf{x}|\mathcal{X}(\mathbf{C}, \mathbf{1}))$ and $\arg\inf_{\mathbf{x}} \gamma(\mathbf{x}) + \delta(\mathbf{x}|\mathcal{X}(\mathbf{C}, \mathbf{q}^{(k)}))$ converges to $\arg\inf_{\mathbf{x}} \gamma(\mathbf{x}) + \delta(\mathbf{x}|\mathcal{X}(\mathbf{C}, \mathbf{1}))$.

Defining $\gamma(\mathbf{x}) = -U(\mathbf{x})$ if $x_p \geq 0$ for each $p \in \mathcal{P}$ and $+\infty$ if $x_p < 0$ for any $p \in \mathcal{P}$, $\gamma(\mathbf{x})$ satisfies the conditions above, proving the second part of the theorem.

Proof of Corollary 1: The corollary follows from interchanging the order of limits.

Proof of Theorem 2: The "if" direction is proved by arriving at a contradiction. Suppose that $\mathbf{x} \in \mathcal{X}(\mathbf{C}, \mathbf{q})$ is max-min fair and assume that there exists a flow p with no bottlenecked conflict. It then follows that for each conflict $i \in \{1, \dots, E + \zeta\}$ such that $p \in i$ and $(BC^{-1}A(\mathbf{q})\mathbf{x})_i = \mathbf{D}_i$, there must exist a flow $p_i \neq p \in i$ such that $x_{p_i} > x_p$. Therefore for every conflict i along the route of flow p we can define the following non-zero quantity

$$\delta_i = \left(\frac{\sum_{e \in r(p_i), B_{i,e}=1} \frac{A_{e,p_i}}{c_e}}{\sum_{e \in r(p), B_{i,e}=1} \frac{A_{e,p}}{c_e}} \right) (x_{p_i} - x_p)$$

if $(BC^{-1}A(\mathbf{q})\mathbf{x})_i = \mathbf{D}_i$ and

$$\delta_i = (\mathbf{D}_i - (BC^{-1}A(\mathbf{q})\mathbf{x})_i) / \left(\sum_{e \in r(p), B_{i,e}=1} \frac{A_{e,p}}{c_e} \right)$$

otherwise. By increasing x_p by $\delta := \min_{i:x_p \in i} \delta_i$, the minimum over all conflicts involving p, and decreasing x_{p_i} by $x_{p_i} - x_p$ at every conflict j along the route of p such that $(BC^{-1}A(\mathbf{q})\mathbf{x})_j = \mathbf{D}_j$, we arrive at a new feasible goodput rate vector where we increase the rate of flow p without decreasing the rate of any flow p' with $x_{p'} \leq x_p$. This contradicts the max-min fairness of $\mathbf{x}$.

The proof in the "only if" direction follows directly from the definition of a bottlenecked conflict.

REFERENCES

[1] L. Tassiulas and S. Sarkar, "Maxmin fair scheduling in wireless ad hoc networks," *IEEE J. Select. Areas Commun.*, vol. 23, pp. 163-173, 2005.

[2] B. Li, "End-to-end fair bandwidth allocation in multi-hop wireless ad hoc networks," in *Proc. IEEE ICDCS*, 2005, pp. 471-840.

[3] Q. Dong, S. Banerjee, and B. Liu, "Throughput optimization and fair bandwidth allocation in multi-hop wireless LANs," in *Proc. IEEE INFOCOM*, Apr. 2006.

[4] L. Tan, X. Zhang, L. Andrew, S. Chan, and M. Zukerman, "Price-based max-min fair rate allocation in wireless multi-hop networks," *IEEE Commun. Lett.*, vol. 10, no. 1, pp. 31-33, 2006.

[5] A. Raniwala, P. De, S. Sharma, R. Krishnan, and T. Chiueh, "End-to-end flow fairness over IEEE 802.11-based wireless mesh networks," in *Proc. IEEE INFOCOM*, 2007, pp. 2361-2365.

[6] M. Wiczanowski, H. Boche, and S. Stanczak, "Unifying characterization of max-min fairness in wireless networks by graphs," in *Proc. IEEE ICC*, 2007, pp. 797-802.

[7] T. Nandagopal, T. Kim, X. Gao, and V. Bharghavan, "Achieving MAC layer fairness in wireless packet networks," in *ACM Mobicom*, Aug. 2000.

[8] S. Kunniyur and R. Srikant, "End-to-end congestion control: utility functions, random losses and ECN marks," *IEEE/ACM Trans. Networking*, vol. 11, no. 5, pp. 689-702, Oct. 2003.

[9] Q. Gao, J. Zhang, and S. Hanly, "Cross-layer rate control in wireless networks with lossy links: Leaky-pipe flow, effective network utility maximization and hop-by-hop algorithms," in *Proc. IEEE INFOCOMM*, 2008.

[10] D. P. Bertsekas and R. G. Gallager, *Data Networks*, 2nd ed. Prentice-Hall, Inc., 1992.

[11] F. P. Kelly, "Charging and rate control for elastic traffic," *European Trans. Telecommun.*, vol. 8, pp. 33-37, 1997.

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

[13] Béla Bollobás, *Modern Graph Theory*, vol. 184 of *Graduate Texts in Mathematics*. New York: Springer-Verlag, 1998.

[14] B. Radunovic and J.-Y. Le Boudec, "A unified framework for max-min and min-max fairness with applications," *IEEE/ACM Trans. Networking*, vol. 15, no. 5, pp. 1073-1083, 2007.

[15] R. Rockafellar and R. Wets, *Variational Analysis*. Springer-Verlag, 1998.

[16] M. Heusse, F. Rousseau, G. Berger-Sabbatel, and A. Duda, "Performance anomaly of 802.11b," in *Proc. IEEE INFOCOM*, 2003.

[17] IEEE, Wirless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications, IEEE std 802.11-1997 edition.

[18] D. Malone, K. Duffy, and D. J. Leith, "Modeling the 802.11 distributed coordination function in non-saturated heterogeneous conditions," *IEEE/ACM Trans. Networking*, vol. 15, no. 1, pp. 159-172, 2007.