

Existence and Uniqueness of Fair Rate Allocations in Lossy Wireless Networks

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

Abstract

There is considerable interest in extending established concepts of fair resource allocation in wired networks to wireless networks. To do this, 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 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. Through examples we illustrate distinctive features of fair solutions that arise when non-congestion based losses occur.

Index Terms

Resource management & QoS provisioning

I. INTRODUCTION

With the increased roll-out of mesh networks, there is considerable interest in achieving fair resource allocation in multi-hop wireless networks. Most current work focuses on identifying algorithms that enable the discovery of max-min fair solutions. A sample of work in this area includes [1]–[6]. Some of this work treats wireless-network specific features such as frequency reuse at non-interfering distances. Notably, however, the only works that we are aware of that explicitly treat the lossy nature of wireless networks, such as [7] and [8], do so by assuming that

Work supported by Science Foundation Ireland grant 07/IN.1/I901. The authors are with the Hamilton Institute, National University of Ireland, Maynooth.

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, so we refer to this as the *last-link loss approximation*. Here we consider the more general case where losses can occur at any (and every) link and this is reflected in the output of each link. We will show that the situation with losses on any link can yield significantly different fairness solutions to those given by the last-link loss approximation.

There are two distinct 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 here, is to retain the notion that the utility of a flow is function of its goodput at the receiver. That is, flows get 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 to treat losses enables us to quickly 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 and thus we recover the well established notions of fairness in the loss-less case. A by-product of our general 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 that it will typically 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. This explains why in proving the convergence of max-min fair solutions, their bottleneck link representation is inappropriate and instead we represent them as a limit of an appropriate sequence of utility fair solutions.

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 to investigate the effects of channel reuse at non-interfering distances. We illustrate new phenomena in the fair solutions that occur in wireless networks and show how the last-link loss approximation can give erroneous solutions in the

presence of lossy links that are not at the last link before the receiver.

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 [9], where the focus was originally on max-min fairness. The introduction of the notion of proportional fairness [10], 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 [11]. Fundamental assumptions in these 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.

While some of these assumptions are appropriate for wireless networks, others need to be reconsidered. In this work we adopt (1) and (2), but adapt (3) and (4) to these: (3') links are potentially lossy (for each link a proportion of every flow passing through it is lost); and (4') links may not be independent. (3') corresponds to losses due to environmental conditions. (4') includes multiple-access channels and primary interference constraints.

In the presence of losses, a flow's input to the network does not equal its network output. 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. We revisit utility fairness and max-min fairness where a flow's utility is a function of its goodput. We establish the existence and uniqueness of solutions. We demonstrate that the generalized treatment is reasonable by proving that, as loss rates go to zero, fair solutions converge to those in the equivalent loss-less system, even though the location of bottlenecked links need not converge.

III. EXISTENCE, UNIQUENESS AND CONVERGENCE OF FAIR SOLUTIONS

Existence and uniqueness of fair solutions. Adopting the notation in [12], 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 $\mathcal{P} = \{1, 2, \dots, P\}$ 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 $r(p)$ from $\vec{\mathcal{E}}$ connecting $s(p)$ to $d(p)$ without a cycle. Note that multiple routes are more likely to occur in wireless networks than in wired networks due to the possibility of multiple non-overlapping frequencies/channels being used in a single physical space. As there is no technical difficulty in having multiple flows taking distinct routes between a source and destination pair, this can be readily 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, we 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' we define $g(r', e) := \infty$. We 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 $g(r', e) \leq g(r', f)$ (so that e is closer to final node in the route r' than f) we 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 $E \times E$ diagonal matrix with zero entries off the diagonal and with diagonal entries $C_{e,e} = c_e$. Treating (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. Thus 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. [9], [11], [13]). The last-link loss approximation, as adopted in [7], [8], would be 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 these losses occur at the flow's last link and nowhere else on each flow's route through the network.

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 as well as interference and spatial reuse, we introduce a

conflict matrix B :

$$B = \begin{bmatrix} I \\ J \end{bmatrix},$$

where I is the $E \times E$ identity and J is a $\zeta \times E$ matrix with $\{0, 1\}$ entries where $\zeta \in \{0, \dots, 2^E - E - 1\}$, so that B is a $(E + \zeta) \times E$ matrix. The I matrix is the individual conflict matrix and gives each link its individual capacity constraint. The J matrix is the joint conflict matrix: if the links e_{i_1}, \dots, 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 or channel reuse at non-interfering distances, as will be illustrated later. 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* $i \in \{1, \dots, E + \zeta\}$ 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 channels at each conflict. Its use will become apparent through its use in the examples of Section IV.

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:

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

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 $\mathbf{D} \neq \mathbf{1}$ so that links have cross dependencies, receivers have additional degrees of freedom and $\mathbf{q} < \mathbf{1}$ (entry-wise) so that links are lossy, then the 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 $\mathcal{X}(\mathbf{C}, \mathbf{q}) \subset \mathbb{R}^P$. 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 for each link $e \in \vec{\mathcal{E}}$ and flow $p \in \mathcal{P}$ it is the case that $A(\mathbf{q})_{e,p} \leq A'(\mathbf{q})_{e,p}$. 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 $U : \mathbb{R}^P \mapsto [-\infty, \infty)$ of the goodput \mathbf{x} , the fair allocation is an optimizer in the solution of the following optimization problem:

$$\sup \{U(\mathbf{x}) : \mathbf{x} \in \mathcal{X}(\mathbf{C}, \mathbf{q})\}. \quad (2)$$

The following proposition follows from the linearity of the constraints in equation (1) [14].

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.*

This establishes the existence and uniqueness of utility fair solutions, but not max-min fair solutions.

Continuity of utility fair solutions. 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 $\mathbf{1}$. We prove this as a consequence of showing that a stronger property holds: a type of set convergence [14] for the regions $\mathcal{X}(\mathbf{C}, \mathbf{q})$. Define the Pompeiu-Hausdorff distance [14] between two non-empty closed sets $D, E \subset \mathbb{R}^P$ as $d_\infty(D, E) := \sup_{\mathbf{x} \in \mathbb{R}^P} |d_D(\mathbf{x}) - d_E(\mathbf{x})|$, 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 \mathbb{R}^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 solution. 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 $\lim_{k \rightarrow \infty} q_e^{(k)} = 1$ for each link $e \in \vec{\mathcal{E}}$. Then the rate regions in the lossy networks converge to the rate region in the corresponding loss-less network, $\lim_{k \rightarrow \infty} d_\infty(\mathcal{X}(\mathbf{C}, \mathbf{q}^{(k)}), \mathcal{X}(\mathbf{C}, \mathbf{1})) = 0$ and, consequently, utility fair solutions converge to the corresponding loss-less utility fair solutions.*

This theorem also holds with the last-link loss approximation in force. Thus, as loss rates con-

verge to zero, both the full lossy network utility fair solutions and the last-link loss approximation solutions converge to the corresponding loss-less solutions. Thus, when loss rates are sufficiently small, the last-link loss approximation is justifiable. 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 $x_p > \bar{x}_p \implies \exists o \in \mathcal{P} \setminus \{p\}$ such that $x_o < \bar{x}_o \leq \bar{x}_p$ [9]. That is, increasing some component \bar{x}_p must be at the expense of decreasing some already smaller or equal component \bar{x}_o . Max-min fairness does not correspond to the solution of (2) for any utility function, but is known to arise as the limit of utility fair solutions. To treat max-min fairness we begin by specializing to a particular class of utility functions. The (w, α) fair solution [11] 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 $\mathbf{w} = (w_1, \dots, w_P)^T$ and $\boldsymbol{\alpha} = (\alpha_1, \dots, \alpha_P)^T$. Lemma 3 in [11] proves that max-min fair solutions arise as the limiting solution as $\alpha \rightarrow \infty$. 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 $\alpha \rightarrow \infty$, giving the following lemma.

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 [13]. The following corollary 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 loss approximation in lieu of the full lossy links formulation.

This corollary is surprising as the original definition of max-min fairness [9] 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.

We have defined max-min fairness for the loss-less case in terms of the goodput rate vectors. In the loss-less case, max-min fairness can, equivalently, be defined in terms of bottlenecked links [9]: 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

$$\sum_{p \in \mathcal{P}: e \in r(p)} x_p = c_e \text{ and } x_p \geq x_{p'}$$

for all flows p' with link e along their routes. Then 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 link with respect to \mathbf{x} (e.g. [9] pg. 527). This alternate definition yields a procedure called the water-filling algorithm to identify the max-min fair solution. The water-filling 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 are increased further until another set of link constraints become active/bottlenecked; and 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 [13].

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 and links need not operate independently. Given a feasible goodput rate vector $\mathbf{x} \in \mathcal{X}(\mathbf{C}, \mathbf{q})$, we say that a conflict $i \in \{1, \dots, E + \zeta\}$ is a *bottlenecked conflict* with respect to \mathbf{x} for a flow p with at least one link in that conflict if

$$(BC^{-1}A(\mathbf{q})\mathbf{x})_i = \mathbf{D}_i \text{ and } x_p \geq x_{p'}$$

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} .*

We will demonstrate in Section IV that bottlenecked conflicts need not converge, even though max-min fair solutions do. However, Theorem 2 does give us a simple algorithm to identify max-min fair solutions in lossy networks: the obvious generalization of the water-filling procedure that is described above. Due to Theorem 2 and the coordinate convexity of the goodput rate region, the generalized water-filling algorithm necessarily identifies the unique lossy max-min fair solution.

IV. EXAMPLES

We illustrate the impact of the lossy nature of wireless networks on utility and max-min fairness through examples. Recalling equation (3), for simplicity we assume that all flows have the same $(w, \alpha) = (1, 1)$ utility function, commonly called proportional fairness. Max-min fair solutions were obtained by the water-filling procedure described in Section III.

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




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.

This example demonstrates that the set of bottlenecked conflicts for each flow need not converge as the loss rates converge to zero on all the links. It illustrates the result in Corollary 1, that despite the behavior of the bottlenecked conflicts, the max-min fair solutions do converge. It also highlights new asymmetries in the lossy max-min fair solution that are not present in the loss-less case.

Consider the basic three flow, two link network in Figure 1, where we assume that $c_2 = 1$. The loss-less version of this network has been used to illustrate max-min fairness [9] in wired networks. For this network, the matrices in equation (1) are $B = I$ and $\mathbf{D} = \mathbf{1}$, and

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

We first show that bottlenecked conflicts do not necessarily converge as loss rates converge to zero. Flow 1 traverses both links while Flow 2 only passes through link 1 and flow 3 only passes through link 2. 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 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

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

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 $\{1\}$. Thus even though the max-min fair solutions are converging by Corollary 1, the bottlenecked conflicts for Flow 1 are not. If instead the sequence of loss rates are $q_1^{(n)} = 1 - 1/n$ and $q_2^{(n)} = 1 - 3/n$ so that again, then $\beta^{(n)} = \{2\}$ for all n , so that again Flow 1's bottlenecked conflicts do not converge. If the sequence of loss rates were such that 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 $\{2\}$ and again it does not converge to the loss-less bottlenecked conflicts $\{1, 2\}$. Thus, despite the convergence of max-min fair solutions, the location of bottlenecked conflicts need not converge.

We now consider the continuity property proved in Corollary 1: that max-min fair solutions converge as the loss rates tends to zero. This example also demonstrates that even though the loss-less max-min solution is symmetric in its goodputs, the lossy max-min fair solutions are not. With $c_1 = 1$, we now 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 $q \rightarrow 1$. From the solutions in Figure 2 it is clear that




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

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. While this may appear unintuitive at first, 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:




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

This example illustrates the differences between max-min and proportionally fair ($w = 1, \alpha = 1$) solutions based on full network dynamics and the last-link loss approximation. Set $c_1 = c_2 = 1$, so that both links have the same capacity. Flow 1 passes through both links in the network and shares link two with Flow 2. 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 $D = \mathbf{1}$, the identity. For the real network, the routing matrix is

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

Irrespective of where the losses occur in the real network, the routing matrix given by the the last-link loss approximation is:

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

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

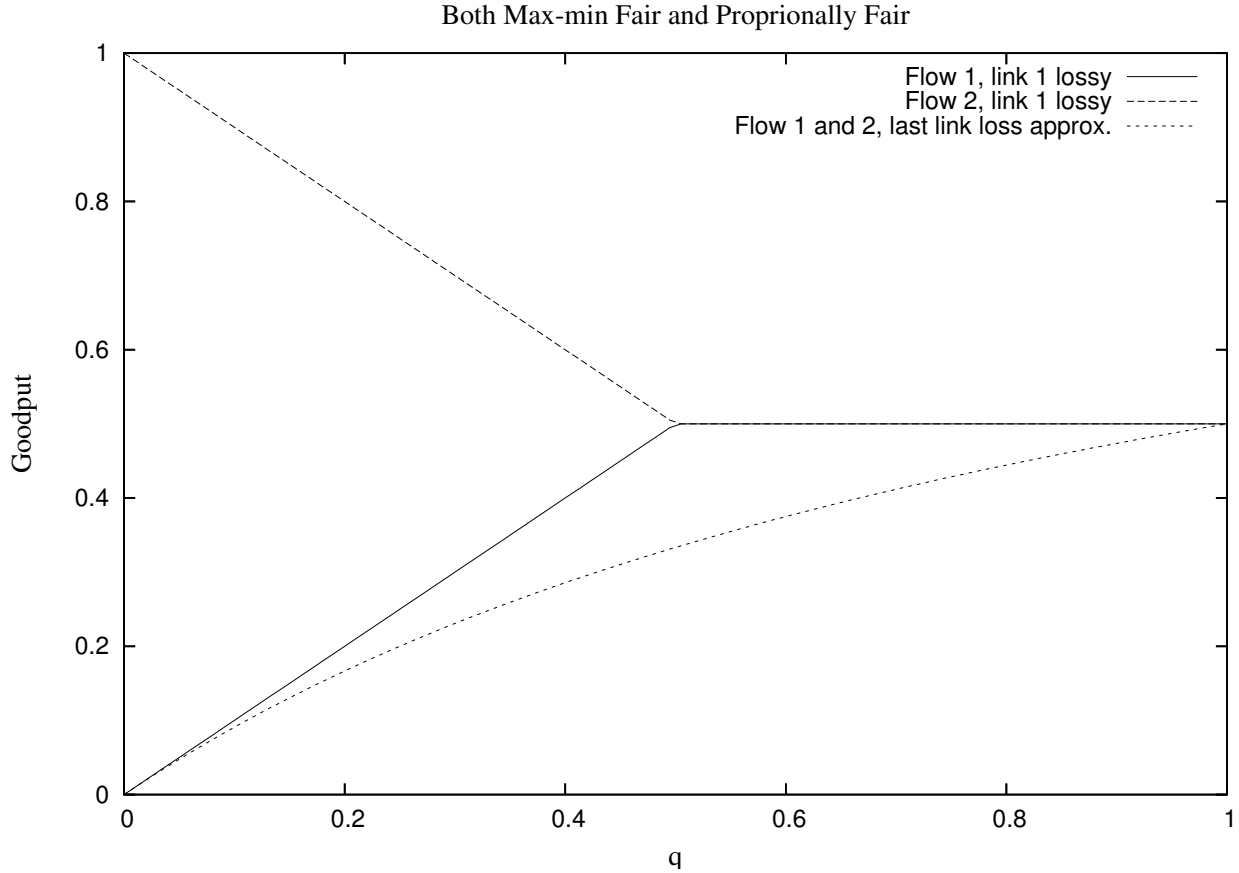


Fig. 4. Network in Figure 3. Difference in max-min fair and proportionally solutions based on loss at first link and last-link loss approximation.

fair solution is

$$x_1 = \min\left(\frac{1}{2}, q\right) \text{ and } x_2 = \max\left(\frac{1}{2}, 1 - q\right),$$

while for the last-link loss approximation, with $A'(q)$,

$$x_1 = x_2 = \frac{q}{1 + q}.$$

These solutions are plotted in Figure 4. It can be seen that the solutions converge to the loss-less ones when $q \rightarrow 1$, as anticipated by Theorem 1 and Corollary 1). However, the solutions diverging 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, Total Loss:

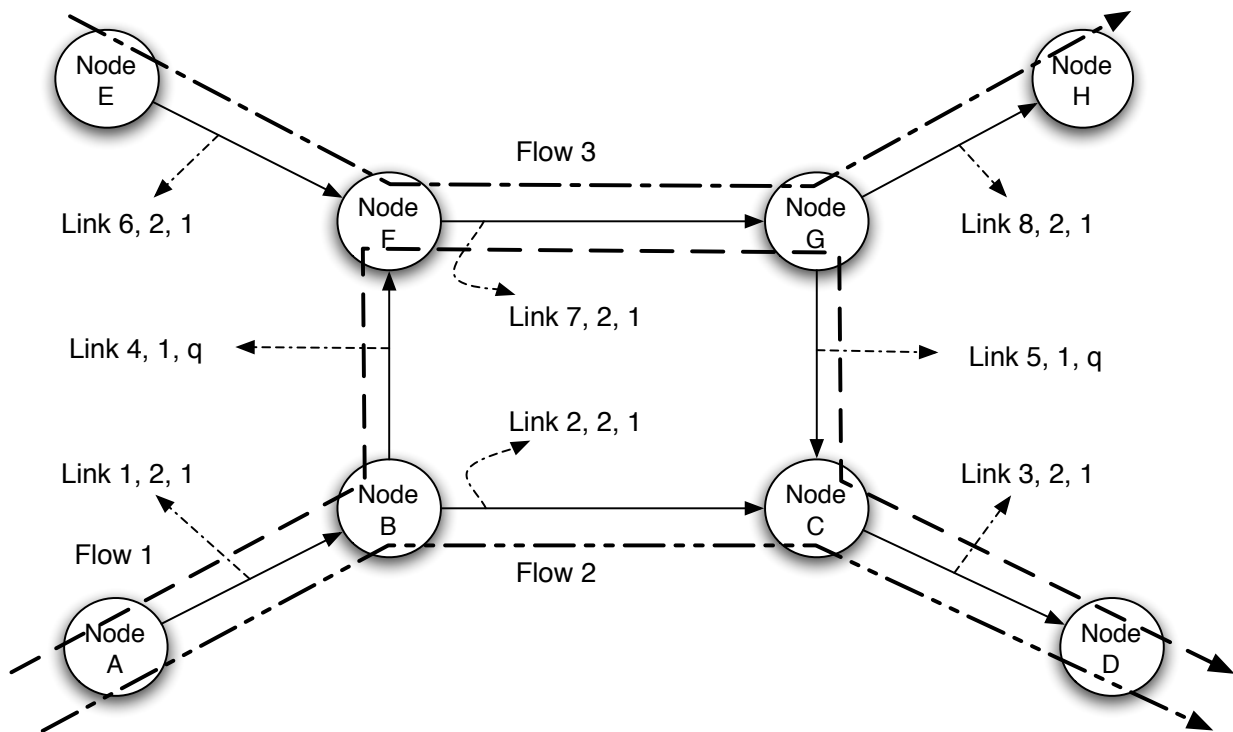


Fig. 5. Three flow, eight link network example to illustrate increased losses in max-min fair solution over proportional fair solution. The pair of numbers on each link indicate its capacity and the proportion of traffic not lost on it.

In all the other examples presented, the max-min fair and proportional fair solutions exhibit equal absolute loss. This example illustrates that it is possible that the max-min fair solution experiences significantly more absolute loss than the proportional fair solution.

Consider the 3 flow network depicted in Figure 5. For this network, the matrices in equation

(1) are $B = I$ and $\mathbf{D} = \mathbf{1}$, $c_1 = c_2 = c_3 = c_6 = c_7 = c_8 = 2$, $c_4 = c_5 = 1$ and

$$A(\mathbf{q}) = \begin{pmatrix} 1/q^2 & 1 & 0 \\ 0 & 1 & 0 \\ 1 & 1 & 0 \\ 1/q^2 & 0 & 0 \\ 1/q & 0 & 0 \\ 0 & 0 & 1 \\ 1/q & 0 & 1 \\ 0 & 0 & 1 \end{pmatrix}.$$

Flows 2 and 3 have routes with no losses, while Flow 1, which share links with both, has a lossy route.

For this network, the max-min fair and proportional fair solutions can be determined as explicit functions of q . For max-min fair, we have

$$x_1 = q^2, x_2 = 1 \text{ and } x_3 = 2 - q,$$

while for proportional fairness,

$$x_1 \frac{2q^2}{1 + \sqrt{q} + q}, x_2 = \frac{2(q + \sqrt{q})}{1 + \sqrt{q} + q}$$

and $x_3 = \frac{2(1 + \sqrt{q})}{1 + \sqrt{q} + q}.$

These solutions are plotted in Figure 6 where it can be seen that the ordering of goodputs of each flow between the max-min fair and proportional fair solution is dependent on q . Note also that the network goodput for the max-min fair solution is $3 - q + q^2$, while for the proportional fair solution the network goodput is $2(q^2 + (1 + \sqrt{q})^2)/(1 + q + \sqrt{q})$. Thus there is no fixed ordering of the network goodputs, which is dependent on q . Figure 7 plots the network loss experienced by each solution. Apart from for $q = 1$, the total loss for the max-min fair solution is significantly higher than with the proportional fair solution.

Example 4, a WiFi network:

The previous three examples are designed to illustrate features that appear in the lossy solutions, but not in the loss-less solutions. Here we show how the model framework can be used to model a network where links do not all operate independently, as conveyed by (4') in our revised hypotheses.




Fig. 6. Network in Figure 5. Goodput for each flow in proportional fair and max-min fair solutions as a function of q

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) [15]. An example of this performance anomaly is the following: 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 will show that this “anomaly” arises as a consequence of that protocol enforcing max-min fairness. We will also show that MIMO gains can overcome this so-called anomaly.

Consider the network depicted in Figure 8. Although each flow has its own link with its own loss rate, the links are coupled by the WiFi Medium Access Controller, typically the IEEE 802.11 DCF [16]. As the links share the wireless resource the Joint Conflict Matrix J has a single row




Fig. 7. Network in Figure 5. Absolute loss in max-min fair and proportional fair solution as a function of q .

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:

$$D = \begin{pmatrix} 1 \\ 1 \\ 1 \\ d \end{pmatrix}.$$

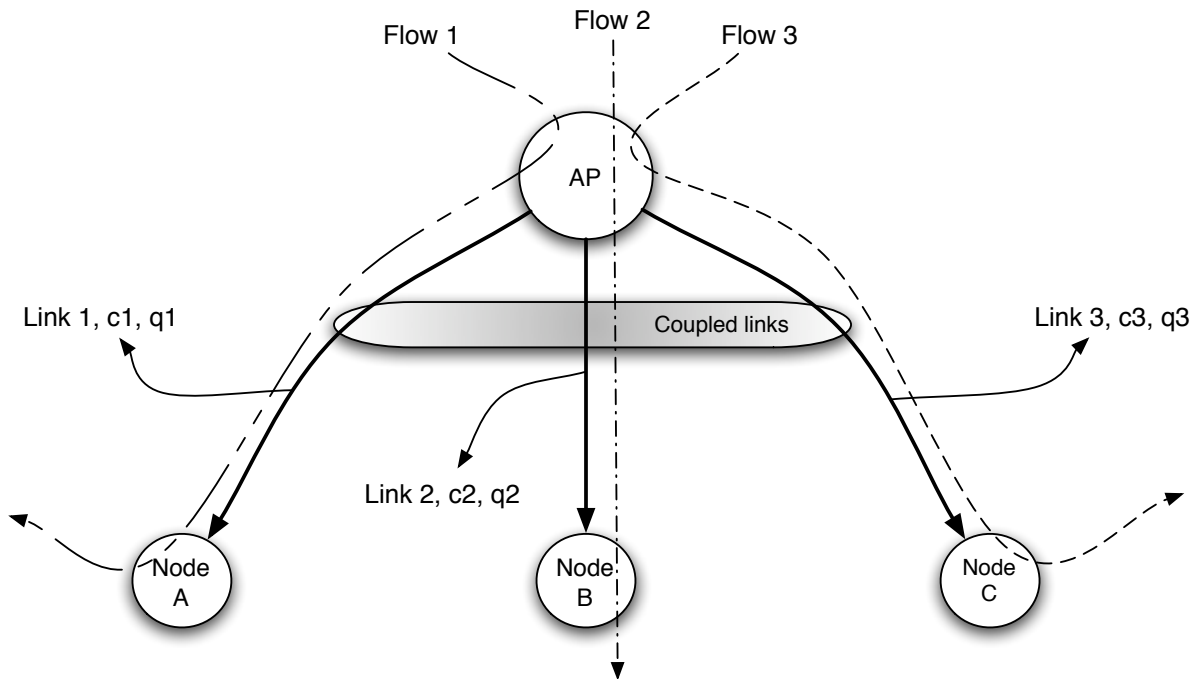


Fig. 8. WiFi network with three flows on three coupled links. Illustrates construction of conflict matrix and the enforcement of max-min fairness by IEEE 802.11

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 $x_1/11 + x_2/11 + x_3 \leq 1$. 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 [17], 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.

The typical solution to this anomaly is to use the TXOP functionality of the revised 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 $x_1/11 + x_2/11 + x_3 \leq 2$, 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.

Note that if the low rate link is lossy, the performance anomaly is even more pronounced: with $d = 1$, if $q_1 = q_2 = 1$, but $q_3 < 1$, then $x_1 = x_2 = x_3 = 11q_3/(2q_3 + 11)$ Mbps. Thus additional MIMO gains would be necessary in this case to overcome the anomaly.

Example 5, Spatial Reuse:




Fig. 9. Two flow, three link network example that illustrates the inclusion of spatial reuse of frequencies

A key feature of wireless mesh networks is channel/frequency reuse at non-interfering distances. In this example we illustrate how spatial reuse can be modeled through an appropriate conflict matrix.

Consider three interference scenarios for the network depicted in Figure 9: (a) all links can operate simultaneously without interfering; (b) neither links 1 and 2 nor links 2 and 3 can operate simultaneously, but links 1 and 3 can; and (c) only one of links 1, 2 and 3 can operate at any given instance. These three scenarios consider increasing levels of radio interference.

For scenario (a), no joint conflicts occur and the conflict matrix $B = I$. For scenario (b), the network structure gives rise to conflicts between the links which can be incorporated into by having a two-row joint conflict matrix J , one of which couple links 1 and 2, the other couples links 2 and 3:

$$B = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \\ 1 & 1 & 0 \\ 0 & 1 & 1 \end{pmatrix}.$$

For scenario (c), we have a single row joint conflict matrix that couples the operation of links 1,

2 and 3:

$$B = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \\ 1 & 1 & 1 \end{pmatrix}.$$

Assuming no MIMO gains, we have that D is the $E + \zeta$ column vectors, all of whose entries are 1, while

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

Choosing $c_1 = c_2 = c_3 = 1$ and $q_1 = q_2 = q_3 = q$, we consider the impact of these scenarios on the proportionally fair and max-min fair solutions.

For this network, the max-min fair solutions can be determined explicitly. In scenario (a) it is

$$x_1 = \begin{cases} q^3 & \text{if } q \leq (\sqrt{5} - 1)/2 \\ q^2(1 - q)/(1 - q^2) & \text{otherwise} \end{cases}$$

and

$$x_2 = \begin{cases} q(1 - q) & \text{if } q \leq (\sqrt{5} - 1)/2 \\ q^2(1 - q)/(1 - q^2) & \text{otherwise.} \end{cases}$$

In scenario (b), with spatial reuse, they become $x_1 = x_2 = q^3/(1 + q + q^2)$. While in scenario (c) with all links interfering, they are $x_1 = x_2 = q^3/(1 + q + 2q^2)$. Each reduction of spatial reuse due to interference reduces each flow's goodput further. These max-min fair solutions are plotted in Figure 10. When there is interference between links, the goodput of Flow 1 is less than half of that in scenario (a). The throughput of Flow 2 also suffers greatly due to the lack of channel reuse.

The proportionally fair solutions can also be calculated explicitly as functions of q . For scenario (a) they are

$$x_1 = \begin{cases} q^3 & \text{if } q \leq 1/2 \\ q^2/2 & \text{otherwise} \end{cases}$$




Fig. 10. Network in Figure 3. Demonstrates impact of spatial reuse on max-min fair solutions.

and

$$x_2 = \begin{cases} q(1-q) & \text{if } q \leq 1/2 \\ q/2 & \text{otherwise.} \end{cases}$$

For scenario (b), they become $x_1 = q^3/(2(1+q))$ and $x_2 = q/2$. For scenario (c), they are $x_1 = q^3(1-q)/(2(1-q^3))$ and $x_2 = q/2$. These are plotted in Figure 11. Unlike the max-min fair solution, the proportionally fair solution protects Flow 2 at the expense of Flow 1. Its goodput never drops below $q/2$, even due to interference constraints.

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.




Fig. 11. Network in Figure 3. Demonstrates impact of spatial reuse on proportionally fair 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.

We demonstrate that due to losses, new asymmetries can arise in fair solutions in wireless networks. We illustrate how the last link loss approximation can lead to incorrect utility fair and max-min fair solutions. We show that max-min fair solutions can experience greater absolute loss than their proportionally fair counterpart. We illustrate how the framework encompasses lossy links, shared channels, frequency reuse at non-interfering distances and other features that are characteristic of wireless networks.

APPENDIX

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

The Pompeiu-Hausdorff distance between $\mathcal{X}(\mathbf{C}(1 - \epsilon), \mathbf{1})$ and $\mathcal{X}(\mathbf{C}, \mathbf{1})$ is (pg. 117 [14])

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

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

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

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

Define the indicator function of a convex set $D \in \mathfrak{R}^P$ 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 [14, 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. $\{\gamma(\mathbf{x}) \leq \eta\}$ is a bounded set for all $\eta \in \mathfrak{R}$). 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 [14, 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 $x < 0$, $\gamma(\mathbf{x})$ 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 first part 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 $p \neq p_i \in i$ 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,

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

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 reverse 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 JSAC*, vol. 23, no. 1, pp. 163–173, 2005.
- [2] B. Li, "End-to-end fair bandwidth allocation in multi-hop wireless ad hoc networks," in *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 *IEEE INFOCOM*, April 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 Comms. Letters*, 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 *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 *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. Network.*, vol. 11, no. 5, pp. 689–702, October 2003.
- [9] D.P. Bertsekas and R.G. Gallager, *Data networks (2nd ed.)*, Prentice-Hall, Inc., 1992.
- [10] F. P. Kelly, "Charging and rate control for elastic traffic," *European Trans. on Telecommunications*, vol. 8, pp. 33–37, 1997.
- [11] 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.
- [12] Béla Bollobás, *Modern graph theory*, vol. 184 of *Graduate Texts in Mathematics*, Springer-Verlag, New York, 1998.

- [13] 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.
- [14] R. Rockafellar and R. Wets, *Variational analysis*, Springer-Verlag, 1998.
- [15] M. Heusse, F. Rousseau, G. Berger-Sabbatel, and A. Duda, "Performance anomaly of 802.11b," in *IEEE INFOCOM*, 2003, pp. 836 – 843.
- [16] IEEE, *Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications*, IEEE std 802.11-1997 edition, 1997.
- [17] D. Malone, K. Duffy, and D. J. Leith, "Modeling the 802.11 Distributed Coordination Function in non-saturated heterogeneous conditions," *IEEE/ACM Trans. Network.*, vol. 15, no. 1, pp. 159–172, 2007.