Some Results on Max-min Fairness for Resource Sharing

Ye Xia

***Abstract*— The max-min fairness is one of the most widely used fairness concepts for resource sharing in communication networks. This paper describes some results related to it.**

I. Introduction

The max-min fairness is one of the earliest, most widely used fairness concepts in communication networks and in several other areas of communications [1]–[3]. There is a vast literature that proposes to allocate network or other communication resources according to this fairness criterion when the resources are shared by multiple users. It is also related to more recent fairness concepts such as the proportional fairness [2]. However, many results about the max-min fairness are scattered in the literature, often without proper proofs. The goal here is to describe some basic results about the max-min fairness rigorously and to correct some misconceptions.

II. Definition and Basic Properties

Let $\mathbb{R}^n_+$ denote the set of nonnegative vectors in $\mathbb{R}^n$. Throughout, we consider only nonnegative vectors, although most of the results hold without this requirement. For a vector $x \in \mathbb{R}^n$, let x_i denote its i^{th} component, for $1 \le i \le n$. That is, $x = (x_1, x_2, \dots, x_n)$. Given a non-empty set $S \subseteq \mathbb{R}^n_+$, a fairness concept supplies a way of designating some vector as the "best" one in S.

Definition 1: A vector $x \in S$ is *max-min fair* (in the set S) if for any other vector $y \in S$, $x_j < y_j$ implies that there is some i, $1 \le i \le n$, $i \neq j$, such that $x_i \le x_j$ and $x_i > y_i$.

The interpretation of this definition is as follows. A vector $x \in S$ is max-min fair if one cannot increase one of its components without decreasing another of its components that is already smaller or equal, while remaining in S. The concept of max-min fairness is well defined by the following uniqueness result.

Lemma 1: When exists, the max-min fairness is unique.

Proof: We will show that $x_j = y_j$ must be true for all j. Suppose not. Without loss of generality, let's assume $x_j < y_j$ for some j. By the definition of the max-min fairness, there must exist an index i, $i \neq j$, such that $y_i < x_i \le x_j$. There must be another index k, $k \neq i$, such that $x_k < y_k \le y_i$. Hence, given $x_j < y_j$, there is k such that $x_k < y_k < x_j < y_j$. If $k = j$, a contradiction occurs. If $k \neq j$, we can continue the same process, which must eventually stop. When it stops, we will have a similar contradiction as above. ■

The notion of max-min fairness is related to the concept of lexicographical order, which is a partial order defined on $\mathbb{R}^n$. A vector $x \in \mathbb{R}^n$ is said to be *lexicographically smaller* than $y \in \mathbb{R}^n$, denoted by $x \preceq^{lex} y$, if either $x = y$ or there exists a k, $1 \le k \le n$, such that $x_k < y_k$ and $x_j = y_j$ for $j < k$.

Given a vector $x = (x_1, x_2, \dots, x_n)$, let $T(x) = (x_{(1)}, x_{(2)}, \dots, x_{(n)})$, where $x_{(i)}$ is the i^{th} order statistics of x. That is, $x_{(1)} \le x_{(2)} \le \dots \le x_{(n)}$. We can define an equivalence relation by $x \sim y$ if $T(x) = T(y)$ for any $x, y \in S$. For each $x \in S$, let $\bar{x}$ denote its equivalent class. Let $\bar{S}$ be the set of equivalent classes in S. That is, $\bar{x} \in \bar{S}$ if and only if some permutation of x is in S. Then, lexicographically smaller defines a total order on $\bar{S}$.

Definition 2: Let $\bar{x}, \bar{y} \in \bar{S}$. We say $\bar{x}$ is lexicographically smaller than $\bar{y}$ if $T(x) \preceq^{lex} T(y)$.

From now on, when without ambiguity, we write $x \preceq^{lex} y$ to mean $T(x) \preceq^{lex} T(y)$, and x being lexicographically smaller than y means $T(x) \preceq^{lex} T(y)$.

Definition 3: A vector $x \in S$ is *lexicographically fair* (in the set S) if for any other vector $y \in S$, $y \preceq^{lex} x$. That is, x is the maximal (also the greatest) element in S with respect to $\preceq^{lex}$.

The following theorem is stated without a proof in [5].

Therorem 2: If $S \in \mathbb{R}^n$ is a non-empty compact set, then a lexicographically fair vector exists.

Proof: Let π be a permutation on the set $\{1, 2, \dots, n\}$. Let $S_\pi = \{y \in S | y_{\pi(1)} \le y_{\pi(2)} \le \dots \le y_{\pi(n)}\}$. Note that the union of S_π's over all possible permutations is S. Hence, at least one of the S_π's is non-empty. Furthermore, each S_π is compact. For each fixed π, let x^π be the maximal element in S_π if the set is non-empty. For a non-empty S_π, x^π exists by the following construction. Let $x_{\pi(1)} = \max\{y_{\pi(1)} | (y_1, y_2, \dots, y_n) \in S_\pi\}$. Since S_π is compact and non-empty, such $x_{\pi(1)}$ exists. Next, we fix the component $y_{\pi(1)} = x_{\pi(1)}$ and maximize the component $y_{\pi(2)}$. This gives $x_{\pi(2)}$. That is, let $x_{\pi(2)} = \max\{y_{\pi(2)} | y_1, \dots, y_{\pi(1)-1}, x_{\pi(1)}, y_{\pi(1)+1}, \dots, y_n) \in S_\pi\}$. In general, once $x_{\pi(1)}$ to $x_{\pi(i-1)}$ are found, we fix $y_{\pi(1)} = x_{\pi(1)}, \dots, y_{\pi(i-1)} = x_{\pi(i-1)}$ and maximize the component $y_{\pi(i)}$. This gives $x_{\pi(i)}$. The final vector is $x^\pi = (x_1, \dots, x_n) \in S_\pi$. Then, the maximal element in the set $\{x^\pi | \pi \text{ is a permutation of } 1, 2, \dots, n\}$ is in S and it is lexicographically fair. ■

The procedure in the above proof is conceptually simple but inefficient to carry out. The Max-Min programming (MP) in [4] is another constructive way of finding the lexicographically fair vector (rather than the max-min fair vector). The rough idea of MP is as follows. In the first step, it solves the following optimization problem, $\max \min_i x_i$ subject to $(x_1, \dots, x_n) \in S$. Let the result be t_1 and suppose the component that achieves the optimum is unique and has the index $\sigma(1)$. The next optimization problem is $\max \min_{i \neq \sigma(1)} x_i$ subject to $x_{\sigma(1)} = t_1$ and $(x_1, \dots, x_n) \in S$. This process continues. The final result $x^\sigma = (x_1, \dots, x_n)$ is one lexico-

The author is with the Computer and Information Science and Engineering Department, University of Florida, 301 CSE Building, P. O. Box 116120, Gainesville, FL 32611-6120. Email: yx1@cise.ufl.edu.

graphically fair vector from its equivalent class. The existence of such a vector is guaranteed because each optimization problem is on a compact set.

It is assumed that, in each of the above sequence of optimization problems, the component that achieves the respective optimum is unique. But, this is not generally the case. The computation becomes much more complicated when the uniqueness condition is not satisfied. In fact, the MP reported in [4] does not handle the non-unique cases correctly. A correct procedure is to envision a tree where each node corresponds to an optimization problem. A node w in the tree has an index set $I(w) \subseteq N$, where $N = \{1, 2, \ldots, n\}$, and a vector $t(w) \in \mathbb{R}^n$. The node corresponds to the problem $\max \min_{j \in N \setminus I(w)} x_j$ subject to $x \in S$ and $x_k = t_k(w)$ for $\forall k \in I(w)$. Let the optimal value be $y(w)$ and the index set that achieves the optimal value be $\Sigma(w)$. For each $i \in \Sigma(w)$, there is a child node of w, denoted by w_i. Record at w_i the index set $I(w_i) = I(w) \cup \{i\}$ and the vector $t(w_i) = t(w)$ with a change at the i^{th} component $t_i(w_i) = y(w)$. To generate such a tree, we initialize the root node r with $I(r) = \emptyset$ and an arbitrary vector $t(r) \in \mathbb{R}^n$. Then, we create children for the nodes until no more nodes can be created, which occurs when all the leaf nodes has an index set $I(w) = N$. The maximal element in S is the maximal element in the finite set $\{t(w)\}$ over all leaf nodes w.

For a vector $x \in S$, let $(l_i)_{i=1}^n$ be a permutation of the index set $\{1, \ldots, n\}$ satisfying $x_{l_1} \leq x_{l_2} \leq \ldots \leq x_{l_n}$. Such an index sequence is called an *increasing index sequence* for x, which may not be unique. The next lemma can be viewed as an equivalent definition of max-min fairness. It is somewhat easier to conceptualize than the original definition.

Lemma 3: A vector $x \in S$ is max-min fair if and only if, for any other vector $y \in S, y \neq x$, there exists an increasing index sequence for x, denoted by $(l_i)_{i=1}^n$, such that $x_{l_i} > y_{l_i}$, where $i = \text{argmin}_{1 \leq j \leq n} x_{l_j} \neq y_{l_j}$.

Proof: Suppose x is max-min fair. We choose an increasing index sequence $(l_i)_{i=1}^n$ that satisfies the following property. If $x_{l_j} = x_{l_{j+1}}$, then $y_{l_j} \leq y_{l_{j+1}}$. That is, when some components of x are identical, not only they are ordered next to each other, but the corresponding components of y are ordered in non-decreasing order.

Suppose $x_{l_i} < y_{l_i}$, where i is as in the statement of the lemma. Since x is max-min fair, there exists some $k \neq l_i$ such that $y_k < x_k \leq x_{l_i} < y_{l_i}$. Suppose $x_k = x_{l_i}$. Then k should be before l_i in the sequence $(l_j)_{j=1}^n$ since $y_k < y_{l_i}$. That is, $k = l_j$ for some $1 \leq j \leq i - 1$. But, for such k, $x_k = x_{l_j} = y_{l_j} = y_k$, which contradicts $y_k < x_k$.

Next, suppose $x_k < x_{l_i}$. Then again, $k = l_j$ for some $1 \leq j \leq i - 1$. But, for such k, $x_k = y_k$, which contradicts $y_k < x_k$. Hence, it must be that $x_{l_i} > y_{l_i}$.

To show the converse, let $(l_i)_{i=1}^n$ be an arbitrary increasing index sequence for x with the property $x_{l_i} > y_{l_i}$, where $i = \text{argmin}_{1 \leq j \leq n} x_{l_j} \neq y_{l_j}$. Then, for any k such that $x_k < y_k$, if such a k exists, it must be true that $k \neq l_j$, $1 \leq j \leq i$. Hence, $x_{l_i} \leq x_k$, and $x_{l_i} > y_{l_i}$. Hence, x is max-min fair. ∎

The next theorem is stated in several sources (e.g., [5]). We supply a proof.

Therorem 4: A max-min fair vector, when exists, is lexicographically fair.

Proof: Let $x \in S$ be max-min fair. Suppose $y \in S, y \neq x$, is another vector in S. Let $(l_i)_{i=1}^n$ be an increasing index sequence for x as in Lemma 3. Let $i = \text{argmin}_{1 \leq j \leq n} x_{l_j} \neq y_{l_j}$. Then, by Lemma 3, $x_{l_i} > y_{l_i}$.

Since $x_{l_j} = y_{l_j}$ for $j = 1, \ldots, i-1$, we have $y_{l_1} \leq y_{l_2} \leq y_{l_{i-1}}$. Then, $y_{l_j} \geq y_{(j)}$ for $j = 1, \ldots, i-1$. Then, we have $x_{(j)} \geq y_{(j)}$ for $j = 1, \ldots, i-1$, where $y_{(j)}$ is the j^{th} smallest component of y. If $x_{(j)} > y_{(j)}$ for some j, $1 \leq j \leq i-1$, then $y \preceq^{lex} x$. Next, consider the case $x_{(j)} = y_{(j)} = y_{l_j}$ for all j, $1 \leq j \leq i-1$. In this case, it must be true that $y_{l_i} \geq y_{(i)}$, because $y_{(i)} = \min\{y_k | 1 \leq k \leq n, k \neq l_1, k \neq l_2, \ldots, k \neq l_{i-1}\}$ but $l_i \neq l_j$ for $j = 1, \ldots, i-1$. Then, $x_{(i)} = x_{l_i} > y_{l_i} \geq y_{(i)}$. Hence, $y \preceq^{lex} x$. ∎

Therorem 5: Suppose $x \in S$ is max-min fair, and hence, lexicographically fair. Let y be any permutation of x and $y \neq x$. Then, $y \notin S$. Therefore, a max-min fair vector, when exists, is the only lexicographically fair vector for the set.

Proof: Let $(l_i)_{i=1}^n$ be an arbitrary increasing index sequence for x. Let $i = \text{argmin}_{1 \leq j \leq n} x_{l_j} \neq y_{l_j}$. Then, $y_{l_i} \geq x_{(j)} = x_{l_j}$ for $1 \leq j \leq i-1$. Hence, $y_{l_i} \geq x_{(i)} = x_{l_i}$. Since equality is not possible by assumption, it must be that $x_{l_i} < y_{l_i}$. By Lemma 3, x cannot be max-min fair, which contradicts the hypothesis of the theorem. ∎

Hence, when a max-min fair vector exists, there is no distinction between the max-min fairness and the lexicographical fairness. In such a situation, much understanding about the max-min fairness, which is a bit difficult to work with, can be achieved by investigating the properties of the lexicographical fairness. For instance, the fact that MP is considered as a generic way of finding the max-min fair vector is because it finds the lexicographically fair vector.

Besides the uniqueness, the max-min fairness has few properties that the lexicographical fairness does not have (But, there is a unique equivalent class that is lexicographically fair.). In contrast, the lexicographical fairness has many more interesting properties that the max-min fairness does not have. Furthermore, the lexicographical fairness is easier to work with. For instance, the condition for its existence is less demanding and often easier to verify. A max-min fair vector may not exist in many compact sets but a lexicographically fair vector always exists in a non-empty compact set. Furthermore, for evaluating the "fairness" of the distribution of a vector's components, the order of these components does not matter. Hence, the lexicographical fairness really captures everything that the max-min fairness embodies. The extra condition required by the definition of the max-min fairness adds little but is only constraining.

Therorem 6: Suppose S is a non-empty compact and convex set. Suppose $x \in S$ is lexicographically fair. Then, x is max-min fair.

Proof: Suppose x is not max-min fair. Then, by Lemma 3 there exists $y \in S$ and $y \neq x$ such that for any increasing index sequence $(l_i)_{i=1}^n$, we have $x_{l_i} < y_{l_i}$, where $i = \text{argmin}_{1 \leq j \leq n} x_{l_j} \neq y_{l_j}$. We choose an increasing index sequence for x that satisfies the following property. If $x_{l_j} = x_{l_{j+1}}$, then $y_{l_j} \leq y_{l_{j+1}}$. Let $z = \alpha x + (1-\alpha) y$ for some $0 < \alpha < 1$. We will choose a particular α such that $z \in S$ and $x \preceq^{lex} z$. This contradicts the assumption that x is lexicographically fair. Hence, x must be max-min fair.

We will show how to choose such an α. Suppose $x_{l_i} = x_{l_{i+1}} = \ldots = x_{l_{i+r}}$, and $x_{l_{i+r+1}} > x_{l_i}$, if $i + r + 1 \leq n$. By the assumption on $(l_i)_{i=1}^n$, $y_{l_i} \leq y_{l_{i+1}} \leq \ldots \leq y_{l_{i+r}}$. Let $\alpha^* \in (0, 1)$ satisfy $\alpha^* x_{l_j} > x_{l_i}$ for $i + r < j \leq n$. Define $z = \alpha^* x + (1 - \alpha^*)y$. By convexity of S, $z \in S$. Furthermore, $z_{l_j} = x_{l_j}$ for $1 \leq j < i$; $z_j = \alpha^* x_{l_j} + (1 - \alpha^*)y_{l_j} \geq \alpha^* x_{l_i} + (1 - \alpha^*)y_{l_i} > x_{l_i}$ for $i \leq j \leq i + r$; and $z_j = \alpha^* x_{l_j} + (1-\alpha^*)y_{l_j} > \alpha^* x_{l_j} > x_{l_i}$ for $i+r+1 \leq j \leq n$. Hence, $z_{(j)} = x_{(j)}$ for $1 \leq j \leq i - 1$, and $z_{(i)} > x_{l_i} = x_{(i)}$. This contradicts the fact that x is lexicographically fair. ■

Corollary 7: A compact and convex set has exactly one lexicographically fair vector.

Proof: The conclusion follows from Theorem 5 and Theorem 6. ■

III. Characterizing Max-min Fairness by Optimization

It is known that many fairness concepts can be characterized by the solutions to optimization problems. For instance, the proportional fairness is defined as the optimal solution to the problem $\max \sum_{i=1}^n \log(x_i)$ subject to $x \in S$, where $S \subseteq \mathbb{R}_+^n$ is a non-empty compact and convex set.

Let $U^\alpha : \mathbb{R} \to \mathbb{R}$ be a set of continuously differentiable, strictly increasing functions, with the property, for every $x = (x_1, \ldots, x_n) \in S$, if $x_i < x_j$,

$$\lim_{\alpha \to \infty} \frac{\dot{U}^\alpha(x_i)}{\dot{U}^\alpha(x_j)} = \infty \tag{1}$$

Consider the sequence of optimization problems.

$$\max \sum_{i=1}^n U^\alpha(x_i) \tag{2}$$

$$\text{s.t.} \quad x \in S \tag{3}$$

Suppose x^α is a solution to the above problem, for each fixed α. The following theorem has been cited in the literature many times. But, there hasn't been a rigorous proof for it.

Therorem 8: Every limit point of the sequence $\{x^\alpha\}$ is max-min fair.

Proof: Let x be a limit point of $\{x^\alpha\}$. Without loss of generality, suppose $x_1 \leq \ldots \leq x_n$, possibly by renaming the axis. We wish to show that x is max-min fair. Consider any vector $y \in S$ and $y \neq x$. Suppose $k = \text{argmin}_{1 \leq j \leq n} x_j \neq y_j$. Suppose $x_k = \ldots = x_{k+s}$ for some $s \geq 0$. Suppose, if $k + s + 1 \leq n$, $x_k \neq x_{k+s+1}$. Also suppose the indices are ordered such that if $x_j = x_{j+1}$, then $y_j \leq y_{j+1}$. Suppose $r = \max_{0 \leq j \leq s} x_{k+j} = y_{k+j}$. Hence, $y_k = y_{k+1} = \ldots = y_{k+r}$, and if $r < s$, $y_{k+r} < y_{k+r+1}$.

Let $\{x^\beta\}$ be a subsequence of $\{x^\alpha\}$ converging to x. Let $\{y^\beta\}$ be a sequence converging to y, with the property that, for each β, $y_j^\beta = x_j^\beta$, for $1 \leq j \leq k - 1$. We will show such a sequence exists. Define the hyperplane H^β and H by,

$$H^\beta = \{(x_1^\beta, \ldots, x_{k-1}^\beta, z_k, z_{k+1}, \ldots, z_n) | z_j \in \mathbb{R}, \ k \leq j \leq n\},$$

$$H = \{(x_1, \ldots, x_{k-1}, z_k, z_{k+1}, \ldots, z_n) | z_j \in \mathbb{R}, \ k \leq j \leq n\}.$$

Note that, as $\beta \to \infty$, H^β approaches H, in the sense that the normal vector to H^β approaches the normal vector to H. Let $T^\beta = H^\beta \cap S$ and let $T = H \cap S$. Clearly, T^β and T are both convex and compact sets. Note that $x, y \in T$. Let $y^\beta = \text{argmin}_{z \in T^\beta} ||y - z||$, where $|| \cdot ||$ is the Euclidean norm. That is, y^β is the projection of y on the set T^β, which exists and is unique. Since T^β becomes T asymptotically and $y \in T$, we have $||y^\beta - y|| \to 0$ as $\beta \to \infty$. Hence, $y^\beta \to y$ as $\beta \to \infty$ [1].

By the optimality condition, for every β,

$$\sum_{j=1}^n \dot{U}^\beta(x_j^\beta)(y_j^\beta - x_j^\beta) \leq 0. \tag{4}$$

Since $y_j^\beta = x_j^\beta$, for $1 \leq j \leq k - 1$,

$$\sum_{j=k}^n \dot{U}^\beta(x_j^\beta)(y_j^\beta - x_j^\beta) \leq 0. \tag{5}$$

Since $x^\beta \to x$ as $\beta \to \infty$ and $x_j > x_k$ for $k + s < j \leq n$, we have $\dot{U}^\beta(x_j^\beta)/\dot{U}^\beta(x_k^\beta) \to 0$ as $\beta \to \infty$. Therefore, the sign of the sum in (5) is determined by $\sum_{j=k}^{k+s} \dot{U}^\beta(x_j^\beta)(y_j^\beta - x_j^\beta)$ when β is large enough. Hence, there exists β_o such that for all $\beta \geq \beta_o$,

$$\sum_{j=k}^{k+s} \dot{U}^\beta(x_j^\beta)(y_j^\beta - x_j^\beta) \leq 0. \tag{6}$$

This implies, for each $\beta \geq \beta_o$,

$$\min_{k \leq j \leq k+s} \dot{U}^\beta(x_j^\beta)(y_j^\beta - x_j^\beta) \leq 0. \tag{7}$$

Since $\dot{U}^\beta(x_j^\beta) \geq 0$, for all j and β,

$$\min_{k \leq j \leq k+s} (y_j^\beta - x_j^\beta) \leq 0, \tag{8}$$

for $\beta \geq \beta_o$.

When $r < s$, $y_k = \ldots = y_{k+r} < y_{k+r+1} \leq \ldots \leq y_{k+s}$. Hence, $y_k - x_k = \ldots = y_{k+r} - x_{k+r} < y_{k+r+1} - x_{k+r+1} \leq \ldots \leq y_{k+s} - x_{k+s}$. Then, there exists β_1 such that for all $\beta \geq \max\{\beta_o, \beta_1\}$, the minimum in (8) is achieved by a term with the index j that always satisfies $k \leq j \leq k + r$. For any such j, $x_j^\beta \to x_k$ and $y_j^\beta \to y_k$ as $\beta \to \infty$. Hence, $y_k \leq x_k$. By the assumption that $y_k \neq x_k$, it must be that $y_k < x_k$. When $r = s$, $x_j^\beta \to x_k$ and $y_j^\beta \to y_k$ for all j, $k \leq j \leq k + s$, as $\beta \to \infty$. From (8), by letting $\beta \to \infty$, we again reach the conclusion $y_k < x_k$. By Lemma 3, x is max-min fair. ■

The uniqueness of the max-min fair vector implies the following.

Corollary 9: The sequence $\{x^\alpha\}$ in Theorem 8 converges.

References

[1] Dimitri P. Bertsekas and Robert Gallager. *Data Networks*. Prentice-Hall, Inc., 2 edition, 1992.

[2] Frank Kelly. Charging and rate control for elastic traffic. *European Transactions on Telecommunications*, 8:33–37, 1997.

[3] Abhay K. Parekh and Robert Gallager. A generalized processor sharing approach to flow control in integrated services networks: The single-node case. *IEEE/ACM Transactions on Networking, Vol 1, NO. 3*, pages 344–357, June 1993.

[4] B. Radunović and J.-Y. Le Boudec. A unified framework for max-min and min-max fairness with applications. *ACM/IEEE Transactions on Networking*, 15(5), October 2007.

[5] S. Sarkar and L. Tassiulas. Bandwidth layers in multicast networks. In *Proceedings of the IEEE Infocom*, 2000.

[1] To argue this rigorously, let $\hat{y}^\beta$ be the projection of y onto H^β. $||\hat{y}^\beta - y|| \to 0$ as $\beta \to \infty$. y^β is the projection of $\hat{y}^\beta$ onto T^β. Since $||y^\beta - y|| \leq ||\hat{y}^\beta - y||$, $||y^\beta - y|| \to 0$ as $\beta \to \infty$.