Maxmin Fairness under Priority for Network Resource Allocation Tasks

Mario Köppen, Kei Ohnishi, and Masato Tsuru
Computer Science and Engineering Department
Kyushu Institute of Technology
680-4 Kawazu, Iizuka, 820-8502 Fukuoka, Japan
Email: mkoeppen,ohnishi,tsuru@cse.kyutech.ac.jp

***Abstract*—Many network control policies can benefit from introducing priorities among users, traffic flows, or service provisions e.g. for QoS improvement or network congestion avoidance. In order to ensure fairness of concomitant resource sharing tasks, generic extensions of maxmin fairness under priority are considered. A critical analysis of existing approaches leads to the definition of two fairness relations based on formal modifications of the maxmin fairness standard of comparison. One is based on using priority functions for internal weighting, the other on priority classes. Experimental evaluation shows that maximizing under these relations indeed gives relations close to maxmin fairness that take priorities into account such that higher priority users receive higher allocations in average. Further studying the wireless channel allocation model problem shows that prioritizing can give solutions of higher efficiency than maxmin fairness.**

Keywords*—*fairness, maxmin fairness, priority fairness, wireless channel allocation

I. Introduction

The rapid growth of communication networks has given the need for smart control policies that also take aspects of the distribution and sharing of network resources into account. This is commonly studied under the theme of fairness. Most popular fairness models here are maxmin fairness [1] and proportional fairness [2]. In these fairness models, it is essential that all users are treated equally in the resource allocation.

However, there are cases where priority among users or flows has to be taken into account. Priorities in wireless networks are often related to the management of QoS classes or as indicator for network congestion. In [3] various queuing schemes for multi-hop wireless networks are compared according to fairness and throughput. It is concluded that for achieving optimal bandwidth utilization, the MAC-layer needs to support different QoS priorities. Also in [4] the improvement especially of multimedia services quality is related to QoS priority classes and an adaptive fairness scheme is considered. In [5] authors propose a priority-based fair medium access control (PMAC) protocol to reflect different weights among traffic flows during approximation of the optimal contention window size. Per-flow weights, related to congestion at nearby nodes, are also at the base of the approach presented in [6]. This scheme approaches long-term fairness, where allocated bandwidths to flows equal their long-termed arrival rate.

In such works there is no explicit model of fairness that can be applied to a wide range of network architectures. A similar problem domain where prioritization of users or flows is needed to establish a suitable control paradigm can be found in related fields: in cognitive radio there is explicit separation between primary and secondary users. Opportunistic networking might take the chance of users occurrence into account and this way maintaining different node priorities. In a practical context, in a disaster situation emergency-case related packets should have higher transport priority, while other traffic is still needed to probably - by distributing new and updated information - adjust the emergency traffic content. Needless to say that all these cases strongly refer to cross-layer control. Nevertheless, in all cases the same need for fairness in the distribution arises. Only few studies consider generic approaches to fairness under priority, including so-called weighted maxmin fairness in [7] or [8]. A common aspect here is the use of priority functions that make priority a function of the allocation (share, rate etc.). This is commonly related to charging for bandwidth, where the user who pays more can expect a specific traffic rate to be higher prioritized than for a user who pays less.

But it was shown that with regard to maxmin fairness all related proposals basically emulate maxmin fairness under a corresponding variable transformation [8]. Here, we study two possible extensions of maxmin fairness that can serve as new fairness relations under priority, one based on priority functions, one based on priority classes. The approaches are based on direct formal modification of the maxmin fairness standard of comparison. Problems with the existing extensions of maxmin fairness will be discussed in Section II. Section III then details the proposed maxmin fairness under priority relations. Section IV provides results of experiments with these two new relations to demonstrate their feasibility. Section V concludes the paper.

II. Extensions of maxmin fairness

For the inclusion of priority in maxmin fairness a more closer look on maxmin fairness itself is needed. Generally, decision-making problems like fair selection can be done by specifying a *standard of comparison* and seeking solutions where each other feasible solution fulfills such a standard. For example, given a set of vectors, the well-known *Pareto efficiency* for some solution x is based on the following:

(**Pareto-efficiency**) *Increasing one component x_i must be at the expense of decreasing some other x_j.*

In [1], based on the analysis of the bottleneck flow control algorithm, the standard of comparison for maxmin fairness was given as:

(**Maxmin fairness**) *Increasing one component x_i must be at the expense of decreasing some other x_j such that $x_i \geq x_j$.*

We may formulate this as a binary vector relation between elements of a set of feasible vectors A.

Definition 1. *Given a domain $A \subseteq R^n$. Vector $x \in A$* maxmin fair dominates *vector $y \in A$, written as $x \geq_{mmf} y$, if and only if*

$$\forall i = 1, \dots, n : (y_i > x_i) \rightarrow \exists_{j \neq i}(x_i \geq x_j \wedge x_j > y_j). \quad (1)$$

Then the maxmin fair (or maxmin equitable) state is a *greatest element* of this relation with regard to the domain A, i.e. an element x' such that for any $x \in A$ it holds $x' \geq_{mmf} x$. Thus, a relation "$x \geq_{mmf} y$" expresses a lacking incentive to change from state x to state y, or generally a preference for the solution x against y.

Obviously, this formulation of maxmin fairness considers all components equally and does not cover any notion of priority. An early statement about a possible inclusion can be found in the same source [1]. There (p.528) the following suggestion was made:

"*Several generalizations can be made to the basic approach described above. ... Next, we can consider ways to assign different priorities to different kinds of traffic and to make these priorities sensitive to traffic levels. If $b_p(r_p)$ is an increasing function representing the priority of p at rate r_p, the max-min fairness criterion can be modified as follows: For each p, maximize r_p subject to the constraint that any increase in r_p would cause a decrease of $r_{p'}$ for some p' satisfying $b_{p'}(r_{p'}) \leq b_p(r_p)$.*"

Also here, the maximization problem can be rephrased as the finding of a greatest element for the relation:

Definition 2. *Given a domain $A \subseteq R^n$ and a set of n (strictly) increasing real-valued priority functions $p_i(x)$. Vector $x \in A$* maxmin fair dominates *vector $y \in A$ under priority if and only if*

$$\forall i = 1, \dots, n : (y_i > x_i) \rightarrow \exists_{j \neq i}(p_i(x_i) \geq p_j(x_j) \wedge x_j > y_j). \quad (2)$$

This relation was later independently introduced as *utility maxmin fairness* in [9]. A special case of this fairness concept can also be found under the name weighted maxmin fairness, for example in [10]:

Definition 3. *Given a domain $A \subseteq R^n$ and a set of n positive real-valued weights w_i. Vector $x \in A$* maxmin fair dominates *vector $y \in A$ under priority weights w_i if and only if*

$$\forall i = 1, \dots, n : (y_i > x_i) \rightarrow \exists_{j \neq i} x_i / w_i \geq x_j / w_j \wedge x_j > y_j). \quad (3)$$

It is easy to see that this definition corresponds to the former one for the choice of priority functions $p_i(x) = x/w_i$. However (see also corresponding remark in [8]) none of these definitions actually defines a new relation. Consider a transformation

$$(x_1, x_2, \dots, x_n) \longrightarrow (p_1(x_1), p_2(x_2), \dots, p_n(x_n)), \quad (4)$$

then the definition of fairness under priority becomes equivalent to the definition of maxmin fairness for the transformed vectors, simply since the priority functions are strictly increasing[1]. Thus, $y_i > x_i$ in the definition of maxmin fairness can be replaced with $p_i(y_i) > p_i(x_i)$ as well as $x_j > y_j$ with $p_j(x_j) > p_j(y_j)$. This fact might limit the applicability of the concept, since maxmin fairness by itself was not designed to represent priority. Assume in a limiting case that the priority functions would not change strongly and as a consequence there is a fixed ranking of users by priority. Then Def. 2 would read as: the rate of the user with constant lowest priority could always grow, while the rates for the user with constant highest priority could only grow if all other rates grow as well. This is not really fitting to a concept of priority.

In [7] weighted maxmin fairness is slightly differently introduced:

Definition 4. *Given a domain $A \subseteq R^n$ and a set of n (strictly) increasing real-valued priority functions $p_i(x)$. Vector $x \in A$* maxmin fair dominates *vector $y \in A$ under weighted priority if and only if*

$$\forall i = 1, \dots, n : (y_i > x_i) \rightarrow \exists_{j \neq i}(p_i(x_i) \geq p_i(x_j) \wedge x_j > y_j). \quad (5)$$

Note that the difference here is to use the same priority function p_i for comparing between components of the vector x. However, also here the strictly increasing character of the priority function makes this equivalent to $x_i \geq x_j$ and the relation is equivalent to maxmin fairness as well. The same would happen if using p_j instead of p_i.

It might be surprising that extending maxmin fairness appears less simple than it seems on first glance. In the next section, we will take up this issue and propose two extensions of maxmin fairness that allow for the inclusion of priority aspects in maxmin fairness and define relations that are different from maxmin fairness.

III. Priority Maxmin Fairness Relations

The former section demonstrated how recent proposals of a maxmin fairness relation under weighting or priority define in fact relations that are homomorph to the original maxmin fairness, with regard to maximum state selection. To avoid this problem, one can formally make a slight adjustment to Defs. 2 or 4 that refers to the comparison between the components of x:

Definition 5. *(Weighted Maxmin Fairness) Given a domain $A \subseteq R^n$ and a set of n (strictly) increasing real-valued priority functions $p_i(x)$. Vector $x \in A$* maxmin fair dominates *vector*

[1] The case of (not necessarily strictly) increasing functions makes some extra formal effort, but will not change the general picture.

$y \in A$ *under weighting* p_i *(written as* $x \geq_{wmmf} y$*) if and only if*

$$\forall i = 1, \ldots, n : (y_i > x_i) \rightarrow \exists_{j \neq i}(p_j(x_i) \geq p_i(x_j) \wedge x_j > y_j). \quad (6)$$

In fact, here the comparison $p_j(x_i) \geq p_i(x_j)$ works as a "promoter" of the comparison between x_i and x_j, driven by the mutual priority functions. One can see this by considering an example: for the vectors $x = (0.4, 0.5, 0.6)$ and $y = (0.5, 0.6, 0.5)$ there would be no maxmin fairness relation: in both cases where a component increases if changing state from x to y there would be no component that becomes smaller (which is the case for index 3) but also already smaller or less than the increasing component (so, 0.6 is larger than 0.4 and 0.5). Now assume that the third component is actually the allocation to a user with higher priority than users 1 and 2. It can be expressed by priority functions $p_1(x) = p_2(x) = x^2$ and $p_3(x) = x$ - so at the same allocation (rate, share etc.) user 3 would always have higher priority than users 1 and 2, if the components are from $(0, 1]$. But now vector x is in weighted maxmin fairness relation to y, since the condition $p_j(x_i) \geq p_i(x_j)$ is fulfilled for the (i, j)-pairs $(1, 3)$ and $(2, 3)$ (for example for $(1, 3)$ $p_3(x_1) = x_1 = 0.4 \geq p_1(x_3) = 0.6^2 = 0.36$). So by taking priorities into account, the modified maxmin fairness relation holds.

We will refer to this relation as *weighted maxmin fairness* in the following, or simply *wmmf.*

But there is also the possibility to take priorities explicitly into account, by posing an additional constraint in the definition. Here, priorities are seen as index classes, independent from any component magnitude. Formally priority values p_i are assigned to each index, representing the priority of the agent receiving the magnitude given by vector component i (and being independent from that magnitude). Then the condition $x_i \geq x_j$ to identify an agent j that envies benefitting agent i in the maxmin fairness relation condition, this agent j must also be of higher priority.

Definition 6. *(Priority Maxmin Fairness) Given a domain* $A \subseteq R^n$ *and a set of real priority values* p_i $(i = 1, \ldots, n)$*. Vector* $x \in A$ *maxmin fair dominates vector* $y \in A$ *under priorities* p_i *(written as* $x \geq_{pmmf} y$*) if and only if*

$$\forall i = 1, \ldots, n : (y_i > x_i) \rightarrow \exists_{j \neq i} x_i \geq x_j \wedge p_j \geq p_i \wedge x_j > y_j). \quad (7)$$

Both definitions refer to relations that are different from maxmin fairness. The difference (with regard to their application) is the way the priority is specified: either as a fixed assignment to the specific user (as it is for example the distinction of primary and secondary user in cognitive radio), or as a weighting that represents the payment of a specific user in relation to payments of other users.

IV. Application to Cross-layer Wireless Architectures

For using fairness relations for network resource allocation problems, we follow the way presented in [11]. For the development of control policies in cross-layer wireless architectures, the problem domains are often discrete. One example would be channel allocation based on Channel State Information (CSI) where a set of transmission channels provided by a Base Station (BS) have to be allocated to Subscriber Stations (SS), while accepting a possible lower transmission via allocated channels due to mobility of the SS, larger distance to BS or obstacles in the cell coverage area. Channels are discrete units. The issue now is that in such cases, often there is no greatest element of the fairness relation (or in economical terms: no equitable solution to the standard of comparison). In order to come up with an effective way of selecting fair allocations, inspiration can be taken from multi-objective optimization: for the Pareto-efficiency, expressed as Pareto-dominance relation, *maximum elements* of the feasible domain are searched, and - since there are usually more than one - a *Decision Maker* (DM) selects finally. If writing relation R as $x \geq_R y$ (for domain elements x and y), then the asymmetric part $P(R)$ is the relation composed of all pairs $(x, y) \in R$ where $(y, x) \notin R$ and the notation $x >_R y$ is used. An element x' of the domain is *maximal* if there is no domain element x such that $x >_R x'$. All maximal elements comprise the *maximum set* of the relation R.

For establishing a fairness relation, the relation should be at least implied by Pareto dominance (this way, the maximum allocations become Pareto efficient), and it should be cycle-free (i.e. there is no sequence of $k \geq 3$ vectors such that $x_1 >_R x_2 >_R \cdots >_R x_k$ and $x_k >_R x_1$). For cycle-free relations, each finite domain has a non-empty maximum set (see e.g. [12] for more details on the relational approach).

Considering Definitions 5 and 6 given in the foregoing section, both relations are obviously implied by Pareto dominance since this means that for each index i $x_i \geq y_i$ and thus there is no index i such that $y_i > x_i$. Regarding cycle-freeness, the answer is easy for the relation *pmmf* of Definition 6: since the exists-part in the condition has an additional constraint (the priority of j is at least as large as the priority of i) it follows that the relation *pmmf* is a subset of the maxmin fairness relation, and since maxmin fairness is cycle-free, *pmmf* is cycle-free as well.

For the relation *wmmf* the answer is not so simple, since there are many choices for the priority functions. The identification of a class of priority functions where the relation is cycle-free is topic of ongoing investigation. At least in all experiments so far using priority functions like exponential functions, sigmoid functions, and their converse, no cycle could be observed. However, for some priority function choices such cycles exist: consider a case of 4 dimensional vectors, using the priority functions

$$p_i(x) = \begin{cases} i & \text{for } i = 1, 2, 3 \\ 4\delta(x - 5) & \text{for } i = 4 \end{cases} \quad (8)$$

where $\delta(x) = 0$ for $x < 0$ and 1 otherwise[2]. Now for the triple of vectors $x = (2, 4, 6, 8), y = (9, 3, 8, 4), z = (3, 7, 6, 4)$ it can be seen that $x >_{wmmf} y \wedge y >_{wmmf} z \wedge z >_{wmmf} x$ (for space reason, we cannot provide the complete calculation here).

[2]Slight variation of these functions also allows to find cases of cycles where the priority functions are strictly monotone

As cycles are - nevertheless - rather sparse, with some caution we can also consider relation *wmmf* as suitable for selection of maximum elements.

V. Experimental Evaluation

In this section, the two maxmin fairness relations under priority, Weighted Maxmin Fairness (*wmmf*, Def. 5) and Priority Maxmin Fairness (*pmmf*, Def. 6) and their relation to Maxmin Fairness will be further investigated by numerical simulations, in order to learn about specific strong and weak points of these definitions with regard to applications. Three types of experiments are performed: first experiments serve to find out about the role as a preference relation in general, then experiments to see how these relations reflect priority, and a final set of experiments consider the efficiency of the maximum selection for the wireless channel allocation problem (the “price of fairness” [13]).

A. Experiment 1: Relational Properties

We randomly sample the frequency of the occurrence of related pairs for *wmmf* and *pmmf* under varying priority functions or values. In case of *wmmf*, the priority functions are selected as equal for all indizes except 1 and 2. For these two, we choose power functions x^n. The *priority index* n gives a larger priority to users 1 and 2 in case of negative n and a lower priority in case of positive n. For $n = 1$ the relation is equivalent to maxmin fairness. The results for 1 Mio. samples of 5-dimensional random vectors with components from $(0, 1]$ are given in Table I.

TABLE I. Numbers of occurrences of maxmin fairness (*mmf*), weighted maxmin fairness (*wmmf*) and their intersection among 1 Mio. 5-dimensional random vectors with components from $(0, 1]$. The priority functions for *wmmf* were $p(x) = x^n$ for index 1 and 2, $p(x) = x$ otherwise. For the case $n = 1$ *wmmf*=*mmf*.

n	occ. *mmf*	occ. *wmmf*	occ. both
-10.0	167149	666058	119826
-9.0	167120	665988	119754
-8.0	167209	665941	119861
-7.0	167131	666055	119802
-6.0	167141	665992	119773
-5.0	167192	665951	119852
-4.0	167142	666043	119809
-3.0	167131	665965	119761
-2.0	167191	665958	119856
-1.0	167158	666070	119822
0.0	167114	749199	154088
1.0	167205	167205	167205
2.0	167133	188835	139377
3.0	167134	210222	133849
4.0	167203	223534	132284
5.0	167128	231326	131769
6.0	167149	236423	131558
7.0	167189	240161	131592
8.0	167129	242568	131540
9.0	167156	244160	131449
10.0	167173	245622	131470

It is obvious that *wmmf* and *mmf* are different relations (column 4 values are lower than both, column 2 and 3 values). The second column shows the distribution of *mmf* occurrence that is rather constant (about 1/6 of the domain). In all cases, *wmmf* is more frequent, with quite high and nearly constant values of about 66% for the case where agents 1 and 2 have higher priority than the others, and values slightly larger than *mmf* (about 24%) where these agents have lower priority (positive priority index n). The number of cases where a random pair is in both relations, *mmf* and *wmmf*, seems rather independent from the priority functions here.

TABLE II. Numbers of occurrences of maxmin fairness (*mmf*), priority maxmin fairness (*pmmf*) and their intersection among 1 Mio. 10-dimensional random vectors with components from $(0, 1]$. The priority values for *pmmf* were 2 for indizes 1 to n and 1 otherwise. For the cases $n = 0$ and $n = 10$ *pmmf*=*mmf*.

n	occ. *mmf*	occ. *pmmf*	both
0	85487	85487	85487
1	85454	53341	53341
2	85489	47126	47126
3	85454	47894	47894
4	85472	50918	50918
5	85505	55112	55112
6	85446	60371	60371
7	85450	65545	65545
8	85514	72156	72156
9	85463	78763	78763
10	85458	85458	85458

The corresponding experiment for *pmmf* gives a different picture. Here, 1 Mio. 10-dimensional random vectors with components from $(0, 1]$ were sampled and for *pmmf* the priority was set to 2 (i.e. higher priority) for the first n users, and 1 otherwise. The results are shown in Table II. The fact that the values in column 4 are equal to the values in column 3 and never larger than the values in column 2 confirms that *pmmf* is always a subset of *mmf*. However, there is a notable influence on the frequency of *pmmf* relation occurrence only in case of a small number of higher priority users.

How will this influence the typical size of maximum sets (which represents the hardness for a DM to select a final solution)? The results for the same priority function or value settings as before are shown in Figs. 1 and 2. In both cases, 30 times 100 random vectors with components from $(0, 1]$ were constructed (5-dimensional for *wmmf*, 10-dimensional for *pmmf*) and the average size of the maximum sets was calculated. In general, a more abundant relation (as long as it stays cycle-free) will produce smaller maximum sets. Figure 1 confirms this. Except for $n = 0$ (a constant priority case) the maximum sets for *wmmf* are indeed smaller than the maximum sets for *mmf* (which have a size of about 3). In general, like mmf *wmmf* appears suitable for effective selection by relational maximality.

This is not the case for *pmmf*. Figure 2 shows that the relation produces maximum sets that are by a factor 2 to 3 larger than the maximum sets of maxmin fairness. Since *pmmf* implies *mmf*, they will never be smaller, but especially for the low-frequency case of few users with higher priority, the maximum sets are much larger than for maxmin fairness.

B. Experiment 2: Representation of Priority

The purpose here is to find out in which way these relations represent priority. Generally it is hard to have a clear-cut criterion for deciding whether priority was taken into account or not. We considered the implicit effect that higher priority users should in average have a higher allocation. Thus we took maximum sets with the same settings as in the experiment on maximum set sizes before (regarding the choice of priority functions or values, and experimental set up) but where looking for the average component magnitude in maximal elements for

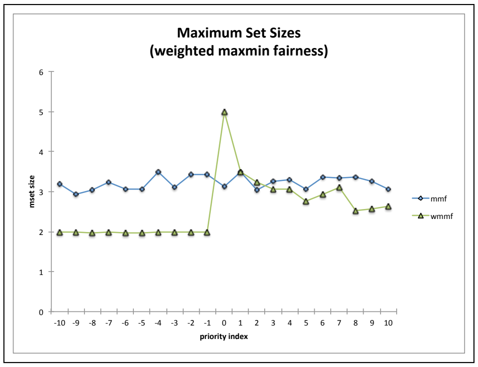

Fig. 1. Comparison of sizes of maximum sets for maxmin fairness and the proposed weighted maxmin fairness in case of 5 dimensions. The priority functions where $p(x) = x^n$ for users 1 and 2, where n is the *priority index*, and $p(x) = x$ for users 3,4 and 5. The maximum sets were computed for 100 random vectors with components from $(0, 1]$, and the average values after 30 experiments are shown.

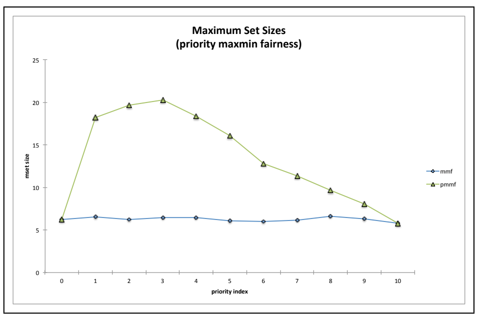

Fig. 2. Comparison of sizes of maximum sets for maxmin fairness and the proposed priority maxmin fairness in case of 10 dimensions. The priorities were 2 for users 1 to n (where n is the priority index here) and 1 for the remaining users. The maximum sets were computed for 100 random vectors with components from $(0, 1]$, and the average values after 30 experiments are shown.

priority users compared to other users. The results are shown in Tables III and IV.

Here, both relations behave quite similar. In all cases higher priority corresponds with higher magnitudes, while this is stronger for *wmmf* (about 60% larger values) than *pmmf* (about 20% larger values). In case of users with lower priority (case of positive n in Table III) users receive notably less, (also) about 60% the value of normal priority users. Both tables also show (second column) that for maxmin fairness there is no such imbalance between users. So these experiments confirm a tendency to take priority into account by higher allocations to higher priority users.

C. Experiment 3: The Price of Fairness

We also wanted to study the relations within a specific network resource allocation problem. As already mentioned, we used the Wireless Channel Allocation problem (WCA) where a number of channels (within a time frame) have to be allocated to users and the performance of the allocation per user is represented by the sum of channel coefficients of allocated channels (a channel coefficient represents the rate at which a user can transmit via an allocated channel). It is a simplified model for the real-world wireless infrastructure, but represents essential aspects of this problem domain.

TABLE III. COMPARISON OF AVERAGE COMPONENTS OF PRIORITY AND NON-PRIORITY USERS FOR MAXIMAL ELEMENTS OF MAXMIN FAIRNESS (MMF) AND PROPOSED WEIGHTED MAXMIN FAIRNESS (*wmmf*) RELATIONS. THE PRIORITY FUNCTIONS WHERE $p(x) = x^n$ FOR USERS 1 AND 2 AND $p(x) = x$ FOR USERS 3,4 AND 5. THE MAXIMUM SETS WERE COMPUTED FOR 100 5-DIMENSIONAL RANDOM VECTORS WITH COMPONENTS FROM $(0, 1]$, AND THE AVERAGE VALUES AFTER 30 EXPERIMENTS ARE GIVEN. FOR $n = 1$ *wmmf*=*mmf*.

n	ratio *mmf*	ratio *wmmf*
-10.0	1.00271	1.50686
-9.0	0.993914	1.64135
-8.0	1.00336	1.59223
-7.0	0.998713	1.36124
-6.0	0.999166	1.77754
-5.0	1.00944	1.57733
-4.0	1.00303	1.60155
-3.0	1.0006	1.752
-2.0	0.988285	1.46511
-1.0	1.00396	1.72591
0.0	1.00252	1.72035
1.0	1.0082	1.0082
2.0	1.01914	0.852477
3.0	1.01702	0.725592
4.0	1.01636	0.735208
5.0	1.00474	0.675472
6.0	1.01269	0.636196
7.0	1.03122	0.632815
8.0	1.01624	0.627458
9.0	1.00453	0.647282
10.0	1.0307	0.602968

TABLE IV. COMPARISON OF AVERAGE COMPONENTS OF PRIORITY AND NON-PRIORITY USERS FOR MAXIMAL ELEMENTS OF MAXMIN FAIRNESS (MMF) AND PROPOSED PRIORITY MAXMIN FAIRNESS (*pmmf*) RELATIONS. THE PRIORITY VALUES WERE 2 FOR INDIZES 1 TO n AND 1 OTHERWISE. THE MAXIMUM SETS WERE COMPUTED FOR 100 10-DIMENSIONAL RANDOM VECTORS WITH COMPONENTS FROM $(0, 1]$, AND THE AVERAGE VALUES AFTER 30 EXPERIMENTS ARE GIVEN.

n	ratio *mmf*	ratio *pmmf*
2	1.00529	1.31717
3	0.99703	1.25191
4	1.01491	1.23798
5	0.991467	1.20122
6	0.997243	1.19441
7	0.968695	1.19406
8	1.02128	1.22074

We selected the problem of allocating 6 channels to 5 users. The channel coefficients were chosen with a uniform radial distribution. Each channel must be allocated to exactly one user, and each user receives at least one channel. This gives 1800 feasible allocations in total. Among them, maximum sets were selected for various fairness relations, including (1) maxmin fairness *mmf*, (2) *wmmf* with priority functions $p(x) = x^2$ for users 1 and 2, $p(x) = x$ otherwise, (3) *pmmf* with priority 2 for users 1 and 2, 1 otherwise, and (4) proportional fairness [2]. In each case, from the maximum set the vector with largest sum of components was selected, and this sum put into ratio with the maximal possible sum among all feasible allocations. The average ratios over 30 experiments were calculated.

The results are shown in Fig. 3. It shows a clear ranking of the relations with regard to efficiency (by “efficiency” here we consider the sum of vector components). We give the additional information that the same ranking appeared in all 30 test cases. The fairness under priority relations appear between maxmin fairness and proportional fairness. From [13], where the phrase “price of fairness” was coined and which inspired related investigations, it is a known fact that generally proportional fairness is more efficient than maxmin fairness.

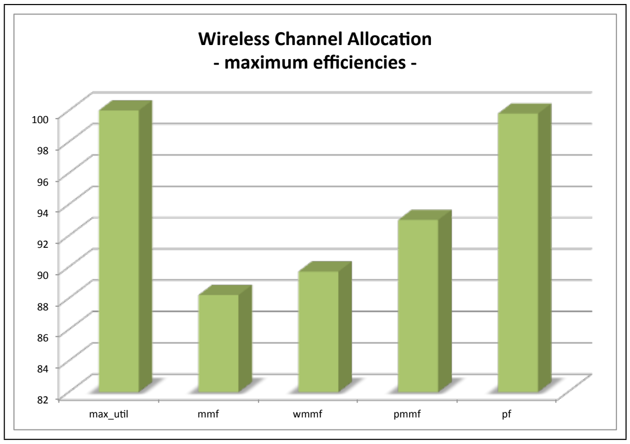

Fig. 3. The "price of fairness" for various fairness relations at a wireless channel allocation problem (5 users, 6 channels, i.e. 1800 feasible allocations). The bar heights correspond with the the average ratio in percent of 100 of the total of maximal allocations (per relation) to the maximum possible total. The values are averages after 30 experiments. See text for details of the relation settings. "pf" stands for proportional fairness.

This is confirmed in this experiment, where mmf gives lowest efficiency (about 88% of maximum possible) and proportional fairness the largest (about 98% of maximum possible). It was also noted in [13] that generally such an efficiency loss, when taking fairness into account, is not very high, hardly exceeding 10%. So even with regard to efficiency, fairness is still an attractive optimization goal. We can now see that priority can help to improve efficiency to some degree, while maintaining "least component preference" - i.e. favouring solutions where the minimum allocation is maximized. The relation *pmmf* performs a little bit better than *wmmf* here, but this can be influenced by the larger maximum sets and thus greater variety of maximal elements of *pmmf*.

VI. Conclusions

Both relations that were introduced in this paper share similarities with maxmin fairness, but also appear distinct from maxmin fairness and each other to a sufficient degree. Both relations modify the maxmin fairness standard of comparison, which states that for any increase x_i there must be a decrease in another x_j such that $x_i \geq x_j$. In case of the proposed weighted maxmin fairness (*wmmf*) relation it is the condition $p_j(x_i) \geq p_i(x_j)$ instead of (just) $x_i \geq x_j$, where the $p_i(x)$ are the priority functions. For priority maxmin fairness (*pmmf*), the additional condition $p_j \geq p_i$ is introduced. Relation *pmmf* is a narration of maxmin fairness, while *wmmf* appears more abundant than maxmin fairness, but closer in decision making tasks like maximum element selection. On the other hand, *pmmf* selects solutions of higher efficiency (at least for the task studied here). Both relations might have their specific application fields: *wmmf* is more attractive if a complex system of priority functions can be specified, while *pmmf* is more attractive if only few priority classes are given by the problem specification, and the focus is on more efficient solutions. Further investigations of the properties of these relations (esp. for the conditions on the priority functions under which *wmmf* appears to be cycle-free, or nearly cycle-free) and theoretical underpinning of the observations made in the experiments are subject of future work.

References

[1] D. P. Bertsekas, R. G. Gallager, and P. Humblet, *Data networks*. Prentice-Hall International, 1992, vol. 2.

[2] F. Kelly, "Charging and rate control for elastic traffic," *Eur. Trans. Telecomm.*, vol. 8, pp. 33–37, Jan./Feb. 1997.

[3] J. Jun and M. L. Sichitiu, "Fairness and QoS in multihop wireless networks," in *Vehicular Technology Conference, 2003. VTC 2003-Fall. 2003 IEEE 58th*, vol. 5. IEEE, 2003, pp. 2936–2940.

[4] M. Malli, Q. Ni, T. Turletti, and C. Barakat, "Adaptive fair channel allocation for QoS enhancement in IEEE 802.11 wireless LANs," in *Communications, 2004 IEEE International Conference on*, vol. 6. IEEE, 2004, pp. 3470–3475.

[5] D. Qiao and K. G. Shin, "Achieving efficient channel utilization and weighted fairness for data communications in IEEE 802.11 WLAN under the DCF," in *Quality of Service, 2002. Tenth IEEE International Workshop on*. IEEE, 2002, pp. 227–236.

[6] L. Tassiulas and S. Sarkar, "Maxmin fair scheduling in wireless networks," in *INFOCOM 2002. Twenty-First Annual Joint Conference of the IEEE Computer and Communications Societies. Proceedings. IEEE*, vol. 2. IEEE, 2002, pp. 763–772.

[7] P. Marbach, "Priority service and max-min fairness," in *INFOCOM 2002. Twenty-First Annual Joint Conference of the IEEE Computer and Communications Societies. Proceedings. IEEE*, vol. 1. IEEE, 2002, pp. 266–275.

[8] B. Radunović and J.-Y. L. Boudec, "A unified framework for max-min and min-max fairness with applications," *IEEE/ACM Transactions on Networking (TON)*, vol. 15, no. 5, pp. 1073–1083, 2007.

[9] D. Rubenstein, J. Kurose, and D. Towsley, "The impact of multicast layering on network fairness," *IEEE/ACM Transactions on Networking (TON)*, vol. 10, no. 2, pp. 169–182, 2002.

[10] M. Allalouf and Y. Shavitt, "Maximum flow routing with weighted max-min fairness," in *Quality of Service in the Emerging Networking Panorama*. Springer, 2004, pp. 278–287.

[11] M. Koeppen, "Relational optimization and its application: From bottleneck flow control to wireless channel allocation," *INFORMATICA*, vol. 24, no. 3, pp. 413–433, 2013.

[12] K. Suzumura, *Rational choice, collective decisions, and social welfare*. Cambridge University Press, 2009.

[13] D. Bertsimas, V. F. Farias, and N. Trichakis, "The price of fairness," *Operations research*, vol. 59, no. 1, pp. 17–31, 2011.