A Generic Approach to Multi-Fairness and Its Application to Wireless Channel Allocation

Mario Köppen, Kaori Yoshida, and Masato Tsuru
Kyushu Institute of Technology
680-4 Kawazu, Iizuka, Fukuoka 820-8502, Japan
Email: mkoeppen@ieee.org, kaori@ai.kyutech.ac.jp, tsuru@cse.kyutech.ac.jp

***Abstract*—In this paper we provide an extension of maxmin fairness to the case of multiple objectives and different agent preferences by the definition of the maxmin multi-fairness relation. This generic relation is based on a formal modification of maxmin fairness according to its implicit comparisons and includes maxmin fairness as a special case. The application of this multi-fairness relation to a selfish, to a collaborative and to a mixed community of agents is presented, as well as a case study on a specific task of fair division of indivisible goods by numerical simulations. All studies demonstrate the suitability of the proposed relation for investigative tasks in shared resource problems, as they are abundant in network design and control.**

***Keywords*-fairness; maxmin fairness; preference modelling; multi-fairness; fair division of indivisible goods; wireless channel allocation**

I. Introduction

For fair division of sharable goods, *maxmin fairness* is defined as a state where no agent can be made better off without making another already equally or worse off agent even more worse off. *Maxmin fairness* has been shown to maximally utilize resource sharing, for example for end-to-end user traffic rates in a communication network, where the *maxmin fair state* can be achieved by the Bottleneck Flow Control algorithm [1]. It can be rationalized as greatest or maximal element of the *maxmin fairness dominance relation*, in the same spirit as Pareto dominance relation can rationalize Pareto efficiency.

Definition 1. *Given a feasible space $A \subseteq R_n$. For two elements (vectors) x and y from A it is said that x* ***maxmin fair dominates*** *y ($x >_{mmf} y$) iff for each i with $y_i > x_i$ there is at least one $j \neq i$ such that (1) $x_i \geq x_j$ and (2) $x_j > y_j$.*

However, in practice, resource division is often not the result of a direct manipulation of the resource, but the result of an allocation or assignment. For example, in user end-to-end traffic in communication networks, the shared resource are links and nodes along with their capacities, but links and nodes cannot be directly assigned to users. This can only be achieved by assigning a routing, i.e. a set of paths linking users, and shared links and nodes will automatically follow. But such an allocation or assignment will also need observables, in order to be represented to the agent. To follow up the communication network example, in addition to putting boundaries onto traffic rates, there are also hop counts and delay times, once a routing is specified. We see that this situation cannot be represented by Definition 1.

The other point: different agents might consider the meaning of "better off" and "worse off" differently. The above specification of a *maxmin fairness* state is transparent to such differences, but Definition 1 does not reflect different agent preferences as well. Especially when there are several observables availabe, thus making the agent justification a multi-criterion task, it needs a clear specification of the meaning of these terms.

We will present a rather natural way to extend mamxmin fairness in a manner that (1) multiple observables can be taken into account, and (2) each agent can follow his or her own interpretation of "better off" and "worse off." For doing so, we start with an observation on the formal structure of Definition 1.

There, the use of $>, \geq$ is not confined to the use of the same relation. Essentially, three relations are taken into account when justifying *maxmin fair dominance*:

$$\begin{tikzcd}
x_1 & x_2 & \cdots & x_i \arrow[rr, Rightarrow, "\geq"] & \cdots & x_j \arrow[d, ">"] & \cdots & x_n \\
y_1 & y_2 & \cdots & y_i \arrow[u, ">"] & \cdots & y_j & \cdots & y_n
\end{tikzcd}$$

These relations are between y_i and x_i (a possible improvement for agent i), x_j and x_i (agent j, who is already equally or worse off compared to agent i) and x_j and y_j (again, agent j, who is made even more worse off, and thus would becomes an "envy" in state y). So, one relation is according to the model of agent i, and two according to the model of agent j. But there are three more relations, ensured by transitivity of the larger relation for real numbers: $x_i > y_j, y_i > y_j$ and $y_i > x_j$. These three are also seen from the perspective of agent j.

If we now replace the real number relations by "agent-specific" transitive vector relations $>_i$ and $>_j$, we get the following scheme:

$$\begin{tikzcd}
x_1 & x_2 & \cdots & x_i \arrow[rr, Rightarrow, "\geq_j"] & \cdots & x_j \arrow[d, ">_j"] & \cdots & x_n \\
y_1 & y_2 & \cdots & y_i \arrow[u, ">_i"] \arrow[urr, ">_j"] & \cdots & y_j & \cdots & y_n
\end{tikzcd}$$

Here, we have four relations that exactly corresponds to the relations of former scheme, and two others follow by transitivity (of $>_j$-relation): $y_i >_j y_j$ and $x_i >_j y_j$.

We will take this observation as a base to define a *maxmin multi-fairness* among multi-vectors. The implications of this approach will be investigated in the following. Section II will provide the formal definitions, and also introduce suitable agent relations $>_i$. Then, section III will provide a study on different settings for the agent relations in an unbiased feasible space, while section IV is focusing on a specific resource sharing problem. We will discuss related work and provide a conclusion in the last two sections of this paper.

The main contributions of present work can be seen in the formal specification of a multi-fairness, directly following *maxmin fairness* (and including *maxmin fairness* as a special case) with design flexibility according to specification of agent relations, and the demonstration of this design flexibility and its implications.

II. DEFINITIONS

We consider a multi-objective distribution problem, where an operator distributes a number of items (their number does not matter here) to a number n of agents (or users). A specific distribution Δ is characterized by m measurements and thus, it is represented by assigning an m-dimensional *objective vector* to each agent. By $(o_i)[\Delta]$ we denote the objective vector assigned to agent i. Its components, i.e. the single measurement values for each agent, are indicated by $(o_i)_j[\Delta]$[1] (we will not give the argument Δ if it is clear from context). For example

$$X = \left(\begin{pmatrix} 8 \\ 6 \\ 9 \end{pmatrix} \begin{pmatrix} 5 \\ 3 \\ 3 \end{pmatrix} \begin{pmatrix} 2 \\ 2 \\ 8 \end{pmatrix} \right)$$

represents an allocation of items to 3 agents, where for each agent the quality of the allocation is represented by three, possibly competing, measured objectives. So, the 3rd measurements qualifies the allocation represented by this multi-vector to agent 1 by the value 9, the highest among the measurements among all agents.

Now we also assume that to each agent i two individual vector relations $>_i$ and $\geq_i$ are assigned. There, the relations both are subsets of $A \times A$ with $A \subseteq R_n$, and $\geq_i$ is understood as the extension of $>_i$ by an equivalence relation. Thus, by virtue of these relations, each agent is able to compare each objective vector with each other (means: not only her "own" objective vector). Basically there are no other requirements on these relations except transitivity. Now we can define generic (maxmin) multi-fairness:

Definition 2. *For two given multi-vectors X and Y of same dimension with corresponding objective vectors (x_i) and (y_i), we say that X is multi-fair dominating Y ($X >_{mf} Y$) iff for each i with $(y_i) \geq_i (x_i)$ there is at least one $j \neq i$ with the following properties:*

1) $(x_i) \geq_j (x_j)$
2) $(y_i) >_j (x_j)$
3) $(x_j) >_j (y_j)$

[1] We do not use the notation o_{ij} to make clear that the set of objective vectors is not primarily a matrix.

By this definition we realize the comparison between two multi-vectors exactly the way as it was discussed in the Introduction. For convenience, we also introduce a few definitions that will be of help to specify the agent-relations $>_i$ and $\geq_i$.

Definition 3. *Given two objective vectors (x) and (y), we say that (x) is (strictly)* ***Pareto dominating*** *(y) ($(x) >_p (y)$) iff for each i $(x)_i \geq (y)_i$ and at least one j $(x)_j > (y)_j$ holds. Correspondingly, $(x) \geq_p (y)$ iff $(x) = (y)$ or $(x) >_p (y)$ (or alternatively: for each i $(x)_i \geq (y)_i$).*

Definition 4. *Given two objective vectors (x) and (y), we say that (x) is* ***marginal dominating*** *(y) by index i ($(x) >_{m(i)} (y)$) iff $(x)_i > (y)_i$. Correspondingly, $(x) \geq_{m(i)} (y)$ iff $(x)_i \geq (y)_i$.*

Given an index set $I = (i_1, i_2, \ldots, i_k)$ (where any $i_l \in \{1, 2, \ldots, m\}$) and an objective vector (x), we consider the *sub-vector* of dimension k of (x) by I as the vector composed of all components of (x) with indizes from I. The notation is $(x)_I$

Definition 5. *Given two objective vectors (x) and (y), we say that (x) is* ***subset-Pareto dominating*** *(y) by index set I ($(x) >_{p(I)} (y)$) iff $(x)_I >_p (y)_I$. Correspondingly, $(x) \geq_{p(I)} (y)$ iff $(x)_I \geq_{p(I)} (y)_I$.*

Obviously, $>_{m(i)}$ corresponds with $>_{p(\{i\})}$. By marginal dominance, we refer to a situation where an agent or the operator is comparing two allocations by the i-th objective alone. By subset-Pareto dominance, an agent or the operator compares by a subset of the objectives.

To give an example, consider two multi-vectors X and Y, where agent i is focussing on the i-th objective alone (i.e. $>_i$ is $>_{m(i)}$ and $\geq_i$ is $\geq_{m(i)}$):

$$X = \left(\begin{pmatrix} 3^* \\ 7 \\ 10 \end{pmatrix} \begin{pmatrix} 10 \\ 9^* \\ 10 \end{pmatrix} \begin{pmatrix} 5 \\ 3 \\ 3^* \end{pmatrix} \right) \quad Y = \left(\begin{pmatrix} 2^* \\ 5 \\ 6 \end{pmatrix} \begin{pmatrix} 8 \\ 2^* \\ 8 \end{pmatrix} \begin{pmatrix} 5 \\ 8 \\ 5^* \end{pmatrix} \right)$$

By the star *, we have indicated the value used in comparisons $>_i$ and $\geq_i$ for each agent's objective. We see that for agents 1 and 2, there is no improvement in Y: neither $(y_1) >_{m(1)} (x_1)$ (i.e. $(y_1)_1 > (x_1)_1$ or $2 > 3$) nor $(y_2) >_{m(2)} (x_2)$ (i.e. $(y_2)_2 > (x_2)_2$ or $2 > 9$) holds. But for $i = 3$, there is an improvement: $(y_3) >_{m(3)} (x_3)$ (i.e. $(y_3)_3 > (x_3)_3$ or $5 > 3$). So, we have to look if there is a $j \neq i$ fulfilling properties (1) to (3) of Definition 2. The choices for j are 1 or 2. In fact, $j = 1$ has these properties: (1) $(x_3) \geq_{m(1)} (x_1)$ (i.e. $(x_3)_1 \geq (x_1)_1$ or $5 \geq 3$); (2) $(y_3) >_{m(1)} (x_1)$ (i.e. $(y_3)_1 > (x_1)_1$ or $5 > 3$); and (3) $(x_1) >_{m(1)} (y_1)$ (i.e. $(x_1)_1 > (y_1)_1$ or $3 > 2$). So, for this example, agent 1 is "envy" about the potential improvement for agent 3, or $X >_{mf} Y$. It can also be seen that this is an example for a 2-cycle, since also $Y >_{mf} X$ holds (agent 3 is an envy for both improvements, the one of agent 1 from 2 to 3, and the one for agent 2 from 2 to 9).

Finally, we recall the concept of a maximum set (or non-dominated set, or set of maximal elements). We provide the simpler definition of Sen [2], as it is sufficient for the present analysis. For a more precise definition, which allows to specify rationalizability of a social choice, see [3]:

Definition 6. *Given a relation $>_r$ as subset of $A \times A$. Then, the **maximum set** of $>_r$, M_r is given by all $a^* \in A$ where there is no other $a \in A, a \neq a^*$ with $a >_r a^*$.*

In the following section, we will analyze the proposed generic multi-fairness with regard to different settings for the agent relations.

III. Comparative Study on User Preference Scenarios

A. Exclusive Marginal Preference Relations

In this setting, each agent is focusing exclusively on a single objective, thus $n = m, >_i \equiv >_{m(i)}$ and $\geq_i \equiv \geq_{m(i)}$. We want to investigate the frequency of the relation occurrence among unbiased multi-vectors, and their maximum sets. Therefore, we perform random sampling of multi-vectors by uniform sampling all components of all objective vectors from $[0, 1]$. Table I gives some results for increasing number of agents (and objectives). Pairs s_1 and s_2 of multi-vectors were sampled 10,000 times. It is shown how often the $>_{mf}$ relation with that setting of the agent relations appears between s_1 and s_2, but also, how often there was a 2-cycle between s_1 and s_2.

TABLE I
Frequency of multi-fairness relation and conflicts for the case that all user focus on different objectives among 10,000 pairs (s_1, s_2) of random samples of dimension n.

n	$s_1 >_{mf} s_2$?	$s_1 >_{mf} s_2 \wedge s_2 >_{mf} s_1$?
2	3382	145
3	2570	143
4	2374	159
5	2306	149
7	2933	467
10	2215	286
15	2911	542
20	3612	933
50	7966	6339
100	9883	9796

It can be clearly seen that with increasing n the relation becomes more and more *abundant*. For large n, nearly each random multi-vector is in relation to any other random multi-vector. However, for n below around 20, it appears to be a rather constant ratio (about 30%). The meaning is that for "selfish" agents, who all focus on their own objective, the chance of having an envy is approximating 1 for increasing number of agents. The number of value-based comparisons is fixed (we check four relations for testing multi-fairness) but the number of candidate envies is increasing. In such a community, it appears hard to achieve fairness.

Table II provides the corresponding analysis of the maximum sets for increasing number of objectives, or users. For the evaluation, a set of 100 random multi-vectors was sampled 100 times, and the maximum sets of the 100 multi-vectors were evaluated according to the minimum size, maximum size, average size, and number of cases where the maximum set was empty (as a consequence of the presence of 2-cycles). Also here, below 10 objectives, the relation appears to be efficient with regard to specifying a maximal state, but then, the size of maximum sets decreases rapidly, and the relation becomes less and less able to specify a fairness state among the agents. Since the relation becomes abundant, it becomes more and more likely that each random multi-vector is dominated by any other multi-vector.

TABLE II
Analysis of the size distribution of maximum sets M_{mf} of 100 random vectors of dimension n, sampled 100 times. Case: all users with different objective.

n	min	$\|M_{mf}\| = 0$?	max	average
2	1	0	9	3.7
3	2	0	11	5.61
5	2	0	10	6.47
10	0	3	9	3.83
20	0	79	2	0.29

B. Shared Sub-Pareto Dominance Relations

In this setting, we consider agents who are each focusing on two objectives, and share exactly one objective with exactly one other agent. So this can be considered a more collaborative community. Formally, all agents compare by sub-Pareto dominance relation: $>_i$ becomes $>_{p(\{i,i+1\})}$ for $i = 1, \ldots, (n-1)$ and $>_n \equiv >_{p(\{n,1\})}$ (and correspondingly for the $\geq_i$ relations). The same evaluations as for the former case were performed. Results for frequencies are given in Table III and for the maximum sets in Table IV.

TABLE III
Frequency of multi-fairness relation and conflicts for the case that all user focus on two objectives, two users share one objective, among 10,000 pairs (s_1, s_2) of random samples of dimension n.

n	$s_1 >_{mf} s_2$?	$s_1 >_{mf} s_2 \wedge s_2 >_{mf} s_1$?
2	5654	2497
3	4377	1250
4	3326	658
5	2581	354
7	1539	231
10	748	4
15	224	0
20	115	0
50	8	0
100	8	0

We note a completely different behavior to the former "selfish" setting of marginal objectives: with increasing number of users, the relation becomes *sparse*, and also the 2-cycles are rapidly vanishing. It has to be related to the fact that Pareto dominance is not a total relation, i.e. there are pairs of vectors x and y where neither $x >_p y$ nor $y >_p x$ holds.

The sparseness of a relation has the advantage that empty maximum sets become more and more unlikely, as can be easily seen in Table IV, but also the disadvantage that they will increase in size, and thus that the relation allows less and less to specify a fair state. For 20 agents, about 50% of random multi-vectors already belong to the maximum set.

TABLE IV
ANALYSIS OF THE SIZE DISTRIBUTION OF MAXIMUM SETS M_{mf} OF 100 RANDOM VECTORS OF DIMENSION n, SAMPLED 100 TIMES. CASE: ALL USERS WITH TWO OBJECTIVES, TWO USERS SHARE ONE OBJECTIVE.

n	min	$\|M_{mf}\| = 0$?	max	average
2	0	95	1	0.05
3	0	87	2	0.15
5	0	31	3	1.02
10	7	0	19	11.41
20	36	0	55	45.08

C. Mixed-Marginal Preference Relations

As a third setting, we consider the sharing of two marginal relations among all agents by different ratios. It means that there are n user, two objectives, and for a ratio r, $n{\cdot}r$ (rounded) agents use the relations $>_{m(1)}$ and $\geq_{m(1)}$, and the remaining agents use the relations $>_{m(2)}$ and $\geq_{m(2)}$. Results of the same analysis as in the cases before are given in Table V for the frequency and in Tables VI, VII and VIII for the maximum sets, all for three choices of the ratio r of agents using objective 1 to agents using objective 2.

TABLE V
FREQUENCY OF MULTI-FAIRNESS RELATION AND CONFLICTS FOR THE CASE THAT A SHARE r OF USERS FOCUSSES ON ONE MARGINAL OBJECTIVE, $(1-r)$ ON A SECOND, AMONG 10,000 PAIRS (s_1, s_2) OF RANDOM SAMPLES OF DIMENSION n. "REL?" STANDS FOR THE TEST $s_1 >_{mf} s_2$ AND "2-CYCLE?" FOR THE TEST $s_1 >_{mf} s_2 \wedge s_2 >_{mf} s_1$.

r	0.1		0.3		0.5	
n	rel?	2-cycle?	rel?	2-cycle?	rel?	2-cycle?
2	3374	0	3350	129	3344	126
3	2498	0	2545	97	2571	117
4	1957	0	2137	75	2180	82
5	1810	63	1934	65	1931	68
7	1359	25	1620	66	1572	67
10	1176	37	1431	75	1447	64
15	952	35	1116	64	1287	74
20	778	23	1029	57	1126	72
50	465	8	797	34	861	34
100	391	0	679	20	729	33

The frequency of occurrence clearly shows that this relation is a good trade-off between the former two cases, the relation neither becomes abundant nor sparse. However, for larger number of users, the relation frequency is also decreasing, but more slowly than in case of sub-Pareto dominance. The number of 2-cycles seems also to decrease with large n. An interesting observation is that for the case $r = 0.1$ (where the multi-fairness is more close to "standard" *maxmin fairness*, as the majority of agents is focusing on the same objective) this value has a maximum for around 10 agents. For different values of r, nevertheless, there are no extreme differences between these frequencies.

The evaluation of the maximum set sizes confirms this picture of a "trade-off" relation. There is no strong dependency from r and the sizes of maximum sets are moderate, between 1 and 10 in most cases. As the average size is always close to the average of maximal and minimal size appearing in the 100 samples, it also indicates a more normal distribution of maximum set sizes. We also note that except one case, the maximum set was always non-empty.

TABLE VI
ANALYSIS OF THE SIZE DISTRIBUTION OF MAXIMUM SETS M_{mf} OF 100 RANDOM VECTORS OF DIMENSION n, SAMPLED 100 TIMES. CASE: 10% OF USERS FOCUS ON ONE OBJECTIVE, 90% ON A SECOND OBJECTIVE.

n	min	$\|M_{mf}\| = 0$?	max	average
2	1	0	2	1.48
3	1	0	3	2.02
5	1	0	9	5.24
10	1	0	14	6.23
20	3	0	14	7.92

Thus, sharing of the same objectives (or comparison relation) among agents, more or less independent of the sharing ratio, appears the optimal way to use the proposed multi-fairness in order to solicit a fair distribution state.

TABLE VII
ANALYSIS OF THE SIZE DISTRIBUTION OF MAXIMUM SETS M_{mf} OF 100 RANDOM VECTORS OF DIMENSION n, SAMPLED 100 TIMES. CASE: 30% OF USERS FOCUS ON ONE OBJECTIVE, 70% ON A SECOND OBJECTIVE.

n	min	$\|M_{mf}\| = 0$?	max	average
2	0	1	7	3.46
3	1	0	8	4.01
5	1	0	10	5.04
10	1	0	8	4.97
20	1	0	7	4.51

TABLE VIII
ANALYSIS OF THE SIZE DISTRIBUTION OF MAXIMUM SETS M_{mf} OF 100 RANDOM VECTORS OF DIMENSION n, SAMPLED 100 TIMES. CASE: 50% OF USERS FOCUS ON ONE OBJECTIVE, 50% ON A SECOND OBJECTIVE.

n	min	$\|M_{mf}\| = 0$?	max	average
2	1	0	8	3.6
3	1	0	10	4.01
5	2	0	12	5.44
10	1	0	11	5.71
20	1	0	9	3.81

D. Discussion

The evaluations for different settings have shown that in fact multi-fairness can behave much differently, and depends strongly on the type and distribution of agent relations. While the qualitative results might be rather obvious for the various settings (for example that under a condition where each agent is focusing on his own objective, envy-free fairness is hard to achieve), the additional specification of a relation allows for further quantitative analysis (to continue the example: however, for a limited number of agents, below 10, fairness can be sufficiently achieved). This can be seen as an advantage of the proposed concept.

The appearance of 2-cycles (not to mention k-cycles with $k > 2$) is a clear disadvantage. For *maxmin fairness*, such cycles are impossible. However, the analysis has also shown that the effect on the existence of maximal elements is also strongly differing, and that there are settings where the chance to have an empty maximum set is practically 0. Moreover, the relation appears to be *scalable*, means that there are ways to modify the representation of agent preferences, expressed by their relations, in order to enforce sparseness of the relation and thus to increase the size of maximum sets.

IV. Case Study: Wireless Channel Allocation

In order to further demonstrate the investigative potential of the proposed generic multi-fairness, we provide a study on the problem of allocation users to channels and timeslots in a wireless communication network. Here, a set C of m cells c_i with $i = 1, \cdots, m$ is given. Each cell represents a channel number and a timeslot, but with regard to the combinatorial nature of the problem, we will ignore this matrix structure and focus on a linear array of cells. There are n users $U = \{u_i \,|\, i = 1, \cdots, n\}$ who want to employ the wireless network. However, each user can employ each cell differently - behind this we can find the physical circumstances of the network, the device and its motion that can be summarized into an utility factor (sometimes also called channel coefficient) represented by a factor f_{ij} (usually a real from $[0, 1]$). Multiplication of the available bandwidth (assumed to be equal for all cells) with that factor gives the usable bandwidth for the particular user u_i, if the cell c_j is allocated to her. The task of the operator is to assign exactly one user to each available cell, i.e. to specify a mapping $a : C \to U$ by which cell c_i is assigned to user $a(c_i)$ (we will call it allocation a and being represented as an ordered m-tupel of users). Then, a user u_i will receive a total utility F_i as sum of the utility factors of all cells that are assigned to her:

$$F_i(a) = \sum_{j, a(c_j) = u_i} f_{ij} \tag{1}$$

There are n^m possible allocations, and the task is to describe an efficient procedure to select one allocation among all allocations under given circumstances[2]. This is the Wireless Channel Allocation (WCA) problem, which appears to be a special case of distribution of indivisible goods (one cell can only be assigned to one user). Several approaches to this problem have already been presented, for example based on auction systems [4] or by direct application of a relational concept of *maxmin fairness* [5]. It can be seen that global optimization can produce unwanted situations. For example, if the operator is considering to maximize the total bandwidth usage, i.e. $\sum_i F_i(a)$, then the solution is to assign the cell to one of the users with largest utility factor:

$$a(c_i) = u_j \text{ where } j = \arg\max_k [f_{ki}] \tag{2}$$

However, in cases like

$$F = \{f_{ij}\} = \{\{0.1, 0.2\}, \{0.3, 0.4\}\}$$

this could lead to the selection of the same user for all cells (user 2 in this example). Then, all other users would not receive any bandwidth.

[2]Note that we are using the term *efficient* here in the sense of efficient cause: the environmental situation, abstracted by the model and its parameters, should always be sufficient to entail a unique allocation.

An agent with total utility 0 can be considered as an envy agent (or envy for short), so the goal of the WCA becomes to find an envy-free allocation. We can use *maxmin fairness* to select an allocation, as well as the proposed *maxmin multi-fairness*. The latter would allow to take user preferences into account. For studying this option, we consider a small scaled problem, because otherwise an exhaustive analysis would not be possible due to the combinatorial explosion of larger-scaled problems.

Assume there are 6 cells, and 3 users. Thus there are $3^6 = 729$ possible allocations. Two users have a preference for sending traffic more early, so the total utility achieved in the first three cells is weighted higher than for the remaining three cells. User 3, on the other hand, is fine with sending traffic later (maybe also by the consideration that more users will compete for the first cells). We assume that the preference is expressed by a weighted average with weights 0.8 and 0.2. So we can specify two objectives:

$$\begin{aligned} L_i(a) &= \sum_{j, a(c_j) = u_i, j \leq 3} f_{ij} && (3) \\ H_i(a) &= \sum_{j, a(c_j) = u_i, j > 3} f_{ij} \\ (o_i)_1(a) &= 0.8 L_i(a) + 0.2 H_i(a) \\ (o_i)_2(a) &= 0.2 L_i(a) + 0.8 H_i(a) && (4) \end{aligned}$$

Users 1 and 2 are comparing allocations with regard to the marginal-larger (or marginal-$\geq$) relation by the first element, and user 3 does the same by the second argument of the objective vector. So we have specified everything needed to find the maximum set of the corr. multi-fairness relation. Then, we want to explore the trade-off for each user caused by taking her preferences into account, by comparison with the corresponding objectives for the maximum set of maxmin fairness (which does not take any preferences into account). The evaluation is as follows:

1) A random instance of a WCA problem for 6 cells and 3 users was created 500 times. Each time, the utility factors f_{ij} for each user and each cell were taken as a uniform random number from $[0, 1]$.
2) For each instance, the maximum set M_{mmf} of *maxmin fairness* was computed, i.e. the set of all allocations a_* such that there is no other allocation a with $F(a) >_{mmf} F(a_*)$.
3) For each such a_*, the objective vectors (o_i) for all three users were computed.
4) For each WCA instance, the maximum set M_{mf} of multi-fairness with above specification of the relations was computed.
5) For all WCA instances, the objective vectors for each element of M_{mmf} were compared to the objective vectors for each element of M_{mf} with regard to the following criteria (with a_{mmf} an allocation from M_{mmf} and a_{mf} an allocation from M_{mf}): the gain in objective

TABLE IX

RESULTS FOR COMPARING THE MAXIMUM SETS OF MAXMIN FAIRNESS AND MULTI-FAIRNESS (VECTORS INDICATE THE CORR. VALUES FOR USERS 1, 2 AND 3), WEIGHTS ARE 0.8 AND 0.2.

Criterion	Value
Number of comparisons	5744
How often could users improve their objective?	$(62\%, 60\%, 87\%)$
How much could users gain for their objective?	$(0.061, 0.041, 0.318)$
How often did users have a loss in total utility?	$(68\%, 68\%, 48\%)$
How large was the loss in total utility?	$(0.239, 0.251, -0.005)$
Empty maximum set for multi-fairness	6 times

TABLE X

RESULTS FOR COMPARING THE MAXIMUM SETS OF MAXMIN FAIRNESS AND MULTI-FAIRNESS (VECTORS INDICATE THE CORR. VALUES FOR USERS 1, 2 AND 3), WEIGHTS ARE 0.6 AND 0.4.

Criterion	Value
Number of comparisons	3028
How often could users improve their objective?	$(67\%, 67\%, 74\%)$
How much could users gain for their objective?	$(0.007, 0.001, 0.039)$
How often did users have a loss in total utility?	$(44\%, 45\%, 43\%)$
How large was the loss in total utility?	$(0.051, 0.054, 0.033)$
Empty maximum set for multi-fairness	38 times

1 for users 1 and 2, or objective 2 for user 3, in a_{mf} against a_{mmf}; the count of occurrences of the event that objective 1 for users 1 and 2, or objective 2 for user 3 in a_{mf} is larger or equal to the corresponding objective in a_{mmf}; the loss in total utility (i.e. objective 1 plus objective 2) in a_{mf} against a_{mmf} for each user; the count of occurrences of the event that there is such a loss; the number of cases where the maximum set of the multi-fairness was empty.

6) The sum of all these values where divided by the sum of $|M_{mmf}| \times |M_{mf}|$ to get their average over the 500 random WCA instances.

The results of these evaluations are shown in Table IX.

The results indicate a trade-off: users have a chance to improve their preferred objective, but may loose in total utility. The evaluation of the results in more detail:

- The maximum sets are rather small in all cases. 5744 comparisons indicate an average size of about 3 elements. Unfortunately, in a few cases, the multi-fairness does not allow for specifying maximal elements.
- Users 1 and 2 have a roughly 60% chance to improve their objective (i.e getting a larger allocation in the first 3 cells), while user 3 has a higher chance, close to 90%. This indicates that users 1 and 2, who share the same objective, in fact *compete* for their objective. If both, user 1 and user 2, have a chance of 60%, then this means that in 84% of the cases (i.e. with probability $1-(1-0.6)^2$), user 1 *or* user 2 will improve, and this value is close to the value for user 3 (who does not share her objective).
- The fact that user 3 does not share the objective seems to give other benefits as well: the average gain for user 3 is about 5 times larger than for users 1 or 2, it looses less frequently in total utility, and the average loss in total utility is close to 0.

To validate the last observation, the corresponding results for the weights 0.6 and 0.4 instead of 0.8 and 0.2 are given in Table X. Now all user start to focus on more similar objectives, and in general, the relation selects in a manner more similar to the *maxmin fairness*, in particular:

- The maximum sets are getting smaller, and the number of empty maximum sets is increasing.
- The gains in objective values get close to 0 (but are still larger for user 3).
- The benefits for user 3 are getting smaller: the chance of loss in total utility is now about 50% for all users, while the amount of loss is nearly 0 in all cases.

So we can confirm a tendency towards *maxmin fairness* for increased similarity of the objectives[3].

In summary, the evaluation of the example case allows for a better understanding of fairness with regard to competing user preferences: user 3 has a clear benefit from her focus on a different (actually exclusive) objective, while users 1 and 2 have to share the benefits as a result of the focus on the same objective. There is also a trade-off for including preferences: a higher chance of increase in an objective is accompanied by a higher chance of loss in the total of objectives. Last but not least, both facts also confirm the further characterization of fairness relations as extremality relations.

V. RELATED WORK

The most closest topic that has been treated in related works is the fair distribution of multiple resources. Recently, Dominant Resource Fairness (DRF) has been proposed [6][7]. There, *maxmin fairness* is extended to the case of distribution of several resources like CPU time and memory demand at the same time. For each agent (task), the maximal share of any resource is selected, and *maxmin fairness* is seeked for these so-called dominant shares. The method is compared to the popular asset fairness and the Competitive Equilibrium from Equal Incomes (CEEI) approach of microeconomics theory [8][9][10] by a catalogue of criteria for such resource sharing regimes. In asset fairness, an equal valued share for each resource has to be identified. Then, assigning equal shares to each agent is formulated as a combinatorial optimization problem. In CEEI, all agents start from an equal share of $1/n$ and then start to perform mutual transactions in order to achieve a Nash equilibrum. The main difference of these approaches and the present one is that resources of multiple qualities can be independently manipulated and that agent intervention is possible in this process. In present approach, the objectives represent different aspects of the same resource

[3]With the exception of the larger number of empty maximum sets.

sharing. For example, in the specification of a routing in end-to-end user traffic in a data network, the same routing gives raise to different objectives like delay time, number of hops, and *maxmin fair* traffic rates. An infrastructure property like delay cannot be subject of a resource sharing alone, since a delay cannot be "shared" or distributed among agents. However, it can be combined with other sharable resources, put under the same umbrella as an additional objective in the presented relation-based framework. Moreover, objectives cannot be manipulated independently. The disadvantage here is that agents cannot directly intervene in the process, and therefore, our focus was giving to the central agency of an "operator". But therefore, we also do not have the problem of deception by agents, probably collaborating, giving wrong information about their demands in order to yield higher shares of resources. It also has to be noted that DRF can be seen as a similar formal extension of the leximin relation, by using the comparison of maxima of objective vectors as agent specific relation - thus, the DFN approach could basically follow the same design principle as our approach. Otherwise, we are not aware of any other direct attempt to generalize *maxmin fairness* to the case of multiple objectives.

The use of relational mathematics in fair division problems is more common. In [11], Bouverte et al. study a *fair division problem of indivisible goods without money transfers*. Here, in addition to the agents and goods, to each agent a preference relation is assigned. The fair division problem then is formulated as the task to achieve Pareto-efficiency and envy-freeness. The innovation in this work is the transition to propositional logic and formal language theory, and also allows for example for analysis of complexity of the fair division problem. In comparison, our approach extends Pareto efficiency (which often appears to be not achievable) to *maxmin multi-fairness* in the same way as Pareto-equilibrum is extended to the *maxmin fair state* by giving an extra focus on agents that already receive less than other agents. Thus we assume that it is possible to handle the accompanying problems within the framework of relational mathematics alone. Complexity analysis and a formal "logic of fairness" are beyond the scope of the present approach, but have to be seen as important questions for future research.

VI. Summary, Conclusions, and Future Work

The extension of *maxmin fairness* to the case of multiple objectives (different observations of the allocation of a shared resource) and multiple agent models (agent-specific relations between the objective vectors) has been introduced, and the implications for various settings of this multi-fairness have been explored. The specific settings were (1) the "selfish community" where each agent is exclusively focusing on a single objective (and where the multi-fairness comes out to be hard to achieve with increasing number of agents), (2) the "collaborating community" where each agent shares an objective's preference with exactly one other agent, and where fairness comes out to be abundant and underspecified, and (3) the mixed setting, where a ratio of agents shares one objective, and the other agents share another objective: this appears to be the trade-off between settings (1) and (2) and allows for efficient specification of fairness. Furthermore, the multi-fairness was also studied on a specific resource sharing task, the Wireless Channel Allocation problem, with a problem scale that allowed for an exhaustive analysis. The trade-off between overlapping user-preferences, as well as between individual user preferences and the indifferent situation of *maxmin fairness* could be demonstrated as well. Multi-fairness allows for a quantitative analysis of qualitatively given user preferences in a convenient way.

One main point has been left open: the task of finding the maxmimum set of the relation in the general setting. Here, only random sampling and exhaustive search have been used as means for investigation. Given the complexity of the search problem, we cannot expect that easy and tractable algorithms like the Bottleneck Flow Control for *maxmin fairness* can be found soon. Instead of this, we promote the use of meta-heuristic search algorithms, as the potential of these algorithms has already been demonstrated for the case of searching the maxmimum set of the Pareto dominance relation by so-called Evolutionary Multi-objective Algorithms, and also been demonstrated for the case of *maxmin fairness* [12].

VII. Acknowledgments

This study was supported in part by the JSPS, KAKEN-HI (B) (21300024).

References

[1] D. Bertsekas and R. Gallager, *Data Networks*. Englewood Cliffs, N: Prentice Hall, 1992.
[2] A. K. Sen, *Collective Choice and Social Welfare*. Holden-Day, San Francisco, 1970.
[3] K. Suzumura, *Rational Choice, Collective Decisions, and Social Welfare*. Cambridge University Press, 2009.
[4] J. Sun, E. Modiano, and L. Zheng, "Wireless channel allocation using an auction algorithm," *IEEE Journal on Selected Areas in Communications*, vol. 24, no. 5, pp. 1085–1096, 2006.
[5] M. Köppen, R. Verschae, K. Yoshida, and M. Tsuru, "Heuristic maxmin fairness for the wireless channel allocation problem," in *Broadband, Wireless Computing, Communication and Applications (BWCCA), 2010 International Conference on*, 2010, pp. 606 –611.
[6] A. Ghodsi, M. Zaharia, B. Hindman, A. Konwinski, S. Shenker, and I. Stoica, "Dominant resource fairness: Fair allocation of multiple resource types," in *Proc. NSDI '11: 8th USENIX Symposium on Networked Systems Design and Implementation, Boston, MA, 2011*, 2011, pp. 323–336.
[7] ——, "Dominant resource fairness: Fair allocation of multiple resource types," Technical Report UCB/EECS-2011-18, EECS Department, University of California, Berkeley, Tech. Rep., March 2011.
[8] H. Varian, "Equity, envy, and efficiency," *Journal of Economic Theory*, vol. 9, no. 1, pp. 63–91, 1974.
[9] H. Young, *Equity: in theory and practice*. Princeton University Press, 1994.
[10] H. Moulin, *Fair Division and Collective Welfare*. The MIT Press, 2004.
[11] S. Bouveret and J. Lang, "Efficiency and envy-freeness in fair division of indivisible goods: Logical representation and complexity," in *Proceedings of the 19th International Joint Conference on Artificial Intelligence (IJCAI-2005)*, 2005, pp. 935–940.
[12] M. Köppen, R. Verschae, K. Yoshida, and M. Tsuru, "Comparison of evolutionary multi-objective optimization algorithms for the utilization of fairness in network control," in *Proc. 2010 IEEE International Conference on Systems, Man, and Cybernetics (SMC 2010), Istanbul, Turkey*, October 2010, pp. 2647 –2655.