Evolutionary Intelligence manuscript No.
(will be inserted by the editor)

Meta-Heuristic Approach to Proportional Fairness

Mario Köppen · Kaori Yoshida · Kei Ohnishi · Masato Tsuru

Received: date / Accepted: date

Abstract Proportional fairness is a concept from resource sharing tasks among n users, where each user receives at least $1/n$ of her or his total value of the infinitely divisible resource. Here we provide an approach to proportional fairness that allows its extension to discrete domains, as well as for the direct application of evolutionary computation to approximate proportional fair states. We employ the concept of relational optimization, where the optimization task becomes the finding of extreme elements of a binary relation, and define a proportional fairness relation correspondingly. By using a rank-ordered version of proportional fairness, the so-called ordered proportional fairness, we can improve the active finding of maximal proportional fair elements by evolutionary meta-heuristic algorithms. This is demonstrated by using modified versions of the Strength Pareto Evolutionary Algorithm (version 2, SPEA2) and Multi-Objective Particle Swarm Optimization (MOPSO). In comparison between proportional and ordered proportional fairness, and by using relational SPEA2, the evolved maximum sets of ordered proportional fairness achieve 10% more dominance cases against a set of random vectors than proportional fairness.

1 Introduction

Optimality in resource sharing tasks is a problem that appears in many application fields. One characteristic of such problems is that efficiency alone, expressed for example by maximality of a real-valued objective function, often not suffices to establish a suitable set of solutions. In addition to the inherent multi-objective nature of such problems, there is also a requirement of a fair usage policy for the resource utilization. While a solution where all resources are allocated to a single agent are still belonging to the Pareto optimal set, such allocations would exclude all other agents from the resource access. Therefore there is also the need for a formal concept of fairness that helps to effectively exclude such solutions. This problem might first have become apparent in the design and control of data communication networks. Maxmin fairness for routing and rate control in wired networks with link capacity constraints was introduced in [1]. It might be one of the first prominent approaches to fairness in a specific resource sharing task and extends the concept of maximization of the smallest allocation. A number of problems with maxmin fairness, especially problems caused by the rather strict focus on small allocations, led to the consideration of proportionality in the allocation as an alternative means for expressing fairness. Kelly et al. proposed proportional fairness in [2] as a state that maximizes the total utility of rate control for elastic traffic in a resource sharing communication network. This has become the most prominent approach to fairness so far.

However, there are still two problems with proportional fairness. One is a lack of an exact algorithm to generally find the proportional fair state within a set of feasible states. As will be detailed in the next section, attempts to provide an algorithmic approach to proportional fairness are usually accompanied by a redefinition of proportional fairness. The most common approach is the replacement of proportional fairness with the sum

All Authors are with
Kyushu Institute of Technology, Network Design and Research Center (NDRC)
680-4 Kawazu, Iizuka, Fukuoka 820-8502, Japan
E-mail: mkoeppen@ieee.org, kaori@ai.kyutech.ac.jp, ohnishi-@cse.kyutech.ac.jp, tsuru@cse.kyutech.ac.jp

of utilities and thus specifying a convex optimization problem. The other problem is the extension of proportional fairness to a discrete domain (i.e. sharing of indivisible resources). When wired communication evolved into wireless communication, problems of sharing indivisible resources prevailed. A main aspect of this problem is the appearance of discrete feasible spaces where no proportional fair state exists.

Here we are going to approach both target problems, by way of specifying proportional fairness as a set-theoretic relation between any pair of feasible allocations x and y, written as $x > y$ and literally read as "with regard to fairness, y appears (proportionally) unfair compared to x," and specifying a corresponding concept of maximality of proportional fair states. Maximality is understood as a state to which no other state is in proportional fair dominance relation. This allows to employ the concept of proportional fairness to any feasible domain as subset of an n-dimensional real space, including finite sets.

It comes out that by considering proportional fairness as a relation (and not primarily a stable state) then also the algorithmic problem can be handled. Recent developments in the field of evolutionary multi-objective optimization (EMO) provide a rich resource of algorithms to approximate Pareto fronts. But seen from relational algebra point of view, Pareto fronts are maximum sets of the Pareto dominance relation. Once proportional fairness is defined as a set-theoretic relation, the task of maximality corresponds with the search for Pareto fronts, after substituting Pareto-dominance relation with the proportional fairness relation (as well as any other relation). By such a substitution, standard EMO algorithms can be modified to make them capable of searching maximum sets of the proportional fairness dominance relation.

For further improving the meta-heuristic search, we also propose ordered proportional fairness as a modification of proportional fairness by ordering the elements of the comparison vectors by their numerical values before comparing. As will be seen, this relation is implied by proportional fairness. This means that maximum sets of ordered proportional fairness are subsets of maximum sets of proportional fairness. Then, EMO algorithms can gain performance from the fact that ordered proportional fairness is more prevalent than proportional fairness. With regard to Pareto dominance, such a method has already been applied in [3] by definition of so-called L-optimality, which is implied by Pareto dominance.

The experiments in this paper are focusing on the wireless communication as application domain. However, these days, we find fairness concepts in many other disciplines, for example in the design of voting systems or the so-called "cake cutting problems" [4] as a general metaphor for a class of problems from social choice theory. It has numerous applications, for example in traffic logistics [5], processor scheduling [6], queue management, multi-class classifier training, and even for harmonics in a musical score [7], and further applications can be envisaged.

The following Section 2 will provide the comments made so far in more detail. Then, Section 3 provides the basic definitions of relations, and a number of their properties are discussed. In Section 4, Monte Carlo simulation results for the proposed relation will be presented. Section 5 presents a comparative study of meta-heuristic approaches to ordered proportional fairness for a problem with a discrete domain. Some conclusions will be drawn in Section 6.

2 Background

2.1 Proportional Fairness

Proportional fairness maximizes the total utility of rate control for elastic traffic in a resource sharing communication network [2]. The proportional fair state is characterized as a state vector x of n positive-valued traffic rates such that for any other state $y \in R_+^n$ the inequality holds:

$$\sum_{i=1}^{n} \frac{y_i - x_i}{x_i} \leq 0 \tag{1}$$

In its original version, proportional fairness was established to model the allocation of traffic rates under a cost model. It reflects the existence of a pricing model that balances the user's willingness to pay differently for different traffic rate allocations and the network's "Best Effort" (BE) promise to optimally allocate the traffic rates weighted by such user payments. This is in contrary to the common "Quality of Service" (QoS) model, where users express minimal traffic rate demands that they wish to be guaranteed. In the latter case, network management needs to set up a "Call Admission Control" for deciding on the possible fulfillment of the QoS. Then, prices usually appear in the boundary conditions of a global optimization task. In BE traffic, the assumption is that users have their own utility functions, covering a range of different demands, while the network is using logarithmic utility. At the end, in [2] it is shown that under some mild conditions it suffices to handle a payment-independent problem, and leave the price negotiation task separated from the network through a network-user interface. It means the network can always operate in a "fixed" BE manner to maximize the network utility without knowing about user payment incentives, as it can be ensured that there is always a fitting price model. This separated task is to achieve *proportional fairness*.

Kelly's seminal work [2] gave a big stimulus to the network control research and found numerous applications. But also, researcher started to expand the concept to more complex network models. We will list a few examples:

In [8], authors study a mixed traffic scenario, where two kinds of price models for traffic rate allocations appear: according to "Best Effort" (BE, proportional fairness applies) and according to "Guaranteed Performance" (GP, a model case for QoS). They study different advanced price models and their mutual influence on BE and GP traffic rates, and from this derive reasonable bounds on prices.

In [9] and subsequent works, authors study a steady-state aspect of resource allocations. In a scenario where users receive a lower bandwidth by some fairness consideration, these users will also tend to stay longer in the system, thus establishing a steady-state of lower bandwidth allocations. The authors consider the general concept of *balanced allocations* under network utilization, where stochastic processes model user activity. The goal is to make the steady-state not depending on specific traffic characteristics. Conditions are identified where maxmin fairness and proportional fairness are balanced. Since these conditions are usually not fulfilled, an alternative notion of *balanced fairness* is introduced. For any flow state, balanced fairness ensures allocation on the boundary of the feasible domain (there given by capacity constraints on traffic flows for a queuing network of processors). The state can be approximated by an iterative procedure.

Ongoing studies in this direction take up the theme of steady-states and cover specific aspects of the network architecture, other networking domains, or improved algorithms to approximate fair states. In [10] the networking domain is file distribution in a file sharing network, and for the same model assumptions as in [9] an algorithm is provided that finds unbounded proportionally fair allocations for the model case of a three-link star network.

Recently, so-called utility fairness has gained some attention. Utility fairness was initially proposed in [11]. In this approach, the fairness conditions are (numerically) based on utility of the traffic rates instead of using traffic rate values directly, along with some necessary adjustments of the formal expressions. With regard to proportional fairness, on first glance the approach seems strange: wasn't it a main contribution of proportional fairness to separate the price model (representing user utility) from the network control? However, the approach seems to have some benefits with regard to congestion control, valuating further explorations of the concept.

For wireless networking, the task often becomes of discrete nature. In [12] authors study cooperative proportional fairness with regard to scheduling of (several) base stations to mobile stations. The utility now becomes dependent on the scheduling, i.e. assignment of base stations to mobile stations. The authors study the utility maximization task and provide a gradient-descent based approximation algorithm to the fair state. The same model can be applied to a wireless network with relays instead of the multiple base stations.

In [13] the strong relation between Nash bargaining and proportional fairness for convex domains is exploited to transfer results from game theory to fairness theory, including a concept for proportional fairness in special non-convex domains.

An information theoretic approach to α-fairness (another formal generalization of proportional fairness) is provided in [14], which is also one of the few works to seek an alternative characterization of proportional fair states.

Many of these works share a common aspect: the goal to express proportional fairness or its expansion in a "computable" fashion, and then to expose a unique state of the network by way of some kind of "computational optimization." Often, we find this reflected as maximizing the product of traffic rates (or other corresponding magnitudes) or the equivalent maximization of the sum of logarithms. This is somehow different from the condition given in Eq. (1).

The examples shown in Fig. 1 may help to explain the reasoning here. With regard to the fair state x, the condition of Eq. (1) for x corresponds to all points y below the tangent to the curve $\prod_i x_i = C$, with C being a constant. If the feasible space F is a convex subset of the positive quadrant, Fig.1(a) shows that the point with maximal product of components, in this example case the point (4,3), coincides with a point where for all other $x \in F$ the condition $(x_1 - 4)/4 + (x_2 - 3)/3 \leq 0$, i.e. the proportional fairness condition holds. Thus, for convex feasible spaces it is equivalent to say that the product of states is maximized, or that the whole feasible space is located below the tangent on the product function isoline at the product maximizing state.

However, the situation changes if the feasible space becomes a finite set of points. A related example is shown in Fig. 1(b). Now, the product maximizing state is not necessarily the same as the set of all points on or below that tangent. The location marked as point A gives an example where the product of components is

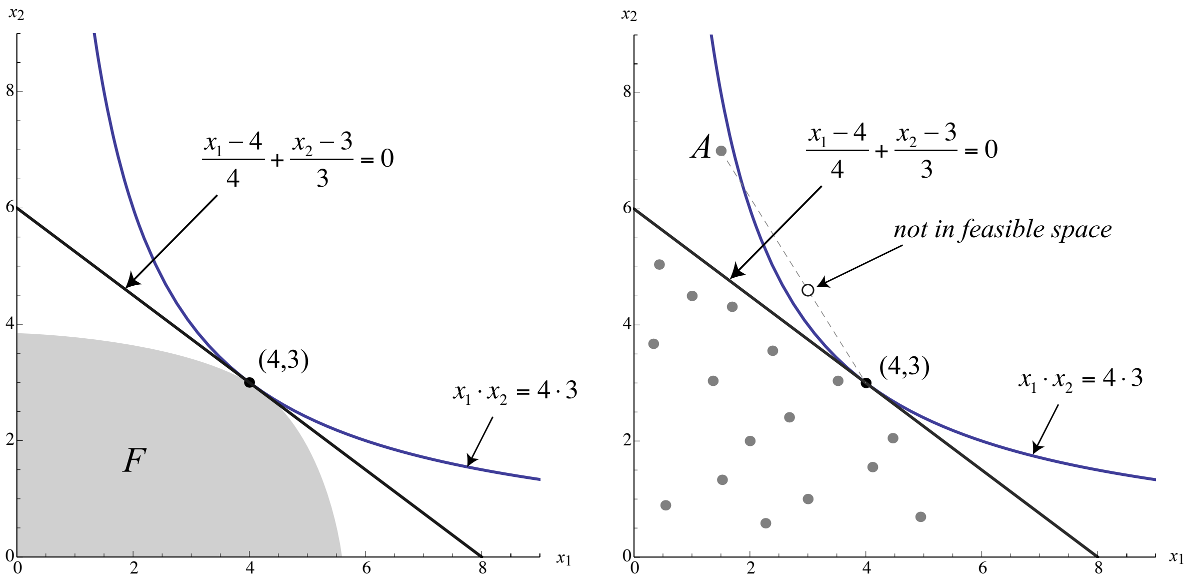

(a) continuous case (b) discrete case

Fig. 1 Proportional fairness for different feasible spaces.

smaller than the maximum value $4 \cdot 3$ among all feasible points, but the point is above the tangent. The reason is that only in a convex feasible space, all points on a line connecting A and $(4,3)$ would belong to the feasible space as well. On this line, there would be points with a larger product than $4 \cdot 3$ – but in the discrete case, they do not belong to the feasible space. Thus, blindly following the requirement of proportional fairness given by Eq. (1) for a discrete domain, there would not always be such a proportional fair state. However, there is always a state with maximum product of components. This is one fact that explains the preference for product maximization as proper definition of a proportional fairness state in expanded application domains (note that the same problem appears for non-convex feasible spaces).

Two other can be shortly mentioned as well. One is that the maximum product is reached for the same state as the maximum sum of logarithms, and this is directly reflecting the total utility, since utility is commonly modeled by the logarithm function (a way to make utility "currency-independent"). So it seems natural to focus on maximizing total utility. The other fact is related to computability. While the product maximization gives a clear objective what to compute and how to compare different states and solutions, the same can't be said about the proportional fair state condition. How can the fulfillment of a condition be used to efficiently compute something other than checking its truth value?

Nevertheless, in [2] and subsequent works we find proportional fairness stated by the tangential condition and not as product maximization. We are not aware of any clear statement made in these works about the reason to make this choice, but at least two arguments can be easily imagined: one is about the bounded trade-off - a condition that seems natural with regard to any model of fairness. If using product maximization, the maximal point can still decide about "un-fairness" of points with infinite components, and the relation would be too far-reaching. On the other hand, the set below the tangent is bounded. The other argument is to allow a more direct comparison with the concept of maxmin fairness [1], where the focus is on maximizing the smallest state. Thus, maxmin fairness tends to ignore improvements of states that are not minimal to any degree, while the proportional fairness condition is to some degree tolerating a decrease of the smallest state, as long as the sum of all relative changes of states remains positive (expressed by the terms $(y_i - x_i)/x_i$).

So, there is some reasoning to prefer the proportional fairness condition as provided by [2], and in the following, we want to recall how the "dilemma" of non-convex feasible spaces can be handled. For doing so, we refer to the specification of proportional fairness as a (set-theoretic, binary) relation.

2.2 Relations and Maximum Sets

Given a feasible space X, a (binary) relation R is a subset of $X \times X$. Thus, for two points x_1 and x_2 from X it is said that x_1 is in relation R to x_2 if the ordered pair (x_1, x_2) belongs to this subset. Relations may represent various things like equality, similarity, dependency etc. depending on the nature of X, and they can be represented in various equivalent forms (list of ordered pairs, tabular form, binary matrix, directed graph, logical property etc.). Here, we are focusing on the interpretation of a relation as a "betterness" relation, and will write $>_R$ to represent that x_1 is of better quality, or preferred to x_2 in some sense.

For each relation (no matter if it represents betterness or not) we may assign two characteristic subsets of X [15].

Definition 1 Given a relation $>_R$ then the **best set** is defined as the set of all x' such that for any other $x \in X$ $x' >_R x$ holds.

Definition 2 Given a relation $>_R$ then the **maximum set** is defined as the set of all x' such that for no other $x \in X$ $x >_R x'$ holds.

The best set is sometimes called set of greatest elements, and the maximum set is also called set of maximal elements or non-dominated set. Both are extreme elements of the relation within X.

Now we can adopt a popular notion in social choice theory, as a sub-discipline of mathematical economics. A social choice is called *rationalizable* if there is a relation such that the choice always corresponds with best set, or maximum set of that relation [16]. With regard to the proportional fairness problem in discrete feasible domains, we can use a similar concept when understanding the proportional fairness condition as a relation between points x and y. Then, the definition of the proportional fair state corresponds to the statement that x is a greatest element of this relation. The problem in a discrete domain is that the best set of greatest elements can be empty. The solution is to focus on the maximum set instead. This is the same approach as studied in multi-objective optimization, where the Pareto dominance takes the role of the binary relation, and the goal is specified as finding its Pareto front of non-dominated elements, i.e. its maximum set.

The approach allows for a general expansion of proportional fairness to any feasible domain, and coincides with proportional fairness in case of convex domains. But there is one relevant aspect (some might consider it a drawback) that has to be taken into account: there can be more than one maximal element. The example in Fig. 1(b) demonstrates that $(4, 3)$ as well as the point A both are maximal. The other aspect is the question how to find such maximum sets efficiently.

Only few works so far consider a strict relational framework to study fairness. With regard to majorities as fairness relation, several approaches have been presented by Ogryczak [17][18][5] as well as Kleinberg [19]. However, consideration of proportional fairness within such a framework has not been undertaken so far.

2.3 Meta-Heuristic Approaches

The other question is about appropriate search algorithms (as long as exact methods are not known, or are not tractable). We want to remind here that for more than a decade, there is already experience to use evolutionary computation for a specific case of such "relational optimization" - it is the use of the Pareto dominance relation in Evolutionary Multi-Objective Optimization (EMO). A large number of fundamental algorithms have already been provided to approximate maximal elements of this relation, each of them embedding the Pareto dominance relation into a specific way. Here we want to do the same but using other "betterness" relations, like proportional fairness. More specifically, we are going to compare corresponding versions of the Strength Pareto Evolutionary Algorithm ("version 2", SPEA2) [20] and the multi-objective version of Particle Swarm Optimization (MOPSO) [21]. It will come out that the embedding of ordered proportional fairness along with its bounded trade-off also allows to better handle a common problem of multi-objective optimization in the case of a larger number of objectives: the rapid decay of the probability of occurrence of the Pareto dominance relation, which often limits the application of EMO algorithms to problems with a smaller number of objectives (usually less than 10). In [22] it was already pointed out that therefore the search should focus on finding at least a few elements of the Pareto front in case of many-objective optimization. The maximum set of ordered proportional fairness will be shown to be a subset of the Pareto front, while the chance of random occurrence is much larger. Thus, ordered proportional fairness also appears as a means to assist the "decision maker" in general multi-objective optimization.

3 Ordered Proportional Fairness

Here, we recall definitions and concepts from [23] for the introduction of ordered proportional fairness and its main properties.

3.1 Definitions

We represent proportional fairness as a *proportional fair dominance relation* defined as follows:

Definition 3 A point $x \in R^n_+$ is proportional fair dominating another point $y \in R^n_+$, written as $x >_{pf} y$, if and only if

$$\sum_{i=1}^{n} \frac{y_i}{x_i} \leq n \tag{2}$$

Then, the best state of this relation, i.e. the state x such that $x >_{pf} y$ for any other y of the feasible space, is generally considered as the proportional fairness state.

We will study a modification of the proportional fairness relation, where the elements of x and y are *sorted* before the condition given by Eq. (2) (also called *indicator expression* in the following) is tested. In the following, a subscript $a_{(i)}$ indicates the i-th smallest element of a set A of real numbers a_i.

Definition 4 A point $x \in R^n_+$ is ordered proportional fair dominating another point $y \in R^n_+$, written as $x >_{opf} y$, if and only if

$$\sum_{i=1}^{n} \frac{y_{(i)}}{x_{(i)}} \leq n \tag{3}$$

There are several reasons to consider this alternative relation. As we will see later on, maximum sets of ordered proportional fairness are much smaller than maximum sets of proportional fairness, while at the same time they only contain maximal elements of the proportional fairness relation. This simplifies the selection, but also promotes convergence of meta-heuristic algorithms. A second reason is that this relation is symmetric with regard to permutations of elements of the comparison vectors. Thus, it better reflects fairness among the users (to which the components of a vector refer). Another fact that will be demonstrated later on is that for states where one or more components become extreme (sometimes called the "corners" of the objective space) this relation behaves more like maxmin fairness, while in a part of the feasible space where all vector components are very similar (and thus the sorting has not much influence), it strongly corresponds with proportional fairness. Thus, as a third reason, ordered proportional fairness can be expected to mediate between proportional and maxmin fairness.

3.2 Basic properties

At first, we will clarify the relations between proportional fairness and ordered proportional fairness. This can be summarized in the following theorem (see [23] for a proof).

Theorem 1 *For any x and y from R^n_+, $x >_{pf} y$ implies $x >_{opf} y$.*

The theorem can also be summarized by the inequality:

$$\sum_{i=1}^{n} \frac{y_{(i)}}{x_{(i)}} \leq \sum_{i=1}^{n} \frac{y_i}{x_i} \tag{4}$$

These relations can be used to better understand the part of space dominated by a point. We recall that the part of R^n_+ that is proportional fair dominated by a point x^* can be seen as the part of R^n_+ below the tangential hyperplane on the hypersurface given by $\prod x_i = \prod x^*_i$. Figure 2 shows the areas dominated by the point $(4,1)$ in the case $n=2$.

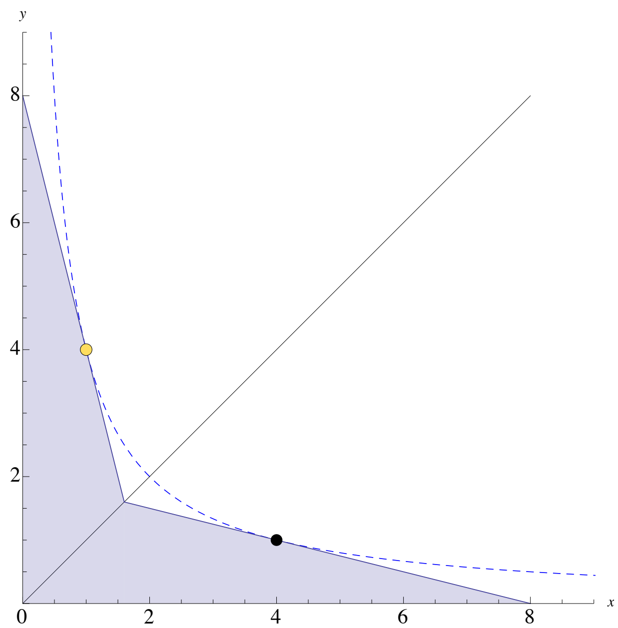

Fig. 2 Set of points *ordered proportional fair dominated* by the point $(4,1)$. It can be seen as the set of points proportional fair dominated by the point $(4,1)$ (black dot in the figure) **or** proportional fair dominated by its mirror point $(1,4)$ (yellow dot in the figure). Note that the blue dashed line shows the curve $x \cdot y = 1 \cdot 4$, and the dominated parts are bounded by tangents to this curve [23].

In this example, it can be seen that a point y is ordered proportional fair dominated by $(4,1)$ iff $(4,1) >_{pf} y$ or $(1,4) >_{pf} y$. From Theorem 1 it can be understood that this is a general characterization of the dominated space.

For given points x and y from R^n_+, we fix y and consider the set X_p of all points that are generated by permutations of the elements of x. Then, if for at least one $x* \in X_p$ the relation $x* >_{pf} y$ holds, by Theorem 1 $x* >_{opf} y$ follows, and then $x >_{opf} y$, since x has the same components like $x*$, only sorted in non-increasing order. Therefore, the value of the indicator expression for ordered proportional fairness does not change. On the other hand, if $x >_{opf} y$ while the components of x are not sorted, must correspond to $x* >_{pf} y$ for the $x* \in X_p$ where for all i $x*_i = y_{(i)}$. In summary, the fact that at least one element of X_p proportional fair dominates y is a necessary and sufficient condition for $x >_{opf} y$.

We remark that both relations are not complete, i.e. there are pairs (x, y) where neither $x >_r y$ nor $y >_r x$ holds. However, as will be demonstrated below, ordered proportional fairness can be seen as a "nearly complete" relation.

Like proportional fairness, also ordered proportional fairness is not transitive. It suffices to provide counterexamples: for ordered proportional fairness, consider the three points $x = (47, 43), y = (36, 53)$ and $z = (5, 91)$. Then we have for the values of the indicator functions $I_{xy} = 53/47 + 36/43 = 1.96 \leq 2, I_{yz} = 1.86 \leq 2$ but $I_{xz} = 2.05 > 2$. So this is a case where $x >_{opf} y$ and $y >_{opf} z$ but not $x >_{opf} z$.

However, both relations are cycle-free, which means that there is no sequence of m points x_i such that $x_i >_{pf|opf} x_{i+1}$ for $i = 1, \ldots, (m-1)$ and $x_m >_{pf|opf} x_1$. For proportional fairness, this can be verified by the fact that $x >_{pf} y$ implies that x has also a larger product of components[1], $\prod_i x_i > \prod_i y_i$ and thus a point with a smaller product of components than $\prod_i x_i$ can never dominate x. This implies that for $x >_{pf} y$ and $y >_{pf} z$ the products of components of z is smaller than $\prod_i x_i$, and z cannot dominate x. The same fact can be seen for sequences of any length in a similar manner.

For the case of ordered proportional fairness, it needs to consider the specific sorting of the components in the indicator expressions. Since in all cases the components are ordered in a non-increasing manner, the singular terms in the indicator expressions are all ordered in the same way. It means if we compare x and y we have an order for the terms $y_{(i)}/x_{(i)}$ which is the same order as for the terms $z_{(i)}/y_{(i)}$ in the comparison of y and z and $z_{(i)}/x_{(i)}$ in the comparison of x to z. Thus, the indicator expressions are formally identical to the expression for proportional fairness, just with all components of all points sorted by size. Then also here, cycle-freeness of proportional fairness implies cycle-freeness of the ordered proportional fairness.

[1] Note that two different points with the same product of components can never dominate each other.

3.3 Links to other relations

As a preparation, we provide a number of related definitions.

Definition 5 For any x and y from R^n it is said that x (strictly) Pareto dominates y, written as $x >_p y$, if and only if

$$\forall_i x_i \geq y_i \wedge \exists_j x_j > y_j \tag{5}$$

(note that there can be a corresponding definition using $<$ and $\leq$). A stronger version of this definition:

Definition 6 For any x and y from R^n it is said that x totally Pareto dominates y, written as $x >_{tp} y$, if and only if

$$\min[x_i] > \max[y_i] \tag{6}$$

We also consider the definition for lexicographic minimum relation.

Definition 7 The leximin relation is defined as a relation between two different vectors from R^n. Given two vectors x and y from R^n, first the coordinates of both vectors are sorted in nondecreasing order, giving vectors x^* and y^*. It is said that $x >_{leximin} y$ if and only if the first coordinate of x^* that is different from y^* is larger than the corresponding coordinate of y^*.

Now we can make some additional statements:

1. If the components of x vary strongly, the ordered proportional fair dominance relation resembles the leximin relation.
2. If the components of x are becoming more similar, the ordered proportional fair dominance relation resemble the proportional fair dominance relation.
3. Pareto dominance implies ordered proportional fair dominance, since Pareto dominance implies proportional fair dominance.

The given arguments also show that ordered proportional fairness can be seen as extension of proportional fairness "up to a permutation". This is formally the same "transformation" of a relation that leads from maxmin fairness to the leximin relation. Thus, we have also achieved a formal way to expand any relation r as a subset of the direct product of two sets A^n and B by a procedure of un-sorting: $x >_{u(r)} y$ holds iff there is at least one permutation of the elements of x such that $x >_r y$ holds. We can do similarly for the processing of over-sorting and requiring $x >_r y$ for *all* permutations. Applied to other relations, over-sorting applied to Pareto-dominance, as well as maxmin fairness, gives total Pareto dominance, and un-sorting gives a relation that has up to our knowledge not been studied so far, based on the comparison $x_{(i)} \geq y_{(i)}$. In summary, we have to acknowledge the fact that fairness relations like proportional fairness are probably not sparse among all relations of practical importance, but establish a rich category with numerous mutual relationships. By following a strict relational framework, we can expand concepts like proportional fairness to various new domains, as long as we can define "set," "subset" and "direct product." This can also be seen as an expansion to the approach presented in [11] as well as [24], where a general utility function was used in order to specify further fairness states (with proportional fairness being the special case for the logarithmic utility function). In all these cases, "unsorting" is applicable to further define new relations.

4 Monte Carlo Simulations

In addition to basic mathematical properties, the important additional practical application aspect is the tractability of the accompanying search problem: given a feasible set of vectors from a specific problem domain, how can we find the maximum sets (also called set of maximal, efficient, or non-dominated elements) for a given relation? In a recent work [25], we have already investigated the options to use meta-heuristic search algorithms, and demonstrated their efficiency. But the study also illustrated the important aspects of the relations itself, in order to pose a tractable search problem, and with regard to the chance of random occurrence of a relation between two random points. It can be easily seen that a sparse relation inhibits the initial explorative stage of such algorithms. Therefore, we want to focus on this aspect when evaluating the presented relations, and provide a number of related Monte Carlo simulations to estimate corresponding probabilities and probability distributions.

4.1 Probability of occurrence

In this part, we want to study the probability of occurrence of relations between random vectors with increasing dimension n. It is known to fall exponentially for the Pareto dominance relation, and can be expected to be 0.5 for complete relations like leximin. Other estimates might be harder to find analytically, so we sampled 100,000 pairs x and y of random points from $(0,1)^n$ and counted the number of occurrences of the relation $x >_R y$ for various relations. Table 1, partially taken from [23] gives an overview of the results for dimensions up to 100.

Table 1 Frequency of occurrence of various relations with increasing dimension n among 100,000 random samples from $(0,1)^n$.

n	$>_p$	$>_{pf}$	$>_{opf}$
2	25032	40344	46784
3	12523	33203	45101
5	3067	23305	43117
10	117	11219	41512
20	0	3103	40853
30	0	994	40628
50	0	217	40744
100	0	0	41720

The exponential decay of the Pareto dominance relation $>_p$ can be confirmed, for dimensions 20 onwards it is virtually not present anymore (a large hindrance for meta-heuristic approaches to multi-objective optimization with larger number of objectives). Proportional fairness $>_{pf}$ also decreases, but at least for less than 100 dimensions not exponentially (the exact decay rule is currently unknown). Then, ordered proportional fairness stays nearly constant at a value around 40%. This means that ordered proportional fairness is close to a complete relation having 50%, and it is even possible that the probability is converging (at least, it seems to decay very slowly). This can be related to the fact that this relation has a much larger space dominated by a point x than for example proportional fairness. In fact, it is fusing an exponentially increasing number of proportionally fair dominated spaces (one for each permutation of the components of x), and this can "compensate" the (conjectured) exponential decrease of the volume of spaces proportionally fair dominated by a point with increasing dimension.

4.2 Conditional Relations

We also studied conditional probabilities, to indicate the relation between ordered proportional fairness, the maximization of product of components, and the maximization of the smallest component. The reason is that in case of convex feasible spaces, the proportional fairness has a best element, which is also maximizing the product of its components. This cannot be expected from the ordered proportional fairness as well, as the subspace dominated by a point is not always convex, but a relation to component product maximization might still exist.

This is confirmed by the result shown in Fig. 3. There, a number of conditional probabilities $p(>_1 \mid >_2)$ have been sampled from random vectors from $(0,1)^n$

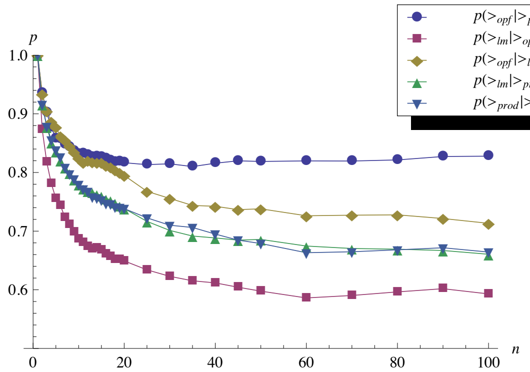

Fig. 3 Sampled conditional probabilities that a random vector x from $(0,1)^n$ $>_1$-dominates another random vector y given that $x >_2 y$, where the relations $>_1$ and $>_2$ are among the relation of having larger product of components($>_{prod}$), leximin, and ordered proportional fairness (100,000 samples for each dimension n). The case $p(>_{prod} \mid >_{opf})$ is not shown, since this is always 1.

(100,000 samples in each case). The used relations were ordered proportional fairness $>_{opf}$, leximin $>_{lm}$ and a further relation $>_{prod}$ which holds between x and y iff $\prod_i x_i > \prod_i y_i$. The case $p(>_{prod} \mid >_{opf})$ is not shown, since ordered proportional fairness always implies a larger product of components (this is the case for proportional fairness, and ordered proportional fairness always corresponds to a particular case of proportional fairness). So, this value is always 1.

A few things can be seen from the result of the simulation: the first is that the probabilities seem to converge, or at least decay very slowly with increasing dimension. Which of these is truly the case might be subject of further investigation. We also can confirm that ordered proportional fairness strongly relates to product maximization. But we can also see the difference between $p(>_{opf} \mid >_{lm})$ and $p(>_{lm} \mid >_{opf})$, with the former being notably larger than the latter. Thus, we can also see that the chance of having ordered proportional fairness among states with larger minimum is rather high, especially higher than the same for states with a larger product of components. This was one goal mentioned in the introduction: for states with small components, select stronger in the sense of maximizing the minimum state, otherwise more in the sense of proportional fairness.

5 Comparison of Meta-Heuristic Algorithms

5.1 Wireless Channel Allocation

We want to consider the problem of finding the maximum set of ordered proportional fairness in a discrete domain, for a task of distribution of indivisible goods. A typical problem here is the allocation of channels of a base station to mobile users in wireless communication. The problem has been studied in [26]. For each time slot, one channel can be given to at most one user. The abstraction of the problem is as follows.

In Wireless Channel Allocation, a blank matrix B of channel-timeslot pairs with a total of M cells b_i, a set $U = (u_i)$ of N users and an $M \times N$ matrix C of *channel coefficients* of real values from $[0,1]$ are given. The task is to enter at most one user into each blank cell in B, i.e. to provide an *allocation* $a : B \to U$ of cells to users with $|\{u \in U \mid a(b) = u\}| \leq 1$ for all cells $b \in B$. Each entry c_{ij} of the matrix C represents the *utility* for user u_i in case of assignment of cell b_j, as a model abstraction of all the physical and logistic circumstances of the wireless access. For a given allocation a, the *performance* for each user is given by

$$p(u_i) = \sum_{j,a(b_j)=u_i} c_{ij} \tag{7}$$

i.e. the sum of channel coefficients for all channels allocated to the user. Channel allocation has to be performed such that, in some sense, all users are “satisfied” with their individual performances as much as possible. The actual problem is to specify the meaning of “satisfied” in an efficient way. For example, considering maximization of the sum of all performances is not a good way to satisfy all users: the optimization problem could be easily solved by selecting for each cell one of the users with maximum channel coefficient. But this way it can happen that then some users will never get any channel allocated, and these users will have no wireless access. Therefore, the economics of WCA becomes relevant, especially aspects of fairness.

5.2 Relational Meta-Heuristics

In a recent study [25] several meta-heuristic algorithms derived from Evolutionary Multi-Objective Optimization (EMO) algorithms have already been investigated for the task of finding the maxmin fair state of traffic rate allocation in a wired network with link-capacity constraints. For this problem, an exact algorithm is known (the so-called Bottleneck Flow Control [27]) and this gave the opportunity for a direct performance comparison of these algorithms. Two algorithms, one derived from SPEA2 [20] and one derived from MOPSO [21] demonstrated the best performance, so we will consider these algorithms here as well. On the other hand, in [28] a comparative study with regard to different relations but same algorithm (relational SPEA2) was provided, already demonstrating the suitability of ordered proportional fairness (or "good performance") compared to other fairness relations, also for the WCA problem. We will adopt the evaluation method to compare with "same effort" (SE) random search and mutual dominance. While the rapid expansion of search space sizes forbids an exhaustive search for problem dimensions beyond 10 (cells or users) we can get at least a measure for the progress of the algorithm towards the maximum set, and also have a base for direct comparison of algorithms with each other.

At first, we use a modification of the Strength Pareto Evolutionary Algorithm SPEA2 [20], where internal processing using Pareto dominance relation is replaced by a general relation.

Relational SPEA2

---

1. Given is a relation R by a set of pairs (a, b), where $(a, b) \in R$ means that a is in relation to b.
2. Initialize the population with random allocation vectors a_i.
3. Repeat the following steps for a fixed number of generations:
4. Compute the performance vectors p_i for all individuals (allocations) a_i.
5. For each individual i, compute the R-value R_i, i.e. the number of other individuals j, for which p_i is in relation to p_j (i.e. $(p_i, p_j) \in R$).
6. For each individual i, compute the S-value S_i, i.e. the sum of the R-values of all individuals j such that p_j is in relation to p_i (i.e. $(p_j .p_i) \in R$).
7. Tournament selection: randomly select two individuals and keep the one with the smaller S-value. If the S-values are equal, take any one of the two. Repeat by selecting again a pair and keeping the one with the smaller S-value. This gives a "mating pair" of individuals.
8. Cross-over: generate a new allocation by randomly selecting assignments from either the first or from the second individual's allocation vector.
9. Mutation: with some probability μ_m, modify each component according to a given distribution.
10. The foregoing three steps establish the children for the next generation. Put children and former population (parents) together into a new intermediate population. Compute all performance vectors and R and then S-values and select the individuals with the smallest S-values to establish the new generation.

The other algorithm is derived from a multi-objective version of Particle Swarm Optimization, where the selection of the global best particle is replaced from a random selection from an archive set, called "leader."

Generally, a PSO maintains a set of particles with position, velocity, and a "memory" of a local best position. In each step, a positional update is performed: the velocities are updated by terms directing along former velocity, partially into the direction of the local best, and partially into the direction of a global best (here a random selection from a set of "leading particles"). The algorithm needs to specify the weights for the update terms. The steps in detail:

Relational MOPSO

---

1. Given is a relation R by a set of pairs (a, b), where $(a, b) \in R$ or $a >_R b$ mans that a is in relation to b.
2. Initialize the swarm S with particles with random allocation vectors a_i (positions) and random velocities v_i from $[0, 1]$. Set the "local best" position l_i of each particle equal to its current position. Initialize the leader set L with all particles, whose performances belong to the maximum set of the relation.
3. Repeat the following steps for a fixed number of positional updates:
4. For each particle i, compute the inertia term $T_1 = w \cdot v_i$ where w is the inertia weight and the multiplication is component-wise; the local update term $T_2 = w_{local} \cdot rand \cdot [l_i - a_i]$; and the global update term: select a random element g_i from the leader set, then $T_3 = w_{global} \cdot rand \cdot [g_i - a_i]$. $rand$ is a random vector with uniform random components from $[0, 1]$.
5. Compute particle velocities as $v_i := T_1 + T_2 + T_3$.
6. Update particle position as $a_i := a_i + v_i$.
7. (Optionally) if $rand < \theta_{mutation}$ then mutate half of the particle positions by adding uniform random numbers from $[-0.5, 0.5]$.
8. Adjust each component a_{ij} of particle positions: $a_{ij} := \max[U \cdot Round(a_{ij} / \max_j[a_{ij}]), 1]$. U is an upper bound for each particle component (here the number of users) and the step ensures integer components of a corresponding to user numbers.
9. Compute the performance vectors p_i for all particles (i.e. allocations) a_i.
10. For each particle, if $p_i >_R p(l_i)$ replace l_i by a_i ($p(l_i)$ denotes the performance of local best position l_i).
11. Update the leader set L as the maximum set of $L \cup S$ with regard to relation R among particle performance vectors.

Two comments are in place: at first we have to note that neither proportional fairness nor ordered proportional fairness are transitive relations. They are "weakly transitive" in the sense of being cycle-free, and this is a natural consequence of the bounded trade-off of these relations. But this makes the popular use of an archive (in order to keep the best solutions found ever during the whole processing of the algorithm) in EMO algorithms complicated. Since the maximum set of the studied fairness relations will be part of the Pareto front, an archive could be used in the common way, but since we are studying problems with larger scales, the archive size would heavily increase due to the exponential decay of Pareto dominance relation occurrence, and we would also have to study countermeasures to keep their size within reasonable bounds. This would somehow interfere with the primary goal of comparing the selection techniques of the various algorithms, and therefore, we are not employing an archive here directly (in some sense we do while using a leader set in relational MOPSO).

Second comment is about the handling of infeasible solutions, i.e. solutions where there are users without a channel allocation. This is a problem since we would have user performance vectors with 0 entries, while the definition of (ordered) proportional fairness is restricted to vectors with positive components. In [28] the "penalty" here was to declare a vector with 0-valued components to not be in relation to any other vector. However, it appeared to lower performance in a few cases where the comparison was made between two vectors, both having 0 components, and thus becoming mutually non-dominating. Especially in case of proportional fairness and a small difference between number of users and number of cells (where such allocations are more likely), solutions with missing user allocations could happen to stay in the maximum set as well, but are easily dominated by random elements. Here, we have replaced 0-valued elements in the indicator expressions by small values (0.0001 specifically). Results confirm that by this small change, allocations where one or more users are excluded are directly removed from the maximum sets, and we achieve a smaller gain in performance.

For evaluating and comparing the performance of relational SPEA2 and MOPSO for the two fairness relations, a suitable measure is the direct comparison with random search. In this paper, we are focusing on three related evaluations: after operating the algorithm for a particular relation, we also sampled 10,000 random vectors. So far, the algorithms would have processed 10 solutions times 1000 generations, thus 10,000 vectors in the searchspace. Then, we compute two measures. Measure M_1 is the number of random vectors that dominate at least one maximal element of the evolved population (or progressed swarm) according to the current relation. Measure $M_2(k)$ has a parameter k. For $k = 1$ we compute the relative number of random vectors that are dominated by at least one vector of the evolved population (or swarm). For $k = 2$ we also compute the relative number of random vectors either dominated by a maximal element of the population, or by at least one random vector that is dominated by at least one element of the maximum set of the population (and which was counted in the measure for $k = 1$). The definition continues in a recursive manner: for larger k we would compute the relative number of random vectors dominated by at least one random vector that was counted in the computation of the measure for $k - 1$. However, we only use $k = 1$ and $k = 2$. Measure $M_2(2)$ is not expected to show a much different statistical distribution than measure $M_2(1)$ but can help to see the "spread" of the dominance relation among random vectors. These measures allow for a quantification of the success of the algorithm to approximate the maximum set of the particular relation.

In all experiments, the relational SPEA2 population had 10 individuals, and the algorithm run for 1000 generations. Before continuous cross-over, tournament selection was used according to the S-value of randomly chosen individuals. Polynomial mutation was used with mutation distribution index 3 and probability 0.3. This basically follows the settings used in [28]. The relational MOPSO also used 10 particles and 1000 update steps, but we have used two parameter sets, as we noted a rather strong sensitivity of this algorithm on parameter choices (somehow in contrary to often made claims in the literature about general simple parameterization of PSO). Parameter set 1 used the inertia weight $w = 0.99$, $w_{local} = 0.5$ and $w_{global} = 0.9$ (this setting seems to achieve best results), parameter set 2 used the same inertia weight, but $w_{local} = 0.8$ and $w_{global} = 0.3$ (this is the same setting as in [25] where relational MOPSO was used for a wired network problem). Mutation probability was 0.1 in both cases.

Table 2 presents a number of average results for different settings of the WCA problem, comparing relational SPEA2 and MOPSO for proportional fairness and ordered proportional fairness. All problem dimensions are between 5 and 15 users, which is a realistic scale for wireless access problems. But the analysis of larger problems is also limited by the growing incapability of random search to find dominating solutions at all. Also we want to remark that a corresponding multi-objective optimization problem with 10 or more objectives already tends to be intractable, and when doing the same but based on a different relation, up to 15 users - thus also 15 objectives - seems to be already a challenging problem. Each average value in the

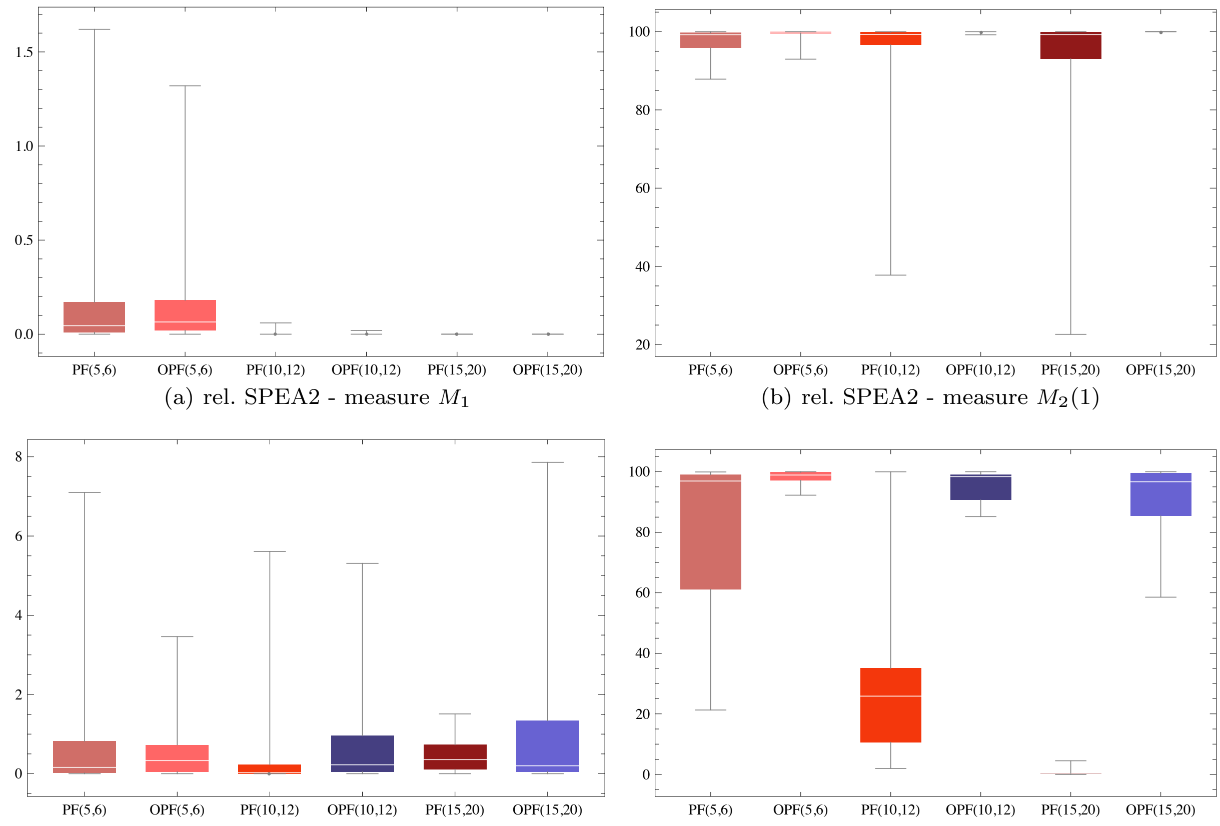

(a) rel. SPEA2 - measure M_1

(b) rel. SPEA2 - measure $M_2(1)$

(c) rel. MOPSO - measure M_1

(d) rel. MOPSO - measure $M_2(1)$

Fig. 4 Box plots of the M_1 and $M_2(1)$ measure for relational SPEA and MOPSO algorithm for proportional (PF) and ordered-proportional (OPF) fairness relation ($R(n, m)$ indicates relation R for n users and m cell WCA problem).

table was computed from 30 random instances of the WCA problem. Figures 4(a) to 4(d) also show the box plots of the statistical distribution of these results (with common elements minimum, 1/4 quantile, median, 3/4 quantile, maximum). Note that small values of M_1 and values of $M_2(1)$ and $M_2(2)$ close to 100% are favorable.

We can make a number of observations (we abbreviate proportional fairness by PF and ordered proportional fairness by OPF, and we will also consider the three measures to represent "performance of a relation"):

- In general, most of the M_1 measure values are below 1% and often are nearly 0 or equal to 0. The largest values appear for smaller problem scales. This can be understood by the fact that in this case (5 users, 6 cells) 10,000 samples come close to exhaustive search (the number of possible mappings is $5^6 = 15,625$, while for the additional requirement that each user appears at least once it would correspond to the number of surjective mappings of m cells to n users, known to be $n!S(m, n)$ with $S(m, n)$ the Stirling number of second kind - the value then is much smaller, 1800).
- For relational SPEA2, the measure values M_1 for PF and OPF are nearly equal, and the general performance for both relations is very good. Figure 4(a) shows a little bit more peaked distribution of the OPF measure values.
- With regard to measure $M_2(1)$, OPF values are notably better than PF values, and the difference grows with increasing problem scale. Figure 4(b) also shows that the fluctuations for OPF are much smaller and tend to stabilize for larger problem scales. The tendency is confirmed by measure $M_2(2)$.
- The results are generally very different for the two parameter settings of the relational MOPSO, indicating a strong sensitivity of relational MOPSO. Results for parameter set 2 are worse - we provide this result here only for illustration of parameter influence on performance, and will not discuss these results any further.
- For relational MOPSO (parameter setting 1), measure values M_1 do not much differ for PF and OPF. This can also be seen in Fig. 4(c), where also other statistical measures appear rather similar. Altogether, the values are all below 1%.

Table 2 Results for relational SPEA2 and MOPSO competing with random search. For each setting of the number of users and cells, and for each relation $>_r$, the three cells indicate (1): the share of random samples among 10,000 samples dominating at least one element of the maximum set of the population according to $>_r$ (measure M_1), (2): the share of random samples dominated by this maximum set (measure $M_2(1)$), in percent of 100, and (3): the share of random samples either dominated by the maximum set, or by a random sample that is dominated by the maximum set (measure $M_2(2)$), in percent of 100.

Relational SPEA2						
WCA	$>_{pf}$			$>_{opf}$		
	M_1	$M_2(1)$	$M_2(2)$	M_1	$M_2(1)$	$M_2(2)$
5 users 6 cells	0.1677	97.64	99.56	0.1573	99.5	99.73
10 users 12 cells	0.002333	91.71	93.44	0.002333	99.93	99.94
15 users 20 cells	0.0	90.63	92.09	0.0	100.0	100.0
Relational MOPSO Parameter Setting 1						
WCA	$>_{pf}$			$>_{opf}$		
	M_1	$M_2(1)$	$M_2(2)$	M_1	$M_2(1)$	$M_2(2)$
5 users 6 cells	0.7373	80.65	84.93	0.5887	98.07	98.78
10 users 12 cells	0.543	28.51	30.91	0.763	95.97	96.36
15 users 20 cells	0.454	0.5117	0.5283	0.948	92.61	92.97
Relational MOPSO Parameter Setting 2						
WCA	$>_{pf}$			$>_{opf}$		
	M_1	$M_2(1)$	$M_2(2)$	M_1	$M_2(1)$	$M_2(2)$
5 users 6 cells	3.44	36.73	43.78	2.732	90.52	92.05
10 users 12 cells	2.355	1.968	2.368	3.979	91.03	91.88
15 users 20 cells	0.8173	0.08133	0.08233	10.437	77.765	78.252

- The situation is different for measure $M_2(1)$: OPF measures are all above 90% while PF values are low, and fall rapidly with increasing problem dimension. Figure 4(d) is another confirmation for this fact, and also measure $M_2(1)$ confirms this trend.
- Nevertheless, with regard to measure $M_2(1)$, the use of relational MOPSO and OPF, there is a scaling effect as the distribution gradually broadens with increasing problem scale (see 2nd, 4th and 6th boxplot in Fig. 4(d)).

The general evaluation gives a clear advantage for the use of ordered proportional fairness (which also allows to find maximal proportional fair solutions), and in combination with the relational SPEA2 algorithm.

We want to add the comment that the choice for the number 10,000 of random samples is basically motivated by the fair comparison of search efforts. Using a larger number of random vectors might be considered to get more stable sample values, but figs. 4(a) to 4(d) already indicate, with the exception of MOPSO performance for measure $M_2(1)$, that the spread of the distribution is small enough to allow for significant conclusions.

We will end this section with a small discussion about the different performances. One question might be about the, on first glance "paradoxical" performance improvement for larger problem scales. In fact, for the problems with 15 users and 20 cells and OPF, there was not a single case that a random allocation dominated the algorithm result, and also allocations from the algorithm result were all dominating each of the 10,000 random points. This is the meaning of $M_1 = 0$ and $M_2(1) = 100$. However, from this we cannot conclude that the algorithm has already arrived at the maximum set, we only see maturing against random search and that it is getting rapidly unlikely to achieve valuable solutions for such problems by pure random search. Beyond this problem scales, our random measure might become infeasible to expose performance differences between algorithms, at least for the WCA problem.

The second question is about the lower performance of MOPSO. The main reason can be seen in the discrete encoding of particle positions. Generally, PSOs (in "single-objective" version) are known to operate better in continuous problem domains. The crucial point is the update of the leaders and the pre-order given between allocations by the relation between the corresponding performances. If we would allow for real-valued components in the allocation, and only round before evaluating a particle, the number of redundant maximum set elements would rapidly increase. We are not aware of any better way to keep the leader set growth bounded than to use integer representation of particles and so avoid duplicates in the leader set (with regard to performance). However, integer representation is not well fitting to the real-valued velocity update terms. In sight of these problems, the performance of relational MOPSO might still be seen surprisingly good.

As a third point, the obvious better performance of ordered proportional fairness has to be relativized by the larger computational effort to compare two vectors. In addition to proportional fairness, this comparison also includes the sorting of the components of vectors before comparison. On the other hand, comparing the computational effort for proportional fairness with e.g. Pareto dominance depends much on the differences in effort to perform a comparison and to perform a division, and no general claims can be made here.

6 Conclusion

In this paper, we have presented an approach to proportional fairness fulfilling two goals: (1) the extension of proportional fairness to discrete problems and to formulate corresponding optimization problems as the problem of finding maximal elements with regard to a proportional fairness relation, and (2) the design of meta-heuristic search algorithms to approximate these sets of maximal elements. The experimental results demonstrate the feasibility of both approaches. Furthermore, the performance of the approximation can be improved by using a sorted version of proportional fairness, introduced here as ordered proportional fairness. Monte Carlo simulation have also demonstrated that ordered proportional fairness can mediate between proportional fairness (in case of less varying vector components) and maxmin fairness (in case of strongly varying vector components).

The application domain here was restricted to wireless access networks. However, there are at least two aspects of the presented approach that are more general: (1) the representation of fairness as a set-theoretic binary relation, and (2) for relations with a subset of R^n as domain (i.e. vector relations) the procedure to define a new relation from a given one by relational sorting and consider the use of the ordered relations to find maximal elements of the original relation.

As future work, we consider the extension of this relational sorting to other relations, esp. the Pareto dominance and the option to improve finding of non-dominated solutions for problems with a larger number of objectives, as well as α-fairness. In addition, there are a number of other allocation problems in wireless access that need further refinements of this concept: mapping of demand vectors to feasible channel allocations, choice of relays, concurrent virtualizations of networks.

Acknowledgment

This work was supported by JSPS KAKENHI Grant Number 24650030.

References

1. D. Bertsekas and R. Gallager, *Data Networks.* Englewood Cliffs, N: Prentice Hall, 1992.
2. F. Kelly, "Charging and rate control for elastic traffic," *Eur. Trans. Telecomm.*, vol. 8, pp. 33–37, Jan./Feb. 1997.
3. X. Zou, Y. Chen, M. Liu, and L. Kang, "A new evolutionary algorithm for solving many-objective optimization problems," *Systems, Man, and Cybernetics, Part B: Cybernetics, IEEE Transactions on*, vol. 38, no. 5, pp. 1402–1412, 2008.
4. L. Dubins and E. Spanier, "How to cut a cake fairly," *The American Mathematical Monthly*, vol. 68, no. 1, pp. 1–17, 1961.
5. W. Ogryczak, "Inequality measures and equitable locations," *Annals of Operations Research*, vol. 167, no. 1, pp. 61–86, 2009.
6. S. Haldar and D. K. Subramanian, "Fairness in processor scheduling in time sharing systems," *SIGOPS Oper. Syst. Rev.*, vol. 25, no. 1, pp. 4–18, 1991.
7. R. W. Hall and D. Tymoczko, "Submajorization and the geometry of unordered collections," *The American Mathematical Monthly*, vol. 119, no. 4, pp. 263–283, 2012.
8. M. Baglietto, R. Bolla, F. Davoli, M. Marchese, and M. Mongelli, "Integration of pricing models between best-effort and guaranteed performance services in telecommunication networks," *Control Engineering Practice*, vol. 11, no. 10, pp. 1209 – 1226, 2003.
9. T. Bonald, L. Massoulié, A. Proutière, and J. Virtamo, "A queueing analysis of max-min fairness, proportional fairness and balanced fairness," *Queueing Systems*, vol. 53, pp. 65–84, 2006.
10. B. Tan, L. Ying, and R. Srikant, "Short-term fairness and long-term QoS in the internet," *Performance Evaluation*, vol. 67, no. 5, pp. 406 – 414, 2010.
11. T. Harks and T. Poschwatta, "Congestion control in utility fair networks," *Computer Networks*, vol. 52, no. 15, pp. 2947 – 2960, 2008.
12. H. Zhou, P. Fan, and J. Li, "Cooperative proportional fairness scheduling for wireless transmissions," in *Proceedings of the 2009 International Conference on Wireless Communications and Mobile Computing: Connecting the World Wirelessly*, ser. IWCMC '09. ACM, 2009, pp. 12–16.
13. H. Boche and M. Schubert, "Nash bargaining and proportional fairness for wireless systems," *IEEE/ACM Trans. Netw.*, vol. 17, no. 5, pp. 1453–1466, Oct. 2009.
14. M. Uchida and J. Kurose, "An information-theoretic characterization of weighted α-proportional fairness in network resource allocation," *Information Sciences*, vol. 181, no. 18, pp. 4009 – 4023, 2011.
15. A. K. Sen, *Collective Choice and Social Welfare.* Holden-Day, San Francisco, 1970.
16. K. Suzumura, *Rational Choice, Collective Decisions, and Social Welfare.* Cambridge University Press, 2009.
17. W. Ogryczak, T. Śliwiński, and A. Wierzbicki, "Fair resource allocation schemes and network dimensioning problems," *Journal of Telecommunications and Information Technology*, vol. 3, pp. 34–42, 2003.

18. W. Ogryczak, A. Wierzbicki, and M. Milewski, "A multicriteria approach to fair and efficient bandwidth allocation," *Omega*, vol. 36, pp. 451–463, 2008.
19. J. Kleinberg, Y. Rabani, and E. Tardos, "Fairness in routing and load balancing," *Journal of Computer and System Sciences*, vol. 63, no. 1, pp. 2–21, 2001.
20. E. Zitzler, M. Laumanns, and L. Thiele, "SPEA2: Improving the Strength Pareto Evolutionary Algorithm," in *EUROGEN 2001, Evolutionary Methods for Design, Optimization and Control with Applications to Industrial Problems*, K. Giannakoglou, D. Tsahalis, J. Periaux, P. Papailou, and T. Fogarty, Eds., Athens, Greece, 2002, pp. 95–100.
21. M. Reyes-Sierra and C. A. C. Coello, "Multi-objective particle swarm optimizers: A survey of the state-of-the-art," *International Journal of Computational Intelligence Research*, vol. 2, no. 3, pp. 287–308, 2006.
22. J. Knowles and D. Corne, "Quantifying the effects of objective space dimension in evolutionary multiobjective optimization," in *Evolutionary Multi-Criterion Optimization*, ser. LNCS 4403. Matsushima, Japan: Springer, 2007, pp. 757–771.
23. M. Köppen, K. Yoshida, and M. Tsuru, "Unsorting the proportional fairness relation," in *Proc. Third International Conference on Intelligent Networking and Collaborative Systems (INCoS 2011)*, Fukuoka, Japan, November 2011, pp. 47–52.
24. M. Köppen, K. Yoshida, M. Tsuru, and Y. Oie, "Fairness-based global optimization of user-centric networks," in *INC, IMS and IDC, 2009. NCM'09. Fifth International Joint Conference on.* IEEE, 2009, pp. 441–443.
25. 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*, 2010, pp. 2647 –2655.
26. 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.
27. J. Jaffe, "Bottleneck flow control," *IEEE Trans. Commun.*, vol. COM-29, July 1981.
28. M. Köppen, K. Yoshida, and K. Ohnishi, "Meta-heuristic optimization reloaded," in *Proc. Third World Congress on Nature and Biologically Inspired Computing (NaBIC2011)*. Salamanca, Spain: IEEE CS Press, October 2011, pp. 562 – 568.