Multi-Jain Fairness Index of Per-Entity Allocation Features for Fair and Efficient Allocation of Network Resources

Mario Köppen, Kei Ohnishi, and Masato Tsuru
Department of Computer Science and Electronics
Kyushu Institute of Technology
680-4 Kawazu, Iizuka
mkoeppen,ohnishi,tsuru@cse.kyutech.ac.jp

***Abstract*—Due to its simplicity and its easy comprehension, Jain's fairness index is still among the most popular measures to compare justness of allocations. However, it was already argued in the original paper that while the way of computing the index is well established, it is not immediately clear to which metric to apply the computation. Thereby, metric stands for a specific choice of a system observable. Here we study the extension of Jain's index to multiple metrics at once. We propose a set of per-entity allocation features to represent justness of an allocation, and to derive corresponding vectors of feature-wise taken Jain's fairness indices. The features give a numerical representation of fulfilling common fairness properties like proportionality, envy-freeness and equity of an allocation. Then, maximizing the smallest index gives an efficient procedure for allocation of goods. We study this procedure for the problem of allocating wireless channels in a multi-user setup and compare the influence of the various feature choices on the efficiency of the solution.**

***Keywords*-fairness, Jain's fairness index, wireless channel allocation, leximin relation**

I. Introduction

These days the networking paradigm became of crucial importance for the modeling and performing of many socio-economic interactions. Along with the spread of this insight among researchers and engineers, the study of related limited resource distribution and allocation problems gained momentum as well. Such problems around the central theme of fairness, justness and efficiency have been explicitly formalized and studied in the economics field for more than 50 years (see e.g. [1]), while at the same time many historical references to related considerations, e.g. in the bible or talmud could be identified (see the excellent monograph of Peyton H. Young for cases [2] or [3] for a more recent textbook). However, it appears to be a problem field with many aspects, many solutions, and also many fundamental drawbacks and impossibilities like the Arrows' Theorem.

There is a general incentive to define measures of corresponding observables of a system and thus being able to express performances, efficiency etc. of a system under observation in as few as possible numerical quantities. These numerical values can be stored as off-line reference to always access the system performance at a definite instance of time and location (cardinality), and, what is more relevant, to allow comparison of current system performance with past states, or judge on the faith of the system in estimates of future states. It is no wonder that this tendency also found its way into the "measurement" of intangible goods like fairness. The temptation to compare the development of economical vectors, monitor the effects of campaigns, follow worrying trends etc. is large, and for example the Gini index already provided a convenient means for judging on the fairness of distribution of wealth in a society.

However, the recent prevalent spread of the internet as a human-crafted and controlled network and its multitude of technical fair division and allocation problems (be it traffic rate assessment, channel allocation, queue management, relay assignment, or carrier selection and many more) was accompanied by the introduction of a new fairness measure. In fact, and oddly enough, in the networking domain the Gini index never played an important role, while the proposal of a rather simple formula in the 80s still astonishes practitioners with its easy comprehension, computation and its endlessly ability to produce charts to demonstrate the fairness aspects of new engineering approaches. We speak about so-called Jain's fairness index as it was introduced in [4].

While hundreds of studies have used this index in simulations or theoretical estimates, the predominant use was always with the performance values directly exposed by the system. Even in the original paper [4] it was mentioned that this is not necessarily the best choice. Fairness can relate to various aspect of an allocation and sometimes the issue can rather become to identify the fitting "allocation metric" than to equalize user performance. In this paper, we want to follow this suggestion by studying multiple fairness indexes at once - one index for one of several specific fairness aspects of the allocation. Then, relational optimization is applied to select an allocation among the vectors of fairness indices. Analysis of the selected allocation gives insight into the mutual influence of different fairness aspects and allows conclusions on the focus of future system design.

In Section II we will recall Jain's fairness index and its essential properties. The following Section III than introduces a multi-Jain fairness index and exemplifies its definition for the case of wireless channel allocation. Then Section IV presents a few experiments using this multi-fairness index and demonstrates the manner of analyzing these results.

II. Jain's Fairness Index

Given a set of n data $x = (x_1, x_2, \dots, x_n)$ where $x_i \in R$ are real numbers and $x_i \geq 0$, then Jain's fairness index is given by the formula [4]

$$J(x) = \frac{\left[\sum_{i=1}^{n} x_i\right]^2}{n \cdot \sum_{i=1}^{n} x_i^2}. \quad (1)$$

Basically it is the square of the ratio of two power means of the x_i, the arithmetic mean and the quadratic mean. By the power means inequality we know that this is a value from $[0, 1]$. Moreover, its value is only 1 if all x_i are equal. In the original paper, this choice was motivated by some wanted features of an index that can represent the fairness of a distribution. These features were in particular:

- Its computation can be performed for any number of arguments.
- It is scale-independent. If we multiply all x_i by a constant factor α the value of J will not change.
- The value of J should be bounded between 0 and 1.
- Continuity, i.e. a small change of any of the x_i should cause a small change of the index J.

It is easy to see that the proposed index fulfills all these requirements. In addition, it has a convenient interpretation: assuming that a distribution is such that M items have to be given to n users. In a distribution commonly seen as "unfair" only k of the n user receive each M/k and all other receive 0. Then the Jain's fairness index computes to

$$J(x) = \frac{[k \cdot M/k]^2}{n \cdot k \cdot (M/k)^2} = \frac{k}{n} \quad (2)$$

i.e. it is equal to the share of users receiving goods. Then, a fairness index of 0.8, for example, can be seen as a situation where only 80% of the users receive goods and all other nothing.

Authors are not aware of any attempt to characterize the class of functions with all these properties, but it can be speculated that the Jain expression will be among the simplest. This can explain the popularity of this fairness index. However, there have been attempts to generalize that index. One example is a recent proposal of so-called "fairness axioms" [5]. Based on a number of functional properties, for some reason called "axioms", a class of functions is identified that fulfills all these properties. It can be written as (taken from [6] where a more compact representation is given):

$$\begin{aligned} F_{\beta,\lambda}(x) &= sgn(1-\beta) \times \\ &\times \left(\sum_{i=1}^{n} \left(\frac{\mu_i x_i}{\sum_{k=1}^{n} \mu_k x_k} \right)^{1-\beta} \right)^{\frac{1}{\beta}} \left(\sum_{i=1}^{n} \mu_i x_i \right)^{\lambda} \end{aligned}$$

which depends on two parameters β and λ. The use of weights w_i is optional. The second factor is remarkable: it was introduced in [5] in order to handle a trade-off with efficiency. In fact, Jain's fairness index gives a strong focus on equality of all x_i but not on their magnitude. To explain this, consider sharing M (dividable) goods among n users such that each user receives M/n goods. For this allocation is $J = 1$. Now compare with a "wasteful" allocation where each user receives $1/100 \cdot M/n$ goods. But also here we have $J = 1$. There might be a hidden agreement that an allocation should be such that all goods are shared, but in many allocation problems (esp. of indivisible goods or for multi-resource tasks) it is not possible to guarantee the distribution of all goods.

Therefore the addition of the λ-dependent factor in Eq. (3) was made. But this breaks one of the basic requirements on the Jain index, to guarantee that the index is bounded by 0 and 1. Also, the expression given by Eq. (3) is *not* a direct generalization of Jain's fairness index, since for the case $\beta = -1$ it computes to $n \cdot J(x)$ and not $J(x)$.

Thus the generalization of Jain's fairness index still remains an issue. It appears that fairness can only be investigated in combination with efficiency measures (like total throughput, delay, maximum rate etc.). But we can pick up another relevant comment from the original paper [4]. It was noted that the value of the fairness index depends on the so-called "allocation metric" where examples are given like: response time, response time per hop, throughput, throughput times hop, power (seen as ratio of throughput and response time), or "fraction of demand" (i.e. proportionality). Thus we promote an approach to consider a variety of fairness indices according to different allocation metrics at once. Instead of *allocation metric*, we will speak about *allocation features*, and for a set of features (each a vector by itself) we compute Jain's fairness index for each feature.

This way, an allocation is represented by a vector of fairness indices, each representing a fairness aspect of that allocation. Then, we can use this to select an allocation by searching the allocation, where the smallest index is maximized (so the lexicographic maxmin, or shortly leximin allocation). Moreover, by investigating the influence of the singular features on other performance measures of the leximin allocation we can learn about the "price" paid with regard to performance to make an allocation fair.

III. Multi-Jain Fairness Index

The original paper of Jain *et al* [4] was considering a wired network. With the emergence of wireless communication, the relevant allocation problem statement changes from rate allocation to channel allocation. This also means a change from a dividable good to an indivisible good. We will focus on the study of wireless channel allocation, and give a formal specification of the allocation task in the next subsection. Then we will list a number of allocation features that can be computed from a given allocation of channels to users.

A. Wireless Channel Allocation

At first we give the formal definition of the Wireless Channel Allocation problem, following [7]: Given a set of n users U (the agents of the allocation) and m channels C (the items of the allocation) and an $n \times m$ matrix CC of *channel coefficients*, i.e. reals from $[0, 1]$. A *channel allocation* is a mapping $A : C \rightarrow U$ where to each channel c_i with $i = 1, \ldots, m$ exactly one user u_j with $j = 1, \ldots, n$ is allocated. The notation is $u_j = A(c_i)$. An allocation is *feasible* if at least one channel is allocated to each user. The *performance* of user u_j in allocation A is $p_j = \sum_{i,A(c_i)=u_j} CC_{ji}$. The task of wireless channel allocation (WCA) is to find a feasible allocation a that "maximizes" the performances for all users. The additional task is to assign an effective meaning to "maximize." The WCA reflects a situation where a wireless infrastructure is composed of a Base Station BS and multiple Subscriber Stations SS, and we consider uplink traffic over one or multiple timeframes where the BS can simultaneously receive data from each user via SS. The "channel" here appears as a virtualization of physical transmission channels, either by channel bonding allowing for assigning more than one channel per timeframe to a user, or by repeated use of same channel(s) at different time slots for same user. The channel coefficients represents the knowledge of the BS about the particular channel states, based on measurements (e.g. using beacon signals) and/or prediction (for example if a user is moving then to predict future channel states). The WCA then is the scheduling task that has to be solved based on available channel state information. Note that standards like IEEE 802.11 (WLan) or IEEE 802.16 (WiMax) do not specify a scheduling itself and such procedures can be used to complement the standard specifications in a real-world application.

However, from economical point of view the WCA appears to be a specific case of fair allocation of indivisible goods.

B. Per-Entity Allocation Features

Given any allocation A (and using terminology of former subsection) that assigns channel c_j to user u_i with channel coefficient CC_{ij} we evaluate the allocation by several feature calculations. The features listed in the following are covering various aspects of the allocations that favorably would be equalized in a "perfect allocation" — even if such allocation is not feasible at all. It should also be noted that the particular selection of features here is partly for demonstration purposes of the general approach. One aspect that we do not quantify here is the channel distribution, for example: assuming the total channels are a smaller set of channels monitored over a number of time frames, then also per time frame allocation features could be considered. For example, an allocation of 5 users to 30 channels could also stand for five allocations of 5 users to 6 channels, one per time frame. However, we will not consider this aspect here. The chosen set of features, where each computes to a vector of - preferably - equal components, are as follows:

1) F_1: performance per user, or $f_i^{(1)} = p_i = \sum_{k,A(c_k)=u_i} CC_{ik}$, as it was mentioned in foregoing subsection. This vector has n components.
2) F_2: channels per user, or $f_i^{(2)} = |\{j \mid A(c_j) = u_i\}|$. It judges whether each user receives about same number of channels. This vector has n components.
3) F_3: equity, or $f_i^{(3)} = \sum_{k,A(c_k)=u_i} \frac{CC_{ik}}{\sum_l CC_{il}}$. The maximum performance of a user is seen as the case where the user would receive all channels. This feature represents the proportions of that maximum performance that each user receives. In case of equality, the allocation would be equitable - however, equity is hard to achieve due to the discrete nature of channels. This feature vector has n components.
4) F_4: delay, where feature $f_i^{(4)}$ is the difference between the last occurrence of user u_i in the allocation and the first appearance, plus 1. Thus, this measures mimics the time a user has to wait if the channels would be processed sequentially. Otherwise, it enforces equidistance in user allocations and might acquire a different meaning. Also here there are n components of this feature vector, one per user.
5) F_5: relative channel utilization. Here we compare the current allocation with the most favorable allocation of channel c_i to some user, i.e. the allocation to a user u_k where $CC_k i$ is maximal. This vector has m components.
6) F_6: absolute channel utilization, which is for channel c_i the channel coefficient CC_{ki} where $A(c_i) = u_k$. Also here, we have m components.
7) F_7: envyness, which compares for each pair of different users u_i and u_j the performance p_i for user u_i of current allocation A with the performance for user u_i in an allocation where users u_i and u_j change roles, i.e. where each allocation of user u_i becomes an allocation of user u_j and vice versa. Thus, we have an indication whether user u_i envies user u_j since the second performance is higher and user u_i would prefer to change with user u_j or not. The comparison is done by dividing performance in current allocation by performance in (i, j)-swapped allocation and the feature has $n(n-1)$ components, one for each (ordered) pair of users.

Once we have computed all feature vectors F_1 to F_7 for an allocation A, we construct the multi-Jain index as follows:

$$J(A) = \{J(F_1), J(F_2), \ldots, J(F_7)\}. \quad (3)$$

As it was mentioned before, from the set of vectors computed from all possible allocations an allocation can be found where the minimum component is maximized, the so-called leximin. We will consider this approach in the next section.

IV. Experiments

For demonstrating the approach we present the results of a few experiments. The general goal is to study the influence of the various features on the results. But before we have to consider the WCA problem instances. The only parameter of a WCA problem is the matrix of channel coefficients. The distribution of their values is strongly problem dependent. We take two cases: in first case (called "i.i.d. distribution" or "random CC" in the following) the channel coefficients are i.i.d. uniform random numbers from $(0, 1)$. This corresponds with a case of no specific knowledge about the connectivity situation in the network. The second case (called "radial distribution") refers to a simple scenario where the users are distributed in a squared area with the base station at its lower-left corner. The channel coefficients are assumed to only depend on the distance to the base station. We take i.i.d. uniform random x- and y-coordinates within a square of side length $1/\sqrt{2}$ (thus the maximum distance to the base station is 1). In this case, the channel coefficients for same user but different channels are all the same. We note that the second case is harder with regard to fairness. The reason is that there is no incentive to serve users that are more remote than the user which is closest to the base station. The service can only be enabled if fairness is taken into account. Moreover, in cases where the number of users is close to the number of channels, typical properties of a fair allocation like proportionality and envy-freeness cannot be achieved.

Now for the experiments: in a first experiment, we compare the different features with regard to minimality. It means if an allocation is selected by prescribed procedure, one component of the multi-Jain index must be the minimal one (actually the minimum that is maximized). We want to know which of the features serves the minimal component most often (and can thus be seen as the "driving" feature that promotes the equality of other features).

Table I
The column F_i lists the number of cases where $J(F_i)$ was the minimal component in the finally selected allocation, out of 100 random WCA problem instances and for different distributions of the channel coefficients.

	F_1	F_2	F_3	F_4	F_5	F_6	F_7
i.i.d. distribution	1	0	0	1	0	2	96
same without F_7	23	3	15	0	14	45	0
radial distribution	21	0	0	0	3	1	75
same without F_7	71	7	7	0	7	8	0

The results are shown in table 1 for the case of 8 channels and 5 users (since the problem domain grows exponentially, this is already a problem of moderate size). The leading role of envy-freeness can be clearly seen. For random CC, nearly all cases are driven by feature F_7, and for the radial distribution, the same happens in 75% of the cases. Second comes the performance, i.e. feature F_1. It is interesting to compare with a situation where feature F_7 is not considered. While in the random CC case all features now start to contribute to the maxmin selection, in the case of radial distribution F_1 takes over the role of F_7. Thus, achieving equality of features appears to be hardest for the feature representing mutual envyness degree, and next for the total allocation per user.

Table II
Relative loss against maximum performance by selecting according to a subset of features. Each row lists a subset where one feature was removed from the total set of features. The last line shows the expected value of performance loss for 1000 random allocations.

removed	i.i.d. distrib.	radial distrib.
F_1	0.631449	0.745444
F_2	0.607061	0.713038
F_3	0.591417	0.737509
F_4	0.595488	0.719371
F_5	0.597669	0.72336
F_6	0.579802	0.711416
F_7	0.5766	0.696861
random	0.581044	0.709548

In experiment 2 we continue to study the influence of the various features, but now on maximum achievable performance. The maximum achievable performance of a WCA is the allocation where each channel is allocated to one of the users with maximum CC. Of course, such an allocation might be infeasible, since there can be users that will never receive any channel (if they happen to never be maximal in any column of the CC matrix). So fair allocations will surely reduce the performance, but by which degree?
Table 2 shows some related results. Here, the same setting was used (5 users, 8 channels). In each case (one row in the table) the maxmin allocation for the set of features except one was selected and its performance compared with the maximum performance. The values shown are sampled over 30 independent runs. The last line gives the related value for 1000 random allocations.

It can be seen that here, except feature F_1 no feature has a significant influence on performance, and esp. there is no significant difference to a random allocation (we can skip the test of significance here). If we ignore the first feature, this is not so surprising: as it was mentioned before, the fairness index cannot resolve the problem of representing both, equality and efficiency at the same time. None of the used features directly represents efficiency (beyond the level that envy-freeness implies proportionality) and so it cannot be expected that selected allocations diverge much from other allocations.

But at the same time, the significant increase in relative performance when feature 1 is not included is notable. It is even more interesting since the ratio is larger for the case of radial distribution (in fact the more realistic scenario).

We can summarize the findings here by the statement that the most relevant feature for establishing fairness appears to be envy-freeness. On the one hand, the increase of the feature promotes the increase of other features, on the other hand, it does not significantly reduce total performance. The opposite claim is about user performance: it also drives the other features, but the price is a significant loss in performance, esp. in the case of radial distribution.

We have to note that these results are more in the sense of demonstrating the analysis that becomes available by the use of a multi-Jain index. In order to establish the facts in more details, a more thorough investigation of the model situation, incl. other relevant aspects has to be performed. This is the scope of future works.

We may conclude with a short discussion: the predominance of envy-freeness is not surprising: it is the most challenging feature of a fair allocation. Let us compare this with other common characteristics of fair allocations. In the following, we consider a general allocation problem where m goods are to be distributed among n users. Each user has values for each good, and these valuations may be different for each user (like the channel coefficients in the WCA problem).

1) Proportionality: an allocation is proportionate if among n users each user receives at least $1/n$-th of the total value of all goods according to her valuation. When summing up the relative valuations of the allocation for all users, the result will be at least 1. The issue here is that nothing is specified beyond the $1/n$-proportionality level: the remaining allocation can be such that all other goods go to one user only, and the allocation would still be proportional fair. In addition it can be easily seen that envy-freeness implies proportionality, while at the same time preventing such unfair allocations.
2) Equity: if proportionality means that each user receives at least $1/n$-th of her total valuation of all goods, equity means that the ratio is the same for all users. This demand can only be fulfilled if the goods are dividable, but it can be approximated in the case of indivisible goods. However, the value of equal ratios does not need to be at least $1/n$ — an allocation providing exactly a $1/1000$-th of their total values to each user is also equitable. Equity can only become strong in combination with other fairness properties.
3) Efficiency, or Pareto efficiency: it states that for a fair allocation, no user can gain more without making some other user worse off. This is in fact a relation, also known as Pareto dominance in the multi-criterion decision making domain, and it says that the fair allocation should be an element of the Pareto set of non-dominated solutions. The problem here, in addition to the fact that Pareto sets can become large with increasing problem dimension, is that a solution where one user receives everything and all other nothing is also efficient. The one user getting everything cannot get more, and in order to give something to any of the users that received nothings means to take away from the user that received everything. Also here, efficiency alone does not guarantee an accepted fairness and needs to be combined with other features.
4) Studies mention other properties of fair allocations and the procedures to find these allocations, giving reference to features related to "consistency" (fair allocations change with change of number of users or goods in a consistent way), "incentive" (or strategy-proofness) where there must be an advantage of confirming on a fair allocation and providing true valuations etc. We note that these features are procedure-dependent and the proposed framework of using a multi-fairness index cannot make any statement about their fulfillment without the exact specification of an allocation procedure.

This short discussion shows that envy-freeness is in fact an attractive feature, as it does not have any of the pitfalls of the other properties. However, it is often not possible to find an envy-free procedure. Therefore, the proposed approach gives a way to approach envy-freeness to some degree.

V. Summary

We have presented a multi-Jain fairness index to describe various fairness related aspects of an allocation. By the example of allocating wireless channels to users in a wireless network infrastructure, seven specific features have been specified. The standard Jain fairness index for each of these features, put together for all features, gives the multi-Jain fairness index. By means of a leximin procedure, an allocation can be selected where the smallest among the Jain fairness indexes takes a largest value. This extends the notion of an allocation where fairness is only achieved for a single allocation metric. By analyzing the influence of each feature on the leximin selection, conclusions can be drawn about the relevance of the different fairness aspects that are all represented at once in a multi-fairness index. In case of wireless channel allocation, this analysis gives indication that in fact the attempt to make user performance as equal as possible decreases available system performance, while the attempt to equalize envyness among all possible pairs of users appears to be of advantage to also increase all other aspects of fairness without reducing system performance. Future work will consider this aspect by taking dynamics of the system into account and also define time-dependent allocation features or to study fairness of multi-resource problems (with the starting point of Dominant Resource Fairness proposed in [8] and improvements as in [9]).

References

[1] B. Knaster, "Sur le probleme du partage pragmatique de h. steinhaus," in *Annales de la Societé Polonaise de Mathematique*, vol. 19, 1946, pp. 228–230.

[2] H. P. Young, *Equity: in theory and practice*. Princeton University Press, 1995.

[3] J. B. Steven and A. D. Taylor, "Fair division: From cake-cutting to dispute resolution," 1996.

[4] R. Jain, D. Chiu, and W. Hawe, "A quantitative measure of fairness and discrimination for resource allocation in shared computer system," Eastern Research Laboratory, Digital Equipment Corp., Tech. Rep., 1984.

[5] T. Lan, D. Kao, M. Chiang, and A. Sabharwal, "An axiomatic theory of fairness in network resource allocation," in *Proceedings of the 29th conference on Information communications*. IEEE Press, 2010, pp. 1343–1351.

[6] C. Joe-Wong, S. Sen, T. Lan, and M. Chiang, "Multi-resource allocation: Fairness-efficiency tradeoffs in a unifying framework," in *INFOCOM, 2012 Proceedings IEEE*. IEEE, 2012, pp. 1206–1214.

[7] R. Verschae, M. Koppen, and M. Tsuru, "A benchmark for evaluating approximated algorithms for fair wireless channel allocation," in *Intelligent Networking and Collaborative Systems (INCoS), 2012 4th International Conference on*. IEEE, 2012, pp. 178–184.

[8] A. Ghodsi, M. Zaharia, B. Hindman, A. Konwinski, S. Shenker, and I. Stoica, "Dominant resource fairness: fair allocation of multiple resource types," in *Proceedings of the 8th USENIX conference on Networked systems design and implementation*. USENIX Association, 2011, pp. 24–24.

[9] M. Köppen, K. Yoshida, and M. Tsuru, "Relational approaches to resource-aware multi-maxmin fairness in multi-valued resource sharing tasks," *International Journal of Space-Based and Situated Computing*, vol. 3, no. 2, pp. 91–101, July 2013.