Knaster Procedure for Proportional Fair Wireless Channel Allocation

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

***Abstract*—For the allocation of indivisible goods to a number of agents, appearing in many wireless infrastructure problems, common characteristics of fairness like proportionality, envy-freeness, or equity are hard to achieve or even impossible. As one of the few feasible allocation approaches, the Knaster procedure for item allocation is based on user bidding for items and adds an additional settlement step where users with higher benefit compensate users with lower benefit by payments. Here, the Knaster procedure is extended to a general allocation procedure, where the payments are considered as a measure for unfairness of a certain allocation. By minimizing the maximal payment, a fair state is uniquely specified. The approach is studied within the domain of wireless channel allocation and it is demonstrated that by following the proposed approach, deviations from proportional fairness can be kept small.**

I. Introduction

In a wireless infrastructure, many allocation tasks of indivisible goods appear, be it in the buffer management and relaying in OFDMA [1], [2], or in the allocation of radio channels to mobile stations [3]. The distribution of indivisible goods is characterized such that the goods to distribute cannot be physically divided (like one relay station) or would loose their value on any splitting attempt. However, the fair division of indivisible goods is essentially an ill-posed problem, as there is no guarantee that basic requirements of fairness like proportionality, envy-freeness or equitability/equity can ever be fulfilled.

To consider a simple example, take the distribution of 3 items among 2 agents. Both agents have their own values for these 3 items. Here, agent 1 assigns values 10, 2, and 4 to these items, while agent 2 assigns values 8, 1 and 1. It means item 1 has the highest value for both agents. But in any allocation of items, item 1 either goes to agent 1 or to agent 2. In either case, the other agent would not achieve a proportional allocation. The total value of all items for agent 1 is $10+2+4 = 16$ and for a proportional allocation, agent 1 would need to gain a value of at least $16/2 = 8$. For agent 2, a proportional allocation would have a value of at least $(8+1+1)/2 = 5$. Then, giving item 1 to agent 1 would leave less than value 2 for agent 2, and giving item 1 to agent 2 would leave less than 6 for agent 1 - no allocation would be proportional fair to all agents.

This simple example demonstrates the general problem with allocation of indivisible goods, especially in the case of smaller numbers of items or agents. It cannot be handled in this original problem representation. Therefore, most work on fair division has focused on divisible goods, and a growing number of successful algorithms have been proposed. With regard to indivisible goods, only a few numbers of approaches have been presented, and they share the common feature to redefine the original allocation problem in a suitable way.

One example is the method referred to as Luca's marker method in [4]. It requires that each agent splits the items into batches (or piles, the number of batches equals the number n of agents) of same total value. At first an independent referee writes the items as a list (from left to right on a piece of paper) in some order (note that while this order is decided independently, it has some effect on the outcome of the procedure). Then, each agent puts n marker lines between items on the list to indicate batches of same value for her. Once accomplished, the procedure goes in rounds agent by agent. In round 1, the agent that put the left-most marker line on the list receives everything up to this marker line. Then, these items are removed, as well as the marker lines of the selected agent. In round 2, among the remaining marker lines the second left-most is selected, and the corresponding agent receives everything between this marker line and her left-most marker line. Again, the items are removed as well as the marker lines of the selected agent. The procedure repeats for the number of agents. In the example before, assume that agent 1 puts the first marker between item 1 and item 2, while agent 2 puts the first marker line between items 2 and 3. This way, agent 1 indicates that item 1 has the same value to her than item 2 and 3 together, and agent 2 indicates same value of items 1 and 2 together as item 3. Then, the first marker line is with agent 1, and she receives item 1. The next marker line would be the one at the end of the list of 3 items, and agent 2 receives item 3. Item 2 remains unassigned.

This procedure ensures a proportional and envy-free allocation. Each agent receives a set of items that has a value of exactly $1/n$ with n the number of agents. By doing assignment in rounds, also envy-freeness can be easily seen. However, the price to achieve a fair solution is amending the original problem. It requires that the agents are able to specify equal-valued batches. In addition, the procedure will not allocate all items to agents, as in the example above, where item 2 was not assigned to any agent. In fact, an attempt to allocate these items to agents as well could easily break envy-freeness.

This longer discussion should illustrate the inherent problems of allocating indivisible items and that they can only be handled by amending the original problem. In a different approach, today known as Knaster procedure [5], the problem is modified by adding a settlement among the agents. In its original form, each agent makes a bid for the goods. Then, for each item it goes to the agent with the highest bidding (or a draw in case of same highest bids by more than one agent). After this, the agents receiving more than their proportional share compensate the agents receiving less. The total of received items bidding value and compensation can be shown to be proportional fair.

With regard to a wireless infrastructure, the most prominent problem of fair allocation is wireless channel allocation. In uplink traffic, and based on available Channel State Information (CSI) a base station has to allocate radio channels to mobile stations. In the abstract problem formulation, one channel can only be assigned to one mobile station. Now, each mobile station can utilize a radio channel as a communication means only to some degree. The most influencing factor here is just the distance to the base station, but also other physical characteristics of the transmission, obstacles, mobility etc. (and should be properly reflected in the CSI). Even worse, as the channel utilization falls with distance, under equal spatial distribution of agents, the number of mobile stations increases with distance (actually, with the square of distance). Thus, the allocation has to respect this fact and cannot favor agents (mobile stations) with high channel utilization. If the base station can provide communication channels with a bandwidth of 100MB/s, within a second, an agent closer to the base station might be able to uplink 80MB per second, while more remote agents can only send 10MB. In such a context, fairness appears to be related to inefficient total infrastructure utilization, but has to be maintained nevertheless as the primary goal of networking is to allow the communication among the agents.

A number of approaches to this problem have already been provided, for example bidding systems [6] as well as relational approaches [7], [8] and ad hoc procedures [9]. A common aspect of these approaches is the primary focus on fulfilling potential user demands and not so much on the manner by which the demand fulfillment is the result of combination of available resources. But then, for example comparing a transmission step for two agents of 80MB and 20MB per second does not relate to the channel states for these agents. It could mean that one agent receives 1 channel utilized with 80MB/s and the other agent 2 channels utilized with 10MB/s, but it could also mean that agent 1 receives 4 channels a 20MB/s and agent 2 one channel a 20MB/s. In the latter case, agent 1 would have access to notably more channels than agent 2, and it would be hard to explain the high allocation of agent 1 to agent 2.

On first glance, the Knaster procedure is not suitable for such a Wireless Channel Allocation (WCA) task, for at least two reasons.

1) The Knaster procedure selects by highest bid. This would always be the agent with highest channel utilization, i.e. agents close to the base station. In general, distant agents would never be selected.
2) A compensation is practically infeasible - how should one receive and utilize a compensation bandwidth?

However, in this paper, we want to present an approach to fair allocation based on this Knaster procedure. For short, we skip the bidding step and consider a compensation for any possible allocation. Then, this compensation is considered to represent unfairness of the allocation: the higher a compensating payment, the higher the excess of some agent against her proportional share. It means, by a minmax criterium we can select an allocation minimizing the highest compensating payment.

In the following section, basic issues about fairness and the Knaster procedure will be recalled. Section III then introduces the proposed concept of Knaster fairness, and some experimental validation of this approach and related discussions are provided in section IV before the paper concludes.

II. Fair Allocations

A. Wireless Channel Allocation

At first we give the formal definition of the Wireless Channel Allocation problem, following [10]: 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. Fairness Aspects

Among all possible allocations, we may focus on certain fairness aspects. In general, the following two aspects are commonly considered:

- *Proportionality*: among n agents, each agent (user) receives at least $1/n$-th of her total value. If we interpret the channel coefficients as such values, it would mean that for all agents U_i we have $p_i \geq 1/n \cdot \sum_j CC_{ij}$ (note that despite similar names, proportionality is not the same as *proportional fairness*, with the latter considering to maximize the product of utilities).

TABLE I. AVERAGE NUMBER OF PROPORTIONAL/ENVY-FREE FAIR ALLOCATIONS (100 SAMPLES EACH).

users	2	3	4	5	6
2 channels	0.39/0.39 of 4				
3 channels	0.9/0.9 of 8	0.74/0.18 of 27			
4 channels	1.83/1.83 of 16	2.3/0.39 of 81	1.42/0.09 of 256		
5 channels	3.53/3.53 of 32	6.56/1.13 of 243	8.61/0.19 of 1024	3.91/0.05 of 3125	
6 channels	7.0/7.0 of 64	18.24/3.83 of 729	26.46/0.33 of 4096	28.44/0.07 of 15625	10.65/0.1 of 46656

- Envy-Freeness: Here each agent A compares her total value with the value that another agent B would receive, if the allocation for B would be evaluated according to the valuation of agent A. In case B receives more than A here, it is said that agent A envies agent B. An envy-free allocation is an allocation where no agent envies any other agent.

We note that there are several other aspects of fair allocations that will not be considered in the following, for example Pareto efficiency (an agent can only receive more if some other agent receives less), strategy-proofness (an agent cannot gain by giving wrong information about her preferences) or equity (the relative gains of all agents according to their own valuation are all equal). The reason to neglect them (despite their undoubted importance) is that even the above two requirements are hard to achieve in a WCA setting.

To see this for proportionality, consider a standard situation where 4 channels are to be allocated to 3 users. The users have the same CC for all channels, assume these are 0.8 for user 1, 0.2 for user 2 and 0.1 for user 3. So, users 2 and 3 might be more distant from the base station than user 1. Now, a feasible allocation must assign one channel to each user, and the 4th channel to one of the 3 users, thus receiving 2 channels. Then for the other two users, the allocation of one channel is less than their proportional share of at least $4/3$ times their CC. Thus, proportional fairness is not possible. The same can be seen for envy-freeness.

To illustrate this in more detail, we report on a simple experiment. We iterate the number of channels from 2 to 6 and the number of users from 2 to the number of channels (since we cannot allocate for $n > m$ without leaving at least one user unserved). Then we average over 100 runs of: (1) setting all CC i.i.d. random values between 0 and 1, (2) computing the number of proportional and envy-free allocations among all possible allocations. Table I shows the results.

It can be clearly seen that the absolute number of proportional and envy-free allocations is rather small, and that there are always cases where no such allocation exists at all. The number of envy-free allocations also falls notably with increasing problem dimension. Thus demanding an allocation to be just proportional fair or envy-free is not sufficient, and additional characterizations are needed (note that proportionality follows from envy-freeness).

C. Fair Allocation by Knaster Procedure

The Knaster procedure was proposed more than 70 years ago [5]. The Introduction already gave a rough overview. Here we want to clarify the procedure by providing an example from which the general procedure should become clear.

TABLE II. EXAMPLE FOR THE KNASTER PROCEDURE TO ALLOCATE 4 ITEMS AMONG 3 AGENTS. FOR AGENT 3, THE ALLOCATION IS LESS THAN PROPORTIONAL, AND AGENT 2 RECEIVES A MUCH HIGHER VALUE THAN AGENT 1, IF COMPARED WITH THEIR IDENTICAL PROPORTIONAL SHARE. AGENT 2 COMPENSATES AGENTS 1 AND 3.

	Agents		
	Agent 1	Agent 2	Agent 3
Valuation			
Item 1	10	4	5
Item 2	6	8	1
Item 3	2	2	3
Item 4	3	7	3
Total valuation	21	21	12
Initial fair share	7	7	4
Item(s) received	Item 1	Items 2, 4	Item 3
Value received	10	15	3
Difference	3	8	-1
Share of surplus	3.33	3.33	3.33
Adjusted fair share	10.33	10.33	7.33
Settlement	-0.33	4.66	-4.33

An example for the Knaster procedure is shown in Table II. Here, 3 agents have to share 4 items, with valuations ("bids") as given in the table. An item is assigned to the highest bidder, and thus each agent receives some value in items. The row "Total valuation" shows the sum of item values per agent, and dividing by 3 (the number of agents) gives the least value for a proportional allocation (shown in the following row).

There is a difference between the value received and the value that would correspond with a proportional share (shown in row "Difference"). These per-agent differences are the base for the calculation of a compensation. At first, the total of all differences (here $3 + 8 + (-1) = 10$) is calculated. The value is called surplus. It will be equally shared among the agents to compute an adjusted fair share as shown in the next row. The adjusted fair share now is compared with the received value. In case that the received value is larger, the agent has to pay compensation to the other agents, and in case it is smaller, the agent will receive compensation. The total of all payments and payoffs is 0.

In the example in Table II agent 3 does not even receive a proportional allocation: value of received items is 3, but total value of all items would be 12 for all items, and $12/3 > 3$. Agent 2 has the same total value like agent 1, but receives a much higher total value (also owned to the fact that agent 2 receives two items). Even if agent 1 receives a proportional allocation, there is some advantage for agent 2. At the end, the settlement is that agent 2 has to compensate to both, agents 1 and 3 for receiving a larger share of items.

III. Knaster Fairness

As already mentioned in the introduction, the Knaster procedure cannot be used directly for the WCA problem. If the channel coefficients were to be interpreted as bids, then in all cases only agents closer to a base station would win, and should somehow compensate the other agents. There is no reasonable way to do this.

But there is a different aspect of this procedure that can give the base for achieving fair allocations:

1) There is no need to base the settlement on the largest bids only. The settlement can be calculated for any allocation of items to agents.
2) The settlement includes payments and payoffs. Large payments are an indication for unfairness, as it means that there are agents that have received a much higher value than other agents.

Independently, the reference point of the Knaster procedure is proportional fairness. In case of exact proportional fairness, all payments would be 0. Everything else indicates an imbalance in the allocation with respect to proportionality. We want to employ this aspect and seek an allocation where the maximal payment is minimized. This serves both aspects: small payment indicate an allocation near proportionality, and a balanced allocation as well. It also solves the problem mentioned in the foregoing section by selecting a single allocation.

The proposed Knaster fairness, given formally in terms of the WCA problem (i.e. agents are users and items are channels) goes as follows: at first, we define the maximal payment of the settlement for each allocation:

1) Given an allocation A with performance p, computed according to Section II.A.
2) Take the total of channel coefficients for each user: $t_i = \sum_{j=1}^{m} CC_{ij}$. Then, t_i/n is the initial fair share of user u_i.
3) Calculate the difference between initial fair share and user allocation: $\delta_i = p_i - t_i/n$.
4) Calculate the surplus $S = \sum_{i=1}^{n} \delta_i$.
5) Adjust the fair share by adding equal shares of the surplus for each user: $f_i = t_i/n + S/n$.
6) The final settlement then is for each user the difference between adjusted fair share and performance: $s_i = p_i - f_i$ (positive values indicate payments, negative values payoffs).
7) The maximal payment $s(A)$ of allocation A is the largest component of the vector s.

Then we select the allocation that minimizes $s(A)$ as Knaster fair allocation.

IV. Discussion

We consider the same WCA problem as in the foregoing section. Three users access 4 channels, user 1 with CC 0.8, user 2 with CC of 0.2 and user 3 with CC of 0.1. Also in this case, there is no proportionally fair allocation. Following the procedure to find the Knaster fair allocation, we yield the minimal payment for the allocation $A = (1, 2, 3, 3)$ with a performance vector $p = (0.8, 0.2, 0.2)$ and a settlement vector of $s = (-0.178, 0.022, 0.156)$. It appears to be close to a proportionally fair allocation, but one may wonder why its actually the weakest user (in the sense of having lowest channel utilization) that pays the bill. It can be explained by the fact that the offset to proportionality can better be reduced by using lower valued items. In the given example, one user will receive two channels. If it would be a user with large channel coefficient, the compensation would become larger than for a user with smaller channel coefficient. This might always happen when the number of users is a little bit smaller than the number of available channels.

TABLE III. Average payments of weakest and strongest user (10000 samples each).

problem dimension	weakest user	strongest user
users: 2 channels: 2	-0.0025	0.0025
users: 2 channels: 3	0.0009	-0.0009
users: 3 channels: 3	-0.011	0.012
users: 2 channels: 4	-0.0012	0.0012
users: 3 channels: 4	0.0065	-0.011
users: 4 channels: 4	-0.011	0.0077
users: 2 channels: 5	-0.0016	0.0016
users: 3 channels: 5	-0.0014	0.0008
users: 4 channels: 5	0.0004	-0.0068
users: 5 channels: 5	-0.0092	0.0038

We study the case in more detail by performing the following experiment:

1) Average over 10,000 times repeating:
2) Set a random CC with uniform random numbers from $[0, 1]$ for its components.
3) Find the user with the lowest total of CC (the weakest user) and the user with the highest total of CC (the strongest user).
4) Find the Knaster fair allocation.
5) Take the payments/payoffs (positive/negative value) of the weakest and the strongest user for averaging.

The result of the procedure is shown in Table III.

It can be seen that in general, the payments are rather small, at most 0.012. As in the example above, in case of m channels and $n = m - 1$ users usually the weakest user pays, so the observation on above example seem to be of general nature. In case of 2 users, the values must add to 0, but except the case of 3 channels (which is a case like just discussed) the weakest user generally receives a payoff.

We can summarize that the Knaster procedure appears to provide a good balance of the allocation with respect to proportionality, i.e. it takes the strength of each user's ability to utilize the allocated channels into account in a fair manner. This makes a difference to pure performance- or demand-based approaches that are commonly studied.

As a final comment, we can see that computational fairness is always reflecting a trade-off between maximality and equality. Each way of specifying fairness in a formal sense will either promote the former or the latter. For example, the popular Jain's Fairness Index

$$J(x) = \frac{(\sum_i x_i)^2}{n \cdot \sum_i x_i^2} \quad (1)$$

of a vector x of n elements x_i takes its maximal value in case all x_i are equal. However, if we allocate only half of the x_i values only, they are still equal and the fairness indicator is still 1. Even if the x_i become very small the value is still 1. Thus, this index is more or less exclusively focusing on equality and not on maximality. On the other hand, common approaches to fairness like maxmin fairness, proportional fairness, alpha-fairness etc. are all implied by Pareto dominance and thus are focusing on maximality and less on equality.

Within the context of this discussion we can see that Knaster fairness belongs more to fairness concepts that are focusing on equality than on maximality. But by taking proportionality as reference point it appears different from e.g. Jain's Fairness Index.

One can envisage a number of further works in this direction. Just to mention a few aspects:

1) Instead of proportionality, envy-freeness can be taken as a reference point. Then, in case of user A is envying user B, user B has to make a compensatory payment to user A. Minimizing the maximal payment would alleviate unfairness caused by envy agents.
2) Moreover, a joint settlement for proportionality and envy-freeness can be introduced as well.
3) Sampling over all possible allocations can rapidly increase the computational effort. Meta-heuristic approaches appear to be suitable to approximate the minmax payment and thus come close to the Knaster fair allocations.
4) Instead of taking the minimum of maximum payments, a vector relation between the settlement vectors can be used. Then, minimal allocations are defined as allocations, whose settlement vector is not in relation to any other settlement vector. A suitable relation extending the minimum of maximal payments would be majorization [11], promoted by the fact that the sum of settlements is 0. The minmax allocation would belong to the minimum set, but other allocations with small maximal payments could belong to the minimum set as well. This approach would allow to better reflect the spread of all payments and payoffs.

V. Conclusions

The Knaster procedure was initially proposed to solve the problem of distribution of indivisible goods by a bidding system, where each item goes to the strongest bidder and afterwards a settlement is done to compensate for unevenness in the distribution. In this form, it cannot be used for resource allocation problems like wireless channel allocation in a wireless infrastructure. However, it can serve as a base for approximating proportional fairness (proportional fairness here in the sense that each user receives at least $1/n$-th of her total value, where n is the number of users). We have proposed to use the Knaster procedure, especially its compensation calculation, for any allocation of goods to agents, and then select by a minmax criterion. In this way, the approach will select an allocation uniquely, with linear search effort (however, the search spaces do grow exponentially), and which comes close to proportional allocations. In addition, the approach can be extended in various ways, for example to solicit envy-free allocations in place of proportional fair allocations.

References

[1] A. Gotsis, D. Komnakos, and P. Constantinou, "Dynamic subchannel and slot allocation for ofdma networks supporting mixed traffic: upper bound and a heuristic algorithm," *Communications Letters, IEEE*, vol. 13, no. 8, pp. 576–578, 2009.

[2] D. Zhang, Y. Wang, and J. Lu, "On qos-guaranteed downlink cooperative ofdma systems with amplify-and-forward relays: Optimal schedule and resource allocation," in *Wireless Communications and Networking Conference, 2009. WCNC 2009. IEEE*. IEEE, 2009, pp. 1–5.

[3] M. Salem, A. Adinoyi, M. Rahman, H. Yanikomeroglu, D. Falconer, Y.-D. Kim, E. Kim, and Y.-C. Cheong, "An overview of radio resource management in relay-enhanced ofdma-based networks," *Communications Surveys & Tutorials, IEEE*, vol. 12, no. 3, pp. 422–438, 2010.

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

[5] 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.

[6] J. Sun, E. Modiano, and L. Zheng, "Wireless channel allocation using an auction algorithm," *Selected Areas in Communications, IEEE Journal on*, vol. 24, no. 5, pp. 1085–1096, 2006.

[7] M. Koppen, K. Yoshida, M. Tsuru, and Y. Oie, "Annealing heuristic for fair wireless channel allocation by exponential ordered-ordered weighted averaging operator maximization," in *Applications and the Internet (SAINT), 2011 IEEE/IPSJ 11th International Symposium on*. IEEE, 2011, pp. 538–543.

[8] M. Koppen, K. Yoshida, and M. Tsuru, "A generic approach to multi-fairness and its application to wireless channel allocation," in *Intelligent Networking and Collaborative Systems (INCoS), 2011 Third International Conference on*. IEEE, 2011, pp. 587–593.

[9] J. B. Ernst and M. K. Denko, "The design and evaluation of fair scheduling in wireless mesh networks," *Journal of Computer and System Sciences*, vol. 77, no. 4, pp. 652–664, 2011.

[10] 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.

[11] A. W. Marshall, I. Olkin, and B. C. Arnold, *Inequalities: theory of majorization and its applications*. Springer Science+ Business Media, 2011.