

Fairness Does Not Imply Satisfaction (Student Abstract)

Andrew Searns, Hadi Hosseini

Department of Computer Science
Rochester Institute of Technology
{abs2157, hhvcs}@rit.edu

Abstract

Fair division is a subfield of multiagent systems that is concerned with object distribution. When objects are indivisible, the Maximin Share Guarantee (MMS) is a desirable fairness notion; however, it is not guaranteed to exist. While MMS allocations may not always exist, a relaxation of MMS is guaranteed to exist. We show that there exists a family of instances for which this relaxation fails to guarantee the MMS value for all but a small constant number of agents.

Introduction and motivation

Fair division encompasses problems within the area of multiagent systems that deal with distributing a set of indivisible objects among a set of agents. These problems arise in practical domains such as course selection and inheritance division (Brams and Taylor 1996; Budish 2011). The gold standard for fair division is envy-freeness where each agent weakly prefers its allocation to those of any other agent. However, for the case of indivisible goods, envy-freeness is not guaranteed to exist: consider a single good which is desired by two agents. One desirable relaxation of envy-freeness is Maximin Share Guarantee (MMS) proposed by (Budish 2011). MMS is a generalization of the *cut-and-choose* protocol (Brams and Taylor 1996) where each agent's MMS guarantee is the amount of value it can guarantee by dividing the goods into bundles and then choosing a bundle last. A surprising construction by (Procaccia and Wang 2014) showed MMS allocations are not guaranteed to exist. In addition, computing MMS values is NP-complete (Woeginger 1997). As a result, recent research has explored multiple approximations of MMS where agents are guaranteed to receive a constant fraction of their MMS value (Garg and Taki 2019). Another MMS approximation technique is optimal-MMS for which an allocation is always guaranteed to exist by definition (Nguyen, Nguyen, and Rothe 2017; Aziz et al. 2017).

In contrast to most approximations of MMS, one can consider scenarios where it is critical that some agents receive their full MMS value. For example, senior college students

must be guaranteed seats in order to avoid graduation delays. We address the question of how many agents receive at least their MMS value in optimal-MMS allocations. Our results imply that no optimal-MMS algorithm can guarantee full MMS value for more than a constant number of agents.

The model

A fair division instance $\mathcal{I} = \langle N, M, V \rangle$ is a set of agents N with $|N| = n$, a set of indivisible goods M , and a valuation profile V . We assume that each agent's valuation function $v_i : 2^M \rightarrow \mathbb{R}_{\geq 0}$ is *additive*. An allocation is an n -partition of the elements $A = (A_1, \dots, A_n)$ where agent i receives bundle A_i . The MMS value of each agent is given by $\text{MMS}_i = \max_{(A_1, \dots, A_n)} \min_j v_i(A_j)$. An allocation satisfies MMS if every agent receives at least their MMS value i.e. $\forall i \in N, v_i(A_i) \geq \text{MMS}_i$.

For a given instance, the optimal-MMS value $\lambda^*(\mathcal{I})$ is the egalitarian optimal approximation of MMS that all agents can be guaranteed. That is $\lambda^*(\mathcal{I}) = \max_{(A_1, \dots, A_n)} \frac{v_i(A_i)}{\text{MMS}_i}$. As the set of possible allocations is finite, an optimal-MMS allocation always exists. Furthermore, an MMS allocation exists iff $\lambda^*(\mathcal{I}) \geq 1$ (Nguyen, Nguyen, and Rothe 2017).

Our results

We show that the notion of optimal-MMS may result in allocations where very few agents receive their MMS guarantees. Specifically, there exists a family of instances such that for any n , only a constant (at most 4) number of agents receive their MMS guarantee. We illustrate our construction for the case where $n = 6$ and $d = 3$ in Figure 1b.

Theorem 1. *For any $n \geq 3$ and any $d \leq \lfloor \frac{n}{2} \rfloor$, there exists an instance with only $O(dn)$ goods where any optimal-MMS allocation guarantees at most $\lceil \frac{n}{d} \rceil + 1$ agents their MMS value.*

Our construction relies upon tensors. An order d tensor S with dimensions $(n \times n \times \dots \times n)$ has elements indexed by a d -tuple (x_1, \dots, x_d) where each $x_i \in [n]$. The element at index (x_1, \dots, x_d) has value $S[x_1, \dots, x_d]$ or $S_{x_1 \dots x_d}$ for short. The $(d - 1)$ -order slices of S along dimension j are given by $S_j(1)$ to $S_j(n)$ where $S_j(i) = \{S_{x_1 \dots x_d} : x_j = i\}$. These slices are shown for an order

3 tensor in Figure 1a. We also define $S_{i;x_j=k}$ to be the element where all indices are i except index $x_j = k$. For example, $S_{i;x_1=n} = S_{n11\dots 1}$. Tensor addition is elementwise $((S + T)_{x_1,x_2,\dots,x_d} = S_{x_1,x_2,\dots,x_d} + T_{x_1,x_2,\dots,x_d})$. We now sketch the proof of Theorem 1.

Proof Sketch. All tensors in our construction are order d with dimensions $(n \times n \times \dots \times n)$.

Our construction first builds a tensor S whose elements are chosen so that the only sets of its elements which sum to 1 correspond to the slices of the tensor. The non-zero elements of S occur along the main diagonal and the diagonal of the last slices in each dimension. The elements along the main diagonal (blue and purple) have value close to 1. The remaining elements (red) ($S_{i;x_1=n}, S_{i;x_2=n}, \dots, S_{i;x_d=n}$) are small enough so that when paired with the related element on the main diagonal ($S_{i\dots i}$), exactly $d - 1$ of them are needed for the set to sum to 1, but they are large enough that taking them with a previous element on the main diagonal has sum more than 1.

We build a tensor T which ensures that every n -partition of the elements of the tensor $M = S + T$ where each set has sum 1 corresponds to aligned slices. The value of non-zero elements of T depend on powers of a small ϵ . The slices of T all have total sum of 0, and all negative elements of T occur only where S has positive value (red). Thus for small enough ϵ , M has only non-negative values. The key component to the construction of T is that exactly one of the (green) elements of $\{T_{i;x_1=(i+1)}, T_{i;x_2=(i+1)}, \dots, T_{i;x_d=(i+1)}\}$ is in the slice $T_j(i + 1)$ for all $j \in [d]$. When a partition picks a set based on slice $M_j(1)$, it also excludes one of $\{S_{1;x_1=n}, S_{1;x_2=n}, \dots, S_{1;x_d=n}\}$. This in turn enforces that the corresponding element $T_{1;x_j=2}$ is in the slice $T_j(2)$.

We next separate agents into d groups of at least two agents each: P_1, P_2, \dots, P_d . Agents within each group perturb M slightly so that they have a unique n -partition of the elements of M' where every set has sum 1. Specifically, agents in group P_j perturb the non-zero elements of the slice $M_j(6)$ so that $M_{nn\dots n}$ is the only element with positive perturbation (if all other perturbations are small enough, M' has only non-negative elements).

We now construct an instance $\mathcal{I} = \langle N, M, V \rangle$ where M is the set of goods with non-zero value from the tensor M' , and v_i is defined by the values of M' for each agent. Since every agent can partition M' so that there are n sets of elements which have sum 1, $\text{MMS}_i = 1$ for all agents. However, in order to receive value 1, each agent must take a slice $M_j(i)$ according to their group (P_j). Since only one set of slices may be taken at a time, only one group and the additional agent who receives the last slice (and thus element $M_{nn\dots n}$) receive their MMS value. An optimal-MMS allocation, guarantees $\lambda^*(\mathcal{I}) \geq 1 - \tilde{\epsilon}$ where $\tilde{\epsilon}$ is the largest negative perturbation made by agents to form M' . If $\tilde{\epsilon}$ is sufficiently small (e.g. $\tilde{\epsilon}$ is smaller than the difference in sum between any pair of sets of elements of M) then the only way all agents can receive value $1 - \tilde{\epsilon}$ is if all agents receive bundles corresponding to slices of M' . However, in this case, we see that only $\lceil \frac{n}{d} \rceil + 1$ agents receive their MMS for any optimal-MMS allocation.




Figure 1: Tensor visualizations.

Lastly, we observe that there are a total of $2dn - 3d + n = O(dn)$ entries of M' which have non-zero value. \square

To ensure that only a constant number of agents receive their MMS value in any optimal-MMS allocation, we simply group agents into $d = \lfloor \frac{n}{2} \rfloor$ groups so that each group has only 2 agents (one group will have 3 agents if n is odd). As only one group and a single additional agent is satisfied by taking slices of the order d tensor along dimension j , a total of 3 (or 4 if n is odd) agents will eventually receive their MMS value in any optimal-MMS allocation.

Theorem 2. For any $n \geq 3$, there exists an instance with only $O(n^2)$ goods where any optimal-MMS allocation guarantees 3 (4 if n is odd) agents their MMS value.

Conclusions and future work

While optimal-MMS is a desirable approximation of the MMS guarantee, it can perform poorly when considering the number of agents receiving their MMS values. One possible extension of this work is to determine how many agents can be guaranteed their MMS if we can sacrifice some agents. Another relevant future direction is to extend our results to chores where agents have negative values for each object.

Acknowledgments

This work was partially supported by NSF grant #1850076.

References

- Aziz, H.; Rauchecker, G.; Schryen, G.; and Walsh, T. 2017. Algorithms for max-min share fair allocation of indivisible chores. In *Thirty-First AAAI Conference on Artificial Intelligence*.
- Brams, S. J., and Taylor, A. D. 1996. *Fair Division: From cake-cutting to dispute resolution*. Cambridge University Press.
- Budish, E. 2011. The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. *Journal of Political Economy* 119(6):1061–1103.
- Garg, J., and Taki, S. 2019. An improved approximation algorithm for maximin shares. *arXiv preprint:1903.00029*.
- Nguyen, N.-T.; Nguyen, T. T.; and Rothe, J. 2017. Approximate solutions to max-min fair and proportionally fair allocations of indivisible goods. In *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems*, 262–271.
- Procaccia, A. D., and Wang, J. 2014. Fair enough: Guaranteeing approximate maximin shares. In *Proceedings of the fifteenth ACM conference on Economics and computation*, 675–692. ACM.
- Woeginger, G. J. 1997. A polynomial-time approximation scheme for maximizing the minimum machine completion time. *Operations Research Letters* 20(4):149–154.