The Complexity of Minimizing Envy in House Allocation

Extended Abstract

Jayakrishnan Madathil
University of Glasgow, UK
jayakrishnan.madathil@glasgow.ac.uk.

Neeldhara Misra
IIT Gandhinagar, India
neeldhara.m@iitgn.ac.in

Aditi Sethia
IIT Gandhinagar, India
aditi.sethia@iitgn.ac.in

ABSTRACT

The house allocation problem asks for m houses to be allocated among n agents so that every agent receives exactly one house. The preferences of the agents over houses can be modeled as weak orders, of which two special cases are strict rankings and binary valuations. Given an allocation ϕ of houses to agents, an agent a *envies* agent b if a receives a house $\phi(a)$ that they like strictly less than $\phi(b)$, the house allocated to b. The *amount of envy* experienced by an agent a is the number of agents b that it envies with respect to ϕ.

We consider the computational problem of finding allocations that minimize: the number of agents who are envious, the maximum envy experienced by any agent, or the total amount of envy experienced by all agents. We investigate the complexity of all three optimization objectives for both strict rankings as well as binary valuations. We show that these problems are FPT when parameterized by the number of houses and the number of agents. When parameterized by solution size, i.e, the value of the optimization objective, we demonstrate W-hardness in the first objective and para-NP-hardness for the last objective. We also consider these questions in the setting of restricted domains and also suggest practical approaches for these problems via ILP formulations.

KEYWORDS

House Allocation; Envy-Freeness; Egalitarian; Utilitarian; Kernel

ACM Reference Format:

Jayakrishnan Madathil, Neeldhara Misra, and Aditi Sethia. 2023. The Complexity of Minimizing Envy in House Allocation: Extended Abstract. In *Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023), London, United Kingdom, May 29 – June 2, 2023*, IFAAMAS, 3 pages.

1 INTRODUCTION

The *house allocation problem* can be thought of as a one-sided matching problem, where a set A of n agents express preferences over a set H of m houses — usually specified as either rankings (with or without ties) or as cardinal utilities.[1] An allocation is an assignment of houses such that every agent receives exactly one house, and every house is assigned to at most one agent.

[1] The setting of $0/1$ utilities, or "dichotomous preferences", where every agent either likes or dislikes a house, is a popular special case.

Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023), A. Ricci, W. Yeoh, N. Agmon, B. An (eds.), May 29 – June 2, 2023, London, United Kingdom.

In one-sided matching problems, the focus is typically on obtaining allocations that optimize for notions of efficiency such as Pareto optimality [1, 6], and rank maximality [7]. On the other hand, one may also view house allocation as a resource allocation problem — specifically, it can be thought of as a constrained version of fair division of indivisible goods, with the constraint being that that every agent receives exactly one item, as opposed to the typical setting where an agent may receive multiple items and even be empty-handed in valid allocations. In this context, the focus is on achieving *fairness*, which is often quantified in terms of *envy*. Given an allocation, say $\phi : A \rightarrow H$, an agent a envies a' if she values $\phi(a')$ more than $\phi(a)$. Finding "envy-free" allocations, i.e, ones where no agent envies another, is one of the main goals in fair division, although it's easily seen to be unattainable in general, motivating a host of relaxed notions of fairness. Nguyen and Rothe [9] and Shams et. al. [10] looked at minimizing the envy when envy-free allocations do not exist.

In the context of the house allocation problem, note that if $n = m$, an envy-free allocation exists if and only if there is a perfect matching in the following bipartite graph: introduce a vertex for every agent and every house, and let the vertex corresponding to an agent be adjacent to all the houses that she values more than any other house. Since every house must be assigned when $n = m$, when an agent is assigned anything short of her best option, she will be envious. Therefore, the existence of an envy-free allocation can be determined efficiently in this situation using standard algorithms for checking if a perfect matching exists.

The question is less obvious when $n < m$, i.e, when there are more houses than agents. Indeed, one could work with the same bipartite graph, but it is possible for the house allocation instance to admit an envy-free allocation even though the bipartite graph does not have a perfect matching. Consider a situation with three houses and two agents, where both agents value one house above all else, and the other two equally. While the graph only captures the contention on the highly valued house, it does not lead us directly to the envy-free allocation that can be obtained by giving both agents the houses that they value relatively less (but equally). It turns out that the question of whether an envy-free allocation exists can in fact be determined in polynomial time even when $n < m$, by an algorithm of Gan et. al. [4] that involves iteratively removing subsets of contentious houses. They also show that an envy-free assignment exists with high probability if the number of houses exceeds the number of agents by a logarithmic factor. Further, Aigner-Holev and Segal-Halevi [2] study the relaxed variant of assigning at most one house to every agent and give an $O(m\sqrt{n})$ algorithm for finding an envy-free matching of maximum cardinality under binary utilities.

<table>
<tr><th rowspan="3"></th><th colspan="5">Cardinal</th><th rowspan="3">Rankings</th></tr>
<tr><th rowspan="2">General</th><th colspan="4">Binary</th></tr>
<tr><th>General</th><th>Extremal Intervals</th><th>d = 1</th><th>d = 2</th></tr>
<tr><td>OHA</td><td>NP-Complete
(by implication)</td><td>NP-Complete (†)
from Clique</td><td rowspan="3" colspan="2">P</td><td>NP-Complete
from Independent Set</td><td>NP-Complete
from Balanced Biclique</td></tr>
<tr><td>EHA</td><td colspan="2">NP-Complete
(by implication)</td><td>NP-Complete (⋆)
from Independent Set</td><td>NP-Complete (⋆)
from Multi-Colored Independent Set</td></tr>
<tr><td>UHA</td><td colspan="2">?</td><td colspan="2">?</td></tr>
</table>

Table 1: A partial summary of our results. Here, d denotes the maximum number of houses approved by any agent. The results marked with a $\star$ refer to reductions that imply hardness even when the standard parameter is a constant, while the result marked with a † is a FPT reduction and also implies W[1]-hardness in the standard parameter.

When an envy-free allocation does not exist at all, a natural question is to find an allocation that "minimizes envy". The problem of finding allocations that minimize the number of agents who experience envy has been articulated and studied in the setting of cardinal utilities. In particular, Kamiyama et. al. [8] showed that it is NP-complete to find allocations minimizing the number of envious agents, even for binary utilities, and it is hard to approximate for general utilities. Beynier et. al. [3] studied a local variant where an agent can envy only those agents who are connected to her in a given social network. In a more recent work, Hosseini et. al. [5] considered minimizing the aggregate envy in the localized setting.

2 OUR CONTRIBUTIONS

We propose to study the issue of "minimizing envy" from a broader perspective, and to this end we consider three natural measures of the "amount of envy" created by an allocation:

(1) the total number of agents who experience envy;

(2) the envy experienced by the *most* envious agent, where envy per agent is simply the number of agents that she is envious of;

(3) the total amount of envy experienced by all agents.

While these notions are natural and standard for fair division problems, as far as we know they have not been considered explicitly in the context of house allocation problems. We refer to the questions of finding allocations that minimize these three measures of envy as the Optimal House Allocation (OHA), Egalitarian House Allocation (EHA), and Utilitarian House Allocation (UHA) problems, respectively. Table 1 contains a summary of our main results and how they are obtained.

We restrict our setting to binary utilities. We show that OHA is NP-complete even on instances where every agent values at most two houses. Further, it is W[1]-hard when parameterized by the number of envious agents. These results require two separate reductions, both from Clique, where the latter result can be viewed as a reformulation of the NP-completeness result of [8]. We also show that the reduction of [8] that was used to show the hardness of approximation of Optimal House Allocation in the setting of general cardinal utilities can be suitably adapted to establish that OHA is NP-complete in the setting rankings (with ties).

EHA turns out to be NP-complete on instances where every agent values at most two houses *and* every house is approved by a constant number of agents by a reduction from Independent Set on cubic graphs. We also show that EHA is NP-complete when the preferences are specified as rankings (without ties) by an intricate reduction from the Multi-Colored Independent Set problem. In both cases, we achieve hardness when the maximum allowed envy is just *one*, establishing that the problem is para-NP-hard in the standard parameter.

We show that both OHA and EHA are FPT when parameterized by the total number of house types and agent types, which intuitively correspond to the number of distinct houses and agents, a parameter that is potentially *much* smaller than $m + n$. This is obtained using an ILP formulation with a bounded number of variables, a result that is also of independent practical value. We also show, using a popular technique known as the expansion lemma, that all three problems admit a linear kernel in n, and can be solved efficiently when agent preferences are "extremal".

On the experimental side, we implemented the ILP for OHA and EHA over synthetic datasets of house allocation problem generated uniformly at random, using Gurobi Optimizer version 9.5.1. The average was taken over 100 trials for each instance. For a fixed number of houses and agents, as the number of agent types, $n^\star$ increases, the number of envious agents and the maximum envy decreases. Instances with identical valuations ($n^\star = 1$) seem to admit more envy than the other extreme ($n^\star = n$). This is due to more contention on the specific subset of goods when $n^\star = 1$. On the contrary, for instances with $m^\star = 1$, envy-free allocations always exist. Indeed, when $m^\star = 1$, all houses are of the same type, which means that an agent either likes all the houses, or dislikes all of them and in either case, she is envy-free no matter which house she gets.

3 CONCLUSION

We introduced and studied three kinds natural quantifications of envy to be minimized in the setting of house allocation: OHA, EHA, and UHA. We leave several questions related to UHA open. Also, finding interesting classes of structured input — beyond extremal intervals — for which these problems are tractable is an important direction for future work.

ACKNOWLEDGEMENT

Jayakrishnan Madathil is supported by the Engineering and Physical Sciences Research Council [EP/V032305/1]. All authors are grateful for support from IIT Gandhinagar.

REFERENCES

[1] David J Abraham, Katarína Cechlárová, David F Manlove, and Kurt Mehlhorn. 2004. Pareto optimality in house allocation problems. In *International Symposium on Algorithms and Computation*. Springer, 3–15.

[2] Elad Aigner-Horev and Erel Segal-Halevi. 2021. Envy-free Matchings in Bipartite Graphs and their Applications to Fair Division. *Information Sciences* (Nov 2021). https://doi.org/10.1016/j.ins.2021.11.059

[3] Aurélie Beynier, Yann Chevaleyre, Laurent Gourvès, Ararat Harutyunyan, Julien Lesca, Nicolas Maudet, and Anaëlle Wilczynski. 2019. Local Envy-Freeness in House Allocation Problems. *Autonomous Agents and Multi-Agent Systems* 33 (09 2019). https://doi.org/10.1007/s10458-019-09417-x

[4] Jiarui Gan, Warut Suksompong, and Alexandros A. Voudouris. 2019. Envy-freeness in house allocation problems. *Mathematical Social Sciences* 101 (2019), 104–106.

[5] Hadi Hosseini, Justin Payan, Rik Sengupta, Rohit Vaish, and Vignesh Viswanathan. 2023. Graphical House Allocation. In *22nd International Conference on Autonomous Agents and Multiagent Systems, AAMAS.*

[6] Aanund Hylland and Richard Zeckhauser. 1979. The efficient allocation of individuals to positions. *Journal of Political Economy* 87, 2 (1979), 293–314.

[7] Robert W. Irving, Telikepalli Kavitha, Kurt Mehlhorn, Dimitrios Michail, and Katarzyna E. Paluch. 2006. Rank-Maximal Matchings. *ACM Trans. Algorithms* 2, 4 (oct 2006), 602–610. https://doi.org/10.1145/1198513.1198520

[8] Naoyuki Kamiyama, Pasin Manurangsi, and Warut Suksompong. 2021. On the Complexity of Fair House Allocation. *Operations Research Letters* 49, 4 (2021), 572–577.

[9] Trung Thanh Nguyen and Jörg Rothe. 2013. How to Decrease the Degree of Envy in Allocations of Indivisible Goods. In *Algorithmic Decision Theory: Third International Conference, ADT 2013, Bruxelles, Belgium, November 12-14, 2013, Proceedings* (Bruxelles, Belgium). Springer-Verlag, Berlin, Heidelberg, 271–284. https://doi.org/10.1007/978-3-642-41575-3_21

[10] Parham Shams, Aurélie Beynier, Sylvain Bouveret, and Nicolas Maudet. 2021. Minimizing and balancing envy among agents using Ordered Weighted Average. In *7th International Conference on Algorithmic Decision Theory*. Toulouse, France. https://hal.archives-ouvertes.fr/hal-03336026