Fine Grained Complexity of Fair and Efficient Allocations

Doctoral Consortium

Aditi Sethia
Indian Institute of Technology
Gandhinagar, India
aditi.sethia@iitgn.ac.in

ABSTRACT

Fair Division is a flourishing field that has garnered a lot of attention in recent times. Allocating a set of valuable resources *fairly* among interested agents along with guaranteeing everyone's satisfaction is a crucial task with a wide range of applications, both routine and high-stakes. This paper presents our existing and ongoing work in the following directions – a) minimizing envy when absolute envy-freeness is unachievable b) identifying the structured instances where fair and efficient allocation problems admit fast algorithms c) quantifying the trade-off between fairness (EF1/EQ1) and efficiency notions (social welfare functions) of an allocation.

KEYWORDS

Computational Social Choice, Fair Division, Resource Allocation

ACM Reference Format:
Aditi Sethia. 2023. Fine Grained Complexity of Fair and Efficient Allocations: Doctoral Consortium. 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 problem of fair division concerns with partitioning a set of valuable resources among interested agents having (possibly heterogenous) preferences over these resources. Allocating scarce medical resources like organs for transplantation, vaccines during a pandemic; dividing territory and resources in a post-conflict scenario are a few high-stakes applications. The theoretical formalism of fair division was initiated by Steinhaus [30] and since then, it has witnessed a vast literature in mathematics, economics and recently computer science [9, 10, 26].

An instance of the fair division problem is specified by a tuple $\langle N, M, \mathcal{V} \rangle$, where $N = \{1, 2, \ldots, n\}$ is a set of $n \in \mathbb{N}$ *agents*, $M = \{g_1, g_2, \ldots g_m\}$ is a set of m indivisible *goods*, and $\mathcal{V} := \{v_1, v_2, \ldots, v_n\}$ is the *valuation profile* consisting of each agent's valuation for the goods. For any agent i, its valuation function $v_i : 2^M \rightarrow \mathbb{N} \cup \{0\}$ specifies its value (or *utility*) for every subset of goods in M. These are said to be *additive* if the value of a set of goods S is simply the sum of the value of the goods in S. An allocation $A = (A_1, \ldots, A_n)$, which is a partition of m goods into n agents and valued at $v_i(A_i)$ by any agent i, is said to be:

- *envy-free* (EF) if every agent values her bundle the best and does not wish to swap with anyone else's. That is, for any pair of agents i and j, we have $v_i(A_i) \geq v_i(A_j)$ [16, 33].
- *equitable* (EQ) if every agent derive equal value from their respective bundles. That is, for any pair of agents i and j, we have $v_i(A_i) = v_j(A_j)$ [15].

When the resources are divisible (money, land, time etc.), EF is guaranteed to exist [8, 31]. But in case of indivisible resources, the agent deprived of the valuable resource is bound to envy the one who gets it. In this realm of non-existence of absolute fairness, the natural objective is to minimize the amount of envy in the allocation. Chevaleyre et. al. [14], Nguyen and Rothe [27] and Shams et. al. [29] have looked at the complexity and approximability of minimizing the degree of envy. As a part of this thesis, we study the complexity of minimizing envy in the context of House Allocation problems, where every agent must get exactly one resource–the house.

Since EF and EQ do not have existential guarantee, in order to define the notions that do exist, there have been advancements towards relaxed versions of fairness. To this end, *Envy-freeness up to one good* (EF1) [11, 23] bounds the envy in the sense that the envious agent can hypothetically remove one good from the envied bundle and then, if she prefers her own bundle, she is said to be envy-free up to one good. Equitability upto one good (EQ1) [18, 19] is defined analogously. Both EF1 and EQ1 allocations exist and are efficiently computable under additive valuations.

While fairness is a compelling notion, it is not the only object of interest. Orthogonal to the notion of every agent getting a fair share, the efficiency of an allocation is determined by the total happiness/welfare of the agents. This is captured by the notions like Pareto optimality (no other allocation makes an agent happier without worsening someone else), Utilitarian welfare (sum of the agent utilities), Egalitarian welfare (the utility of the least happy agent), Nash welfare (geometric mean of the agent's utilities).

There has been significant progress towards the question whether it is possible to simultaneously achieve fairness and efficiency. Even if fair and efficient allocations exist in a setting, the computation part is hard [5, 17]. As a part of this thesis, we propose to identify structures in the valuations of agents that facilitate polynomial time algorithms for finding the desired allocation.

Caragiannis et. al. [13] showed that when valuations are additive, all MNW allocations are both EF1 and PO, indicating the compatibility of EF1 and PO. On the other hand, fairness and welfare maximization may not fare well together. For instance, allocating all the items to an agent who values them the most seems to be an efficient allocation (both PO and maximizes utilitarian welfare) but is blatantly unfair. Not allocating anything to anyone is fair in the sense that no agent envies anyone, but it is inefficient and has worst possible welfare. Recent works [2, 4] have looked at the hardness and approximations for computing the utilitarian optimal allocation among the set of fair allocations. The prominent question here is to

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.

quantify how much welfare is sacrificed in order to achieve a fair question. This is captured by the *price of fairness* [7, 12], which is defined as the supremum ratio of the maximum welfare achievable by any allocation to the maximum welfare achievable by a *fair* allocation. The bounds on price of EF1, Nash welfare and other fairness criterion are discussed in fairly recent papers [2, 3, 6]. These works capture the loss of utilitarian welfare under the said fairness constraint. The dimension that is not captured in the previous works is the price one has to pay in terms of Nash and other general welfare functions in order to achieve EF/EQ allocations. An immediate corollary from the previous works is that the price of EF1 for Nash welfare is essentially 1, that is to say that Nash welfare remains intact while achieving EF1. While this is not the really case with the other welfare/fairness combinations, we propose to close this gap and study the price of fairness in this broader sense.

Based on above discussed notions, the work in this thesis revolves around the following narratives, which are discussed in detail in the respective sections:

- Minimizing the envy when absolute envy-freeness cannot be guaranteed (Section 2).
- Identifying the parameters (agent/item-types) and structures in the valuation profile (degeneracy/distance) that facilitate the computation of fair and efficient allocations (Section 3).
- Quantifying the fairness and welfare trade-off for general welfare functions (Section 4).

2 ENVY-MINIMIZATION

When EF allocations do not exist, a natural question is to find an allocation that minimizes "envy". Kamiyama et. al. [22] showed that it is hard to find allocations that minimize the number of envious agents, even for binary utilities. Under the relaxed variant of assigning at most one house to every agent, Horev et. al. [1] gave an algorithm for finding a maximum envy-free matching under binary utilities. In a more recent work, Hosseini et. al. [21] consider minimizing the aggregate envy where an agent envies only those who are connected to her in a given social network.

In our recent work [24], we look at minimizing three different notions of envy – a) the total number of agents who experience envy, b) the envy experienced by the *most* envious agent, where the amount of envy experienced by an agent is simply the number of agents that she is envious of, and c) the total amount of envy experienced by all agents. We propose efficient algorithms for instances with binary valuations that admit extremal structure. We show tractability, for the former two notions, parameterised by agent/house-types, which intuitively correspond to the number of distinct agents/houses, a parameter that is potentially *much* smaller than number of agents/houses. This is obtained using an ILP formulation with a bounded number of variables, a result that is also of independent practical value. Extending the interesting classes of structured input for which these problems are tractable is an important and interesting direction for future work.

3 STRUCTURAL ASPECTS

Sandomirskiy and Halevi [28] demonstrated that finding fair allocations with minimum sharing of goods is tractable when valuations are non-degenerate, a notion which captures scenarios that are "far from identical". This result holds for any fixed number of agents. Building up on this, in a recent work [25], we show that the usefulness of non-degeneracy does not scale to the setting of many agents. In particular, we demonstrate that the problem of finding (fractionally) PO and EF allocations is hard even for instances with constant degeneracy and no sharing. The idea of capturing the 'extent of similarity in the valuations' and the dependence of existence of desired allocations on this parameter is worth-exploring. One way this can be captured is by considering some well-defined notion of *distance* between two valuation vectors. It is clear that when this distance is large, EF and PO allocations can be easily constructed. To capture the extent to which it helps in the computation of desired allocations is an interesting direction of future work.

4 FAIRNESS AND WELFARE TRADE-OFF

A social welfare function is an aggregation of the agent's utilities under a given allocation. For any $p \in (-\infty, 1]$, the *p-mean welfare* is the generalized p-mean of utilities of agents under A, i.e., $W_p(A) := \left(\frac{1}{n}\sum_{i\in N} v_i^p(A_i)\right)^{\frac{1}{p}}$. With $p = 1$, the p-mean welfare function corresponds to the Utilitarian welfare. As $p \rightarrow 0$ and $p \rightarrow \infty$, it tends to the Nash and Egalitarian welfare respectively.

Sun et. al. [32] proposed bounds (as a function of number of agents) on the price of EQ1/EQX in terms of egalitarian/utilitarian welfare. Building up on this, in an ongoing work, we propose bounds on price of EQ1 in the context of generalized p-mean welfare functions (in terms of agent-types). It is interesting to see that under binary additive valuations, although EF1 is compatible with Nash Welfare [20], and the price of EF1 in this setting (and hence, for generalized p-means) is essentially 1, but for EQ1, there are instances where loss of p-mean welfare is inevitable, and hence we have non-trivial lowerbounds.

There are quite a few compelling future directions here to pursue. First is to explore the fairness-efficiency trade-off in the setting of chores and/or a mixture of goods and chores. An interesting dual question to price of fairness is the price of welfare-maximization, which is to quantify how much fairness is to be sacrificed to achieve a welfare optimal allocation. Bei et. al. [6] look at the bounds on Utilitarian social welfare achieved by PO allocations, comparing the trade-off of two efficiency notions. Along similar lines, an interesting direction would be to explore the trade-off of two fairness notions, say capturing the extent of loss of EQ while ensuring EF.

5 CONCLUSION

We hope to provide new insights on scenarios where desired fair and efficient algorithms exist and are computable. Tractability parameterised by agent-types and item-types in the context of house allocation algorithms builds hope for exploring these parameters in more general settings. Decoding the structure hidden in the valuations of agents seems to be a promising forward step. Whether low price of fairness guarantee can be instrumental in designing such algorithms is also one direction, which we feel, is worth exploring.

REFERENCES

[1] 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

[2] Haris Aziz, Xin Huang, Nicholas Mattei, and Erel Segal-Halevi. 2022. Computing welfare-Maximizing fair allocations of indivisible goods. *European Journal of Operational Research* (2022). https://doi.org/10.1016/j.ejor.2022.10.013

[3] Siddharth Barman, Umang Bhaskar, and Nisarg Shah. 2020. Optimal Bounds on the Price of Fairness for Indivisible Goods. In *Proceedings of the 16th International Conference on Web and Internet Economics*. 356–369.

[4] Siddharth Barman, Ganesh Ghalme, Shweta Jain, Pooja Kulkarni, and Shivika Narang. 2019. Fair Division of Indivisible Goods Among Strategic Agents. In *Proceedings of the 18th International Conference on Autonomous Agents and Multi-Agent Systems* (Montreal QC, Canada) *(AAMAS '19)*. International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 1811–1813.

[5] Siddharth Barman, Sanath Kumar Krishnamurthy, and Rohit Vaish. 2018. Finding Fair and Efficient Allocations. In *Proceedings of the 2018 ACM Conference on Economics and Computation* (Ithaca, NY, USA) *(EC '18)*. Association for Computing Machinery, New York, NY, USA, 557–574. https://doi.org/10.1145/3219166.3219176

[6] Xiaohui Bei, Xinhang Lu, Pasin Manurangsi, and Warut Suksompong. 2021. The Price of Fairness for Indivisible Goods. *Theory of Computing Systems* (2021), 1–25.

[7] Dimitris Bertsimas, Vivek F. Farias, and Nikolaos Trichakis. 2011. The Price of Fairness. *Operations Research* 59, 1 (2011), 17–31. https://EconPapers.repec.org/RePEc:inm:oropre:v:59:y:2011:i:1:p:17-31

[8] Steven J Brams and Alan D Taylor. 1995. An envy-free cake division protocol. *The American Mathematical Monthly* 102, 1 (1995), 9–18.

[9] Steven J Brams and Alan D Taylor. 1996. *Fair Division: From cake-cutting to dispute resolution*. Cambridge University Press.

[10] Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D Procaccia. 2016. *Handbook of computational social choice*. Cambridge University Press.

[11] Eric Budish. 2011. The Combinatorial Assignment Problem: Approximate Competitive Equilibrium from Equal Incomes. *Journal of Political Economy* 119, 6 (2011), 1061–1103.

[12] Ioannis Caragiannis, Christos Kaklamanis, Panagiotis Kanellopoulos, and Maria Kyropoulou. 2012. The Efficiency of Fair Division. *Theory of Computing Systems* 50, 4 (2012), 589–610.

[13] Ioannis Caragiannis, David Kurokawa, Hervé Moulin, Ariel D Procaccia, Nisarg Shah, and Junxing Wang. 2019. The unreasonable fairness of maximum Nash welfare. *ACM Transactions on Economics and Computation (TEAC)* 7, 3 (2019), 1–32.

[14] Yann Chevaleyre, Ulle Endriss, Sylvia Estivie, and Nicolas Maudet. 2007. Reaching Envy-Free States in Distributed Negotiation Settings. In *Proceedings of the 20th International Joint Conference on Artifical Intelligence* (Hyderabad, India) *(IJCAI'07)*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 1239–1244.

[15] Lester E Dubins and Edwin H Spanier. 1961. How to Cut a Cake Fairly. *The American Mathematical Monthly* 68, 1P1 (1961), 1–17.

[16] Duncan Foley. 1967. Resource Allocation and the Public Sector. *Yale Economic Essays* (1967), 45–98.

[17] Rupert Freeman, Sujoy Sikdar, Rohit Vaish, and Lirong Xia. 2019. Equitable Allocations of Indivisible Goods. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence* (Macao, China) *(IJCAI'19)*. AAAI Press, 280–286.

[18] Rupert Freeman, Sujoy Sikdar, Rohit Vaish, and Lirong Xia. 2019. Equitable Allocations of Indivisible Goods. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence*. 280–286.

[19] Laurent Gourvès, Jérôme Monnot, and Lydia Tlilane. 2014. Near Fairness in Matroids. In *Proceedings of the 21st European Conference on Artificial Intelligence*. 393–398.

[20] Daniel Halpern, Ariel D. Procaccia, Alexandros Psomas, and Nisarg Shah. 2020. Fair Division with Binary Valuations: One Rule to Rule Them All. In *Web and Internet Economics: 16th International Conference, WINE 2020, Beijing, China, December 7–11, 2020, Proceedings* (Beijing, China). Springer-Verlag, Berlin, Heidelberg, 370–383. https://doi.org/10.1007/978-3-030-64946-3_26

[21] 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*.

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

[23] Richard J Lipton, Evangelos Markakis, Elchanan Mossel, and Amin Saberi. 2004. On Approximately Fair Allocations of Indivisible Goods. In *Proceedings of the 5th ACM Conference on Electronic Commerce*. 125–131.

[24] Jayakrishnan Madathil, Neeldhara Misra, and Aditi Sethia. 2023. On the Complexity of Minimizing Envy in House Allocation (Extended Abstract). In *22nd International Conference on Autonomous Agents and Multiagent Systems, AAMAS*.

[25] Neeldhara Misra and Aditi Sethia. 2021. Fair Division Is Hard Even for Amicable Agents. In *SOFSEM 2021: Theory and Practice of Computer Science: 47th International Conference on Current Trends in Theory and Practice of Computer Science, SOFSEM 2021, Bolzano-Bozen, Italy, January 25–29, 2021, Proceedings* (Bolzano-Bozen, Italy). Springer-Verlag, Berlin, Heidelberg, 421–430. https://doi.org/10.1007/978-3-030-67731-2_31

[26] Hervé Moulin. 2004. *Fair division and collective welfare*. MIT press.

[27] Trung Thanh Nguyen and Jörg Rothe. 2014. Minimizing Envy and Maximizing Average Nash Social Welfare in the Allocation of Indivisible Goods. *Discrete Applied Mathematics* 179 (2014), 54–68.

[28] Fedor Sandomirskiy and Erel Segal-Halevi. 2019. Fair Division with Minimal Sharing. *CoRR* abs/1908.01669 (2019). arXiv:1908.01669 http://arxiv.org/abs/1908.01669

[29] Parham Shams, Aurélie Beynier, Sylvain Bouveret, and Nicolas Maudet. 2019. Minimizing and balancing envy among agents using Ordered Weighted Average. In *13èmes Journées d'Intelligence Artificielle Fondamentale (JIAF 2019)*. Association française pour l'Intelligence Artificielle (AFIA), Toulouse, France, 175–182. https://hal.science/hal-02302957

[30] H Steinhaus. 1948. The problem of fair division. *Econometrica* 16 (1948), 101–104.

[31] Walter Stromquist. 1980. How to Cut a Cake Fairly. *The American Mathematical Monthly* 87, 8 (1980), 640–644. https://doi.org/10.1080/00029890.1980.11995109 arXiv:https://doi.org/10.1080/00029890.1980.11995109

[32] Ankang Sun, Bo Chen, and Xuan Vinh Doan. 2022. Equitability and Welfare Maximization for Allocating Indivisible Items. *Autonomous Agents and Multi-Agent Systems* 37, 1 (dec 2022), 45. https://doi.org/10.1007/s10458-022-09587-1

[33] Hal Varian. 1974. Equity, efficiency, and envy. *Journal of Economic Theory* 9 (1974), 63–91.