Fair Division is Hard even for Amicable Agents*

Neeldhara Misra and Aditi Sethia

Indian Institute of Technology, Gandhinagar
{neeldhara.m,aditi.sethia}@iitgn.ac.in

Abstract. We consider the problem of distributing a collection of indivisible objects among agents in a manner that satisfies some desirable notions of fairness and efficiency. We allow agents to "share" goods in order to achieve efficiency and fairness goals which may be otherwise impossible to attain. In this context, our goal is to find allocations that minimize the "amount of sharing". We follow up on recent work demonstrating that finding fair allocations with minimum sharing 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. 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 Pareto optimal and envy-free allocations is NP-complete even for instances with constant degeneracy and no sharing. We also demonstrate an alterate approach to enumerating distinct consumption graphs for allocations with a small number of sharings.

Keywords: Fair Division · Indivisible Items · NP-completeness

1 Introduction

The task of fairly distributing indivisible goods among interested agents is challenging already for the simplest possible scenario: one object valued by two or more people. We deal with the case when m *objects* are to be allocated amongst n *agents*, respecting certain notions of *fairness* and *efficiency*. Every agent attributes a *value* to each object, stating the extent to which he wants the object.

A natural and well-studied notion of fairness is *envy-freeness*, where everyone values their bundle of objects at least as much as they value others'. Some notion of efficiency are: *completeness*, which requires all items to be allocated and *fractionally Pareto Optimal*(fPO), where no agent can be made better off without making another worse off. The opening example already shows that there are instances where no allocation is simultaneously complete and envy-free(EF). This has led to several notions of "workarounds": approximate envy-freeness (e.g,

 Supported by the Indian Institute of Technology, Gandhinagar and DST-SERB.

requiring allocations to be envy-free up to the removal of one good [2,7], or any good [3], or using hidden goods [6]), subsidy (introducing money to compensate for envy [1]), donating items (this involves giving up on completeness, but to a limited extent [4]), and sharing (wherein we allow for some goods to be shared between agents [9,8]). Our focus is on the settings, where sharing goods appears to be the most reasonable of all workarounds (for instance, when high valued goods are involved), and the question of interest is to find allocations that meet our goals of fairness and efficiency with minimum sharing.

In a recent development, Sandomirskiy and Segal-Halevi [8] show that the case of identical valuation is in fact the “hardest” — they propose a notion of degeneracy(d) which captures the degree of similarity across agent valuations. We refer to the setting of low degeneracy, the ones where valuations are generally dissimilar, hence, less conflicting, as a scenario involving *amicable agents*. One of the key results in [8] is that finding allocations that are both fPO and EF is tractable for a constant number of amicable agents. In contrast, it was shown that the problem remains NP-hard for instances of high degeneracy.

Our Contributions. We investigate the complexity of finding fPO and EF allocations for amicable agents, from the perspective of the number of agents. For example, can the running time be improved to $(n + m)^{O(d)}$, which would increase the realm of tractability to scenarios with any number of agents and constant degeneracy, or more ambitiously, $O(2^{O(d)} \cdot (m + n)^{O(1)})$, which would make the problem tractable for instances with any number of agents and degeneracy logarithmic in $(n + m)$? Our main contribution here is to show that even the former goal is unlikely to be achievable: when the number of agents is unbounded, the problem of finding allocations that are fPO and EF remains *strongly* NP-complete for instances with degeneracy one, even for the specific question of allocations with no sharings. Our result also has consequences for the problem of finding EF allocations, which is weakly NP-complete by a reduction from PARTITION [8]. It turns out that the arguments in the reverse direction of our reduction do not require the allocation in question to be fPO, allowing us to obtain a stronger hardness result.

We also revisit the algorithm for finding fPO+EF allocations from [8]. The algorithm relies on enumerating certain *consumption graphs* corresponding to fPO allocations that fix the sharing structure of a potential solution. We propose an alternate method for generating the relevant consumption graphs that takes advantage of the upper bound on the number of sharings upfront. This leads to a slightly different bound that leads to a better exponential term at the cost of a worse polynomial factor. Although the difference in the bound is not significant, we believe our approach lends additional understanding to the structure of class of graphs based on fPO allocations.

2 Preliminaries

Allocations and Sharing. We use $\mathcal{A} = \{a_1, \ldots, a_n\}$ to denote a set of agents and $\mathcal{G} = \{g_1, \ldots, g_m\}$ to denote a collection of objects. A *bundle* of objects is a vector $\mathbf{b} = (b_j)_{j\in[m]} \in [0,1]^m$, where the component b_j represents the portion of g_j in the bundle. The total amount of each object is normalized to one. An *allocation* $\mathbf{z}$ is a collection of bundles $(\mathbf{z}_i)_{i\in[n]}$, one for each agent, with the condition that all the objects are fully allocated. Note that an allocation can be identified with the matrix $\mathbf{z} := (z_{i,j})_{i\in[n],j\in[m]}$, where $z_{i,j}$ take values from [0,1].

If $z_{i,j} \neq 1$, then the object g_j is *shared* between two or more agents. We define *total number of sharings* as the number of times that an object is shared, i.e:

$$\#s^\star(\mathbf{z}) = \sum_{j\in[m]} \left(\left|\{i \in [n] : z_{i,j} > 0\}\right| - 1 \right).$$

Value and Utility. For every $i \in [n], j \in [m]$, $v_{i,j}$ denotes agent a_i's *value* for the entire object g_j. In the setting of *additive* utilities, the valuations naturally lead us to an utility function over bundles defined as $u_i(\mathbf{b}) = \sum_{j\in[m]} v_{i,j} \cdot b_j$. The matrix $\mathbf{v} = (v_{i,j})_{i\in[n],j\in[m]}$ is called the *valuation matrix.*

We recall the notion of degeneracy that was proposed in [9,8]. To this end, we say that two goods g_p, g_q are valued *similarly* by a pair of agents i, j if there exists a constant r such that $v_{i,p} \cdot v_{j,q} = v_{i,q} \cdot v_{j,p} = r$. If the valuations in question are all non-zero, then the value-ratios of goods g_p and g_q by the agents i and j equals the constant r. Now, we define the *similarity* between a pair of agents i and j as:

$$s_\mathbf{v}(i,j) = \max_{r>0} \left|\left\{k \in [m] \,:\, v_{i,k} = r \cdot v_{j,k}\right\}\right| - 1.$$

Note that the similarity of a pair of agents captures the notion of the largest number of goods that the agents value similarly when considered pairwise. This finally leads us to the notion of *degeneracy* defined as $d(\mathbf{v}) = \max_{i,j\in[n],i\neq j} s_\mathbf{v}(i,j)$. Valuations for which $d(\mathbf{v}) = 0$ are called *non-degenerate*. Also, note that if any two agents have the same valuations for all goods, then $d(\mathbf{v}) = m - 1$.

Fairness and Efficiency. An allocation $\mathbf{z} = (z_i)_{i\in[n]}$ is called *envy-free (EF)* if every agent prefers her bundle to the bundles of others. Formally, for all $i, j \in [n]$: $u_i(\mathbf{z}_i) \geqslant u_i(\mathbf{z}_j)$. An allocation $\mathbf{z}$ is Pareto-dominated by an allocation $\mathbf{y}$ if $\mathbf{y}$ gives at least the same utility to all agents and strictly more to at least one. An allocation $\mathbf{z}$ is *fractionally Pareto-optimal (fPO)* if no feasible $\mathbf{y}$ dominates it.

Computational Questions. For the fairness concept EF and an efficiency concept fPO, Minimal Sharing problem is the following. Given $(\mathcal{A}, \mathcal{G}, \mathbf{v}, t \in \mathrm{N})$ as input, the question is if there exists an allocation where the total number of sharings is at most t.

3 Hardness for Instances of Constant Degeneracy

We define here, a structured version of SATISFIABILITY problem called LINEAR NEAR-EXACT SATISFIABILITY (LNES), which is known to be NP-complete [5]. An LNES instance consists of $4p$ *core* clauses ($p \in \mathbb{N}$) and p *auxiliary* clauses: $\mathcal{C} = \{U_1, V_1, U'_1, V'_1, \cdots, U_p, V_p, U'_p, V'_p\} \cup \{C_1, \cdots, C_p\}$. The set of variables consists of p *main variables* $x_1, \ldots, x_p$ and $4p$ *shadow variables* $y_1, \ldots, y_{4p}$. Each core clause consists of two literals and $\forall\, i \in [p], U_i \cap V_i = \{x_i\}$ and $U'_i \cap V'_i = \{\bar{x}_i\}$. Each shadow variable occurs as a positive literal in an auxiliary clause and as a negative literal in a core clause. An auxiliary clause consists of four literals, each corresponding to a positive occurrence of a shadow variable. We will use u_i, v_i, u'_i, and v'_i to refer to the shadow variables in the core clauses U_i, V_i, U'_i, and V'_i, respectively. The LNES problem asks whether, given a set of such clauses, there exists an assignment τ of truth values to the variables such that *exactly one* literal in every core clause and *exactly two* literals in every auxiliary clause evaluate to TRUE under τ. The main result of this section is the following:

Theorem 1. *(EF,fPO)-*MINIMAL SHARING *is NP-hard even when restricted to inputs with bounded valuations, degeneracy one, and no sharing.*

Proof. (Sketch) We reduce from LNES. Let $\mathcal{C}$ be an instance of LNES as described above. For each main variable x_i we introduce three agents: $\{a_i, \bar{a}_i, d_i\}$, and the goods $\{g_i, \bar{g}_i, h_i\}$. We refer to d_i as the *dummy* agent and a_i and $\bar{a}_i$ as the *key* agents associated with x_i. Also, we refer to h_i as the *trigger* good and g_i and $\bar{g}_i$ as *consolation* goods. For the shadow variables u_i, v_i, u'_i, v'_i, we introduce four *shadow agents* agents: b_i, c_i, b'_i, c'_i and four *essential* goods: r_i, s_i, r'_i, s'_i. Finally, for each auxiliary clause C_j, we introduce two *backup* goods f_j^1 and f_j^2. Note that our instance consists of $7p$ agents and $9p$ goods. Thus the size of the valuation matrix is $N := 63 \cdot p^2$. We let $L = 4000 \cdot p^5$.

Let $\mathbf{w} = (w_{i,j})_{i\in[n],j\in[m]}$ denote the $(7p \times 9p)$ matrix whose entries are given by $w_{i,j} = (i-1) \cdot m + j$. Intuitively, we can think of these values as being small enough to be negligible, and we will obtain our final valuation matrix by starting from $\mathbf{w}$ and "overwriting" some entries to reflect the fact that certain goods are valued highly by certain agents. This is done to ensure that the final valuation matrix has low degeneracy, and hence, the agents are amicable. We modify the matrix $\mathbf{w}$ as described below. For $i \in [p]$:

- The dummy agent d_i has a high value L for the consolation goods g_i and $\bar{g}_i$. Also, they value the four essential goods associated with them at zero.
- The first key agent a_i has a somewhat high value $\frac{L}{3}$ for the consolation good g_i and the essential goods r_i and s_i, and a high value L for h_i.
- The second key agent $\bar{a}_i$ has a somewhat high value $\frac{L}{3}$ for the consolation good $\bar{g}_i$ and the essential goods r'_i and s'_i, and also has a high value L for the trigger good h_i.

	g_i	$\overline{g_i}$	h_i	r_i	s_i	r'_i	s'_i	f_j^1	f_j^2
a_i	L/3	$\star$	L	L/3	L/3	$\star$	$\star$	$\star$	$\star$
$\overline{a_i}$	$\star$	L/3	L	$\star$	$\star$	L/3	L/3	$\star$	$\star$
d_i	L	L	$\star$	$\star$	$\star$	$\star$	$\star$	$\star$	$\star$
b_i	$\star$	$\star$	$\star$	L	$\star$	$\star$	$\star$	L	L
c_i	$\star$	$\star$	$\star$	$\star$	L	$\star$	$\star$	$\star$	$\star$
b'_i	$\star$	$\star$	$\star$	$\star$	$\star$	L	$\star$	$\star$	$\star$
c'_i	$\star$	$\star$	$\star$	$\star$	$\star$	$\star$	L	$\star$	$\star$

Fig. 1. The overall schematic of the construction in the proof of Theorem 1. The entries $\star$ indicate small values. In this example, the literal corresponding to the agent b_i, i.e, u_i, belongs to the auxiliary clause C_j corresponding to the backup goods f_j^1 and f_j^2.

- The shadow agents have a high value L for their associated essential goods and the backup good which represents an auxiliary clause that contains the associated shadow variable. Also, b_i (respectively, c_i) values the consolation good g_i and the essential good s_i (respectively, r_i) at zero. And b'_i (respectively, c'_i) values the consolation good $\bar{g}_i$ and the essential good s'_i (respectively, r'_i) at zero.

The valuation matrix is depicted in the Figure 1. We ask if this instance admits an allocation with zero sharing. Note that the degeneracy of the valuation matrix is indeed one, contributed by the values that shadow agents have for the backup goods. We claim that if there exist an assignment for the LNES instance, then we have an EF and fPO allocation for the above fair division instance, and vice-versa. We defer to the detailed arguments in the full version of the paper. □

4 Concluding Remarks

We demonstrated the hardness of finding fPO+EF allocations even for instances with constant degeneracy, with an unbounded number of agents. We note that running times of the form $d^{O(n)} \cdot \mathrm{poly}(m, n)$ are "weakly ruled out" because of the hardness result in [8]. However, all the hardness results combined so far do not rule out the possibility of an algorithm with a running time of $c^{O(d+n)} \cdot m^{O(1)}$, which would imply strongly polynomial running times for instances where $(d+n)$ is bounded by $O(\log m)$. One framework to rule out such a possibility would be parameterized complexity, where one might attempt demonstrating W-hardness in the combined parameter (n, d). On a related note, we show that instances that have bounded degeneracy and a bounded number of values in the valuation matrix are essentially bounded — we refer the reader to the full version of the paper for a more detailed discussion on bounded valuations as well as remarks on enumerating the consumption graphs.

References

1. Brustle, J., Dippel, J., Narayan, V.V., Suzuki, M., Vetta, A.: One dollar each eliminates envy. In: Proceedings of the 2020 ACM Conference on Economics and Computation (2019)
2. Budish, E.: The combinatorial assignment problem: approximate competitive equilibrium from equal incomes. In: Proceedings of the Behavioral and Quantitative Game Theory BQGT. p. 74:1. ACM (2010)
3. Caragiannis, I., Kurokawa, D., Moulin, H., Procaccia, A.D., Shah, N., Wang, J.: The unreasonable fairness of maximum nash welfare. ACM Trans. Economics and Comput. **7**(3), 12:1–12:32 (2019)
4. Chaudhury, B.R., Kavitha, T., Mehlhorn, K., Sgouritsa, A.: A little charity guarantees almost envy-freeness. In: Proceedings of the 2020 ACM-SIAM Symposium on Discrete Algorithms, SODA 2020. pp. 2658–2672. SIAM (2020)
5. Dayal, P., Misra, N.: Deleting to structured trees. In: International Computing and Combinatorics Conference. pp. 128–139. Springer (2019)
6. Hosseini, H., Sikdar, S., Vaish, R., Wang, J., Xia, L.: Fair division through information withholding. In: Proceedings of the Thirty-Fourth AAAI Conference on Artificial Intelligence (AAAI) (2020)
7. Lipton, R.J., Markakis, E., Mossel, E., Saberi, A.: On approximately fair allocations of indivisible goods. In: Proceedings of the 5th ACM Conference on Electronic Commerce (EC). pp. 125–131. ACM (2004)
8. Sandomirskiy, F., Segal-Halevi, E.: Fair division with minimal sharing. CoRR **abs/1908.01669** (2019), http://arxiv.org/abs/1908.01669
9. Segal-Halevi, E.: Fair division with bounded sharing. CoRR **abs/1912.00459** (2019), http://arxiv.org/abs/1912.00459