Fair Division of Indivisible Goods Among Strategic Agents

Siddharth Barman
Indian Institute of Science
Bangalore, India
barman@iisc.ac.in

Ganesh Ghalme
Indian Institute of Science
Bangalore, India
ganeshg@iisc.ac.in

Shweta Jain
Indian Institute of Technology
Bhubaneswar, India
shwetajains20@gmail.com

Pooja Kulkarni
University of Illinois
Urbana-Champaign, USA
poojark2@illinois.edu

Shivika Narang
Indian Institute of Science
Bangalore, India
shivika@iisc.ac.in

ABSTRACT

We study fair division of indivisible goods among strategic agents in a single-parameter environment. This work specifically considers fairness in terms of *envy freeness up to one good* (EF1) and *maximin share guarantee* (MMS). We show that (in a single-parameter environment) the problem of maximizing welfare, subject to the constraint that the allocation of the indivisible goods is EF1, admits a polynomial-time, 1/2-approximate, truthful auction. Under MMS setup, we develop a truthful auction which efficiently finds an allocation wherein each agent gets a bundle of value at least $(1/2 - \varepsilon)$ times her maximin share and the welfare of the computed allocation is at least the optimal, here $\varepsilon > 0$ is a fixed constant. Our results for EF1 and MMS are based on establishing interesting majorization inequalities.

KEYWORDS

Fair Division; Social Welfare; Approximation Algorithms; Auctions

ACM Reference Format:
Siddharth Barman, Ganesh Ghalme, Shweta Jain, Pooja Kulkarni, and Shivika Narang. 2019. Fair Division of Indivisible Goods Among Strategic Agents. In *Proc. of the 18th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2019), Montreal, Canada, May 13–17, 2019,* IFAAMAS, 3 pages.

1 INTRODUCTION

Fairness is a fundamental consideration in many real-world allocation problems. Arguably, the two most prominent notions of fairness in this line of work are *envy freeness up to one good* (EF1) and *maximin share guarantee* (MMS). These notions have been well substantiated by the development of existential results, efficient algorithms, and implementations, such as Course Match and Spliddit [2–8, 10, 11].

We focus on settings in which indivisible goods have to be auctioned off among strategic bidders/agents in *single-parameter environments.* The valuation of each strategic agent i (over goods) decomposes into the agent's private valuation parameter, $v_i \in \mathbb{R}_+$, and a (global) *public value summarization function* $w : 2^{[m]} \mapsto \mathbb{R}_+$, here m is the set of goods. We assume that w is identical across all agents.

Our problem formulations are broadly motivated by the fact that fairness is an important concern in many such applications of single-parameter environments. For example, in ad auctions it is relevant to consider fairness both from a quality-of-service standpoint and for regulatory reasons. These formulations, by construction, provide fairness guarantees which can be independently validated i.e. even when an agent is not privy to the payments charged to (and the valuation parameters of) the other agents. Besides fairness, our objectives conform to the quintessential desiderata of algorithmic mechanism design: to develop a computationally efficient, truthful auction for maximizing social welfare.

We design fair auctions (FA) under following two notions of fairness.
(i) FA-EF1: Envy freeness up to one good (EF1) was defined by Budish [3]. An allocation is said to be EF1 iff, under it, every agent values her bundle at least as much as any other agent's bundle, up to the removal of the most valuable good from the other agent's bundle. Interestingly, an EF1 allocation always exists and can be computed efficiently, even under general, combinatorial valuations [8]. The goal here is to find an allocation of the goods which achieves as high a social welfare as possible while ensuring that no agent is envious of any other, up to the removal of a good from the other agent's bundle.
(ii) FA-MMS: Maximin share guarantee (MMS) is a threshold-based notion defined by Budish [3]. This notion deems an allocation to be fair iff every agent gets a bundle of value at least as much as an agent-specific fairness threshold called the maximin share. These shares correspond to the maximum value that an agent can guarantee for herself if she were to (hypothetically) partition the goods into n subsets and, then, from them receive the minimum valued one; here n is the total number of agents. Our goal is to develop a truthful social-welfare maximizing auction subject to the constraint that each agent receives a bundle of value at least her maximin share. As computing the maximin share is NP-hard, this paper considers a bi-criteria approximation guarantee.

Our algorithms for the EF1 and MMS formulations are completely combinatorial and can be implemented in sorting time. The approximation results for the EF1 and MMS formulations rely on proving interesting *majorization inequalities*. In particular, for EF1 we show that all EF1 partitions $\frac{1}{2}$−majorize each other. For the MMS problem, we design an efficient algorithm which finds a $\left(\frac{1-\varepsilon}{2}\right)$-approximate MMS allocation that majorizes an optimal allocation.

Proc. of the 18th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2019), N. Agmon, M. E. Taylor, E. Elkind, M. Veloso (eds.), May 13–17, 2019, Montreal, Canada.

2 PRELIMINARIES AND MAIN RESULTS

We denote an instance of the fair-auction setting $\mathcal{I}$ with n bidders, $[n] = \{1, 2, \ldots, n\}$, and m indivisible goods, $[m] = \{1, 2, \ldots, m\}$ by a tuple $\langle [m], [n], w, (v_i)_{i \in [n]} \rangle$. The private preference of each agent i is represented by a single parameter $v_i \in \mathbb{R}_+$. In addition, the weight of a subset of goods $S \subseteq [m]$, is specified through a publicly-known summarization function, $w : 2^{[m]} \mapsto \mathbb{R}_+$. The valuation of bidder $i \in [n]$ for a subset of goods S is defined to be $v_i\, w(S)$. Throughout, we will consider w to be additive, i.e., $w(S) := \sum_{g \in S} w(g)$, where $w(g)$ denotes the weight of good $g \in [m]$. Write $\Pi_n([m])$ to denote the set of n-partitions of the set $[m]$. An *allocation* $\mathcal{A} = (A_1, A_2, \ldots, A_n) \in \Pi_n([m])$ refers to an n-partition of $[m]$ in which subset A_i is assigned to agent i.

Given a fair division instance $\mathcal{I} = \langle [m], [n], w, (v_i)_i \rangle$ and an allocation $\mathcal{A}$, if for every pair of agents $i, j \in [n]$ there exists a good $g \in A_j$ such that $v_i(A_i) \geq v_i(A_j) - v_i(g)$ then the allocation $\mathcal{A}$ is said to be EF1. The maximin share, μ, is defined as $\mu := \max_{(P_1, \ldots, P_n) \in \Pi_n([m])} \min_{j \in [n]} w(P_j)$. An allocation $\mathcal{A}$ is said to be MMS iff $w(A_i) \geq \mu$ for all agents $i \in [n]$. The maximin share of an agent i is defined as $\text{MMS}_i := \max_{(P_1, \ldots, P_n) \in \Pi_n([m])} \min_{j \in [n]} v_i(P_j)$. Note that $\text{MMS}_i = v_i \mu$ and we get that an allocation is MMS iff each agent i receives a bundle of value at least $v_i \mu$. We will also consider allocations which satisfy the MMS requirement approximately: for $\alpha \in (0, 1]$, an allocation $\mathcal{A}$, which satisfies $w(A_i) \geq \alpha \mu$ for all $i \in [n]$, is said to be α-approximate MMS.

An auction (A, p) is given by an allocation rule $A : \mathbb{R}^n_+ \mapsto \Pi_n([m])$ which maps the bids, $(b_i)_{i \in [n]}$, to a partition of goods, and a payment rule, p, which specifies the payment p_i charged to agent $i \in [n]$. We rely on the foundational result of Myerson [9] which asserts that for DSIC mechanisms it suffices to construct *monotone* allocation rules. We develop monotone allocations by first computing a partition, $\{P_i\}_{i \in [n]}$, of the indivisible goods $[m]$ and then allocating the jth highest (with respect to $w(\cdot)$) bundle of the partition to the jth highest bidder. We call such allocations sorted. A sorted allocation of $\mathcal{P}$ ensures that the resulting allocation rule not only satisfies the desired approximation guarantee, but is also monotone. Therefore, via Myerson's Lemma, we obtain a DSIC mechanism. These observations imply that the underlying mechanism design problem reduces to developing bid-oblivious algorithms which find fair partitions with above-stated approximation guarantee. The notable property of our algorithms is that they find a partition $\mathcal{P}$ which provides a "universal" approximation guarantee: as long as we perform a sorted allocation of $\mathcal{P}$ the stated approximation ratio is achieved, independent of the bids per se. We first introduce β-Majorization.

Definition 2.1 (β-Majorization). A sequence $(x_i)_{i=1}^n$ is said to $\beta \in \mathbb{R}_+$ majorize another sequence $(y_i)_{i=1}^n$ iff $\sum_{i=1}^k x_{(i)} \geq \beta \sum_{i=1}^k y_{(i)}$ for all $1 \leq k \leq n-1$ and $\sum_{i=1}^n x_i = \sum_{i=1}^n y_i$. Here, $x_{(i)}$ and $y_{(i)}$ denote the ith largest element in the two sequences, respectively.

It is relevant to note that FA-EF1 is NP-hard. We show that any two EF1 allocations $\frac{1}{2}$ – *majorize* each other. In particular, social welfare of a round robin allocation [4] is atleast half that of an optimal solution of FA-EF1. The following approximation guarantee holds for any EF1 allocation.

FA-EF1

$$\max_{\substack{(S_1,\ldots,S_n)\\ \in \Pi_n([m])}} \sum_{i=1}^n v_i\, w(S_i)$$

s.t. $w(S_i) \geq w(S_j) - w(g)$ for all $i, j \in [n]$ some $g \in S_j$

FA-MMS

$$\max_{\substack{(S_1,\ldots,S_n)\\ \in \Pi_n([m])}} \sum_{i=1}^n v_i\, w(S_i)$$

s.t. $w(S_i) \geq \mu$ for all $i \in [n]$ (μ is MMS value under w)

Theorem 2.2. *There exists a polytime, DSIC mechanism that achieves an approximation ratio of* $1/2$ *for FA-EF1.*

We further complement the approximation guarantee of Theorem 2.2 by showing that it is NP-hard to obtain an m^δ-approximation for the analogous problem (of maximizing social welfare subject to EF1 constraints) in *general* single-parameter environments. In the context of FA-MMS, we consider a bi-criteria approximation guarantee and establish the following result.

Theorem 2.3. *There exists a polynomial time, DSIC mechanism which computes a* $(1/2 - \varepsilon)$*-approximate* MMS *allocation with social welfare at least as much as the optimal value of FA-MMS, here* $\varepsilon \in (0, 1)$ *is a fixed constant.*

The detailed analysis of the proposed algorithm is given in full version of the paper [1]. Here we provide a brief overview of the proposed algorithm. First group the items into three categories, i.e. large, medium and small goods, based on their weights. The medium weight goods — with weight between MMS value and $\frac{1}{2}$MMS value — are allocated as singleton bundles (atmost n) in the first step. If the first step creates n bundles add all the remaining items in the bundle with largest weight and return the partition. If not, allocate small valued goods — with weight $< \frac{1}{2}$MMS — next. In this second step add the small weight goods in a bundle until the total weight of the bundle exceeds $\frac{1}{2}$MMS value. If the number of bundles formed till now reaches n, add the remaining goods to the largest bundle and return the partition. Else, in step 3, allocate large valued bundles as singletons until n bundles are formed. Any leftover goods from step 2 and/or step 3 are added to the bundle with highest weight.

We show that the above procedure creates exactly n bundles in polytime and that the weight of each bundle is atleast $\frac{1}{2}$MMS. Further we prove that the social welfare of the resulting allocation is atleast that of optimal value of FA-MMS.

3 CONCLUSION AND FUTURE WORK

This paper develops truthful and efficient mechanisms which are (approximately) fair in terms of EF1 and MMS. Going forward, it would be interesting to address revenue maximization. Note that in the standard Bayesian framework the *virtual valuations* of the agents can be negative. Hence, it is not clear if bid-oblivious algorithms exist when the objective is to maximize expected revenue. Considering non-additive public value summarization functions (e.g., submodular) is also an interesting direction for future work.

ACKNOWLEDGEMENTS

Siddharth Barman gratefully acknowledges the support of a Ramanujan Fellowship (SERB - SB/S2/RJN-128/2015) and Pratiksha Trust Young Investigator Award.

REFERENCES

[1] Siddharth Barman, Ganesh Ghalme, Shweta Jain, Pooja Kulkarni, and Shivika Narang. 2019. Fair Division of Indivisible Goods Among Strategic Agents. *arXiv preprint arXiv:1901.09427* (2019).

[2] Siddharth Barman, Sanath Kumar Krishnamurthy, and Rohit Vaish. 2018. Greedy Algorithms for Maximizing Nash Social Welfare. In *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS'18)*. ACM, 7–13.

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

[4] Ioannis Caragiannis, David Kurokawa, Hervé Moulin, Ariel D. Procaccia, Nisarg Shah, and Junxing Wang. 2016. The Unreasonable Fairness of Maximum Nash Welfare. In *Proceedings of the 2016 ACM Conference on Economics and Computation (EC '16)*. ACM, 305–322.

[5] Richard Cole and Vasilis Gkatzelis. 2015. Approximating the Nash Social Welfare with Indivisible Items. In *Proceedings of the Forty-seventh Annual ACM Symposium on Theory of Computing (STOC '15)*. ACM, 371–380.

[6] Mohammad Ghodsi, MohammadTaghi Hajiaghayi, Masoud Seddighin, Saeed Seddighin, and Hadi Yami. 2018. Fair Allocation of Indivisible Goods: Improvement and Generalization. In *Proceedings of the 2018 ACM Conference on Economics and Computation (EC'18)*. ACM, 539–556.

[7] Jonathan Goldman and Ariel D Procaccia. 2015. Spliddit: Unleashing Fair Division Algorithms. *ACM SIGecom Exchanges* 13, 2 (2015), 41–46.

[8] 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 (EC'04)*. 125–131.

[9] Roger B. Myerson. 1981. Optimal Auction Design. *Math. Oper. Res.* 6, 1 (Feb. 1981), 58–73.

[10] Abraham Othman, Tuomas Sandholm, and Eric Budish. 2010. Finding Approximate Competitive Equilibria: Efficient and Fair Course Allocation. In *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems (AAMAS'10)*. 873–880.

[11] Ariel D. Procaccia and Junxing Wang. 2014. Fair Enough: Guaranteeing Approximate Maximin Shares. In *Proceedings of the Fifteenth ACM Conference on Economics and Computation (EC '14)*. ACM, 675–692.