Report from Dagstuhl Seminar 24401

Fair Division: Algorithms, Solution Concepts, and Applications

Evangelos Markakis[*][1]**, Ruta Mehta**[*][2]**, and Yair Zick**[*][3]

1   **Athens University of Economics and Business, GR.** `markakis@gmail.com`
2   **University of Illinois – Urbana-Champaign, US.** `rutamehta@cs.illinois.edu`
3   **University of Massachusetts Amherst, US.** `yzick@umass.edu`

---- **Abstract** ----

This report documents the program and the outcomes of Dagstuhl Seminar 24401 "Fair Division: Algorithms, Solution Concepts, and Applications". The main goal of the seminar was to bring together leading scientists in the field of fair division so as to discuss current challenges and research directions. We summarize organizational aspects of the seminar, the talk abstracts, and the problems that were discussed in the open problem sessions.

1   Executive Summary

Evangelos Markakis (Athens University of Economics and Business, GR)
Ruta Mehta (University of Illinois – Urbana-Champaign, US)
Yair Zick (University of Massachusetts Amherst, US)

Fair division concerns the allocation of resources to a set of interested entities, according to some fairness criteria. Fair division has remained an active research area over the years, attracting the attention of several scientific disciplines, such as mathematics, computer science, game theory, and political science. Especially within the last two decades, this area has gradually drawn significantly more attention, due to the emergence of novel solution concepts, algorithmic techniques, and promising applications. It has also gained increasing popularity within computer science, since many of the field's key questions are inherently algorithmic. Consequently, there has been a notable growth of the relevant literature by now, spanning numerous fascinating research directions.

The main objective of the seminar was to bring together a leading set of researchers and discuss fundamental economic and computational challenges in fair division. Consequently, the seminar focused on three main categories of research topics, grouped as follows: 1) fundamental questions on the existence and the efficient computation of solution concepts in fair division; 2) the interplay of fair division with other fields such as mechanism design and machine learning; and 3) applications of fair division.

---

* Editor / Organizer

Most of the participants were academics from computer science departments or from research centers. Some were from other disciplines such as economics and mathematics. We also had the pleasure to host 2 HLF (Heidelberg Laureate Forum) participants, that is, 2 young researchers in Computer Science, in line to Dagstuhl's cooperation with the HLF.

The seminar started on Monday with an introductory session, in which the participants shared their main research interests. This session was very well-received by the participants, and it initiated discussions directly from the start. The subsequent program included 3 tutorials of roughly one hour each, along with 22 contributed talks that were solicited from the participants and lasted around 25 minutes each. In many of these talks there were lively discussions. On most days, the program also provided free time between lunch and the afternoon talks, which was used for research and collaborative meetings.

The program of the tutorials was as follows:
1. On Monday, H. Moulin provided a tutorial on the concept of "Guarantee" in fair division.
2. On Tuesday, A. Psomas gave a turorial on the theory and practice of fair food allocation.
3. On Wednesday, we had our last tutorial by Y. Zick on the theory and practice of course allocation.

Regarding the contributed talks, a large percentage of the presentations focused on solution concepts in fair division and questions related to their existence, their complexity, and their approximability. There was a substantial interest on the EFX notion, which is currently the most popular and at the same time the most intriguing fairness criterion. For this notion, G. Christodoulou presented results on EFX allocations on graphs, a special class of instances defined using an appropriate graph theoretic representation. A. Sgouritsa also presented new results on pushing the frontier on approximate EFX allocations, showing cases where one can achieve a 2/3-approximation guarantee (an open problem for general instances). E. Markakis also presented families of special cases where one can obtain 2/3-approximation algorithms or even better guarantees. Following a different direction, N. Rathi discussed epistemic notions of EFX allocations and their existence.

Beyond the EFX criterion, there were several talks on other fairness notions. U. Bhaskar presented results on the existence of EF1 allocations and H. Akrami talked about breaking the barrier of 3/4 for approximate MMS allocations. Both of these criteria, EF1 and MMS, have been studied extensively within the last years. Furthermore, V. Bilo presented results on achieving envy-freeness via selling items. In a different direction, H. Aziz discussed the pursuit of *best of both world* fairness, where the goal is to achieve both *ex ante* and *ex post* guarantees. Another model of fair division concerns the case where the items to be allocated are chores instead of goods. There were two presentations covering this topic, namely by R. Vaish, whose talk focused on a model of interval scheduling for chores, and by A. Igarashi, who also presented an actual application for sharing household chores, implemented in Japan. Yet another model was presented by A. Hollender, who discussed open problems for the continuous setting in fair division, referred to as cake cutting. In terms of applying all these notions in relevant areas, S. Roy and N. Benabbou talked about fairness as well as group-fairness in house allocation problems. Finally, E. Elkind talked about incorporating fairness concepts into temporal voting.

Another focus area was the study of incentives for strategic agents and mechanism design aspects. The seminar had 2 sessions on these topics. A. Eden presented results on the recently introduced paradigm of learning-augmented mechanism design for truthful fair division mechanisms via predictions. G. Shahkarami also talked about mechanism design with predictions for facility location problems. G. Amanatidis focused on an equilibrium analysis of the famous Round-Robin algorithm with strategic agents. In yet another direction,

S. Bouveret presented results for strategyproof and fair algorithms based on picking sequences. Finally, T. Walsh talked about the derandomization of mechanisms that retain good normative properties regarding the effects of strategic behavior.

Other directions that were covered by the talks included problems that are motivated by the further development of machine learning and AI. N. Maudet discussed a model of explainable fair division, regarding locally envy-free allocations and R. Mehta talked about fairness and incentives in federated learning.

Finally, J. Lang gave a presentation towards achieving a Sandewallian taxonomy for fair division papers. Essentially this amounts to classifying the works in this field, based on a variety of features, such as the topic or the flavor of the obtained results.

Overall, the seminar was a big success. We believe it will stimulate new and very fruitful collaborations in the near future. We also received very enthusiastic feedback from many participants, which is also reflected in the survey conducted by Dagstuhl.

2   Table of Contents

Open problems

3   Overview of Talks

3.1   Breaking the 3/4 Barrier for Approximate Maximin Share

Hannaneh Akrami (MPI für Informatik – Saarbrücken, DE)

We study the fundamental problem of fairly allocating a set of indivisible goods among n agents with additive valuations using the desirable fairness notion of maximin share (MMS). MMS is the most popular share-based notion, in which an agent finds an allocation fair to her if she receives goods worth at least her MMS value. An allocation is called MMS if all agents receive at least their MMS value. Since MMS allocations need not exist when $n > 2$, a series of works showed the existence of approximate MMS allocations with the current best factor of $3/4 + O(1/n)$. However, a simple example in [DFL82, BEF21, AGST23] showed the limitations of existing approaches and proved that they cannot improve this factor to $3/4 + \Omega(1)$. In this paper, we bypass these barriers to show the existence of (3/4+3/3836)-MMS allocations by developing new reduction rules and analysis techniques.

3.2   Allocating Indivisible Goods to Strategic Agents

Georgios Amanatidis (University of Essex – Colchester, GB)

We consider the problem of fairly allocating a set of indivisible goods to a set of strategic agents. We show that, for agents with submodular valuations, the celebrated Round-Robin ordinal algorithm always has (approximate) PNE for every instance and that these induce allocations with fairness guarantees. Further, when Round-Robin is implemented as a decentralized sequential protocol instead, by following simple greedy policies, agents have solid approximation guarantees as well as approximate EF1-type guarantees.

3.3   Best of Both Worlds Fairness

Haris Aziz (UNSW – Sydney, AU)

In this talk, I discuss the best of both worlds fairness paradigm where the goal is to design desirable randomized algorithms that achieve axiomatic properties such as fairness/stability in both in an ex-ante sense and an ex-post sense.

References

1 Aziz, H.; Freeman, R.; Shah, N.; Vaish, R., *Best of Both Worlds: Ex Ante and Ex Post Fairness in Resource Allocation.* Operations Research, 72(4), pp. 1674-1688, 2024.

2 Aziz, H.; Ganguly, A.; and Micha, E., *Best of Both Worlds Fairness under Entitlements.* In Proceedings of the 22nd International Conference on Autonomous Agents and Multiagent Systems, AAMAS 2023, pp. 941–948, London, UK, 29 May 2023 – 2 June 2023.

3 Aziz, H.; Lu, X.; Suzuki, M.; Vollen, J.; and Walsh, T., *Best-of-Both-Worlds Fairness in Committee Voting.* In Proceedings of the 19th Conference on Web and Internet Economics, WINE 2023. Extended version available at `https://arxiv.org/abs/2303.03642`.

3.4 Group-fairness in House Allocation Problems

Nawal Benabbou (Sorbonne University – Paris, FR)

We address important extensions to the house allocation problem: The agents are partitioned into disjoint groups (based on ethnicities or socio-economic status) and we want the allocation to respect some notion of diversity and/or fairness with respect to these groups. We study two specific incarnations of this general problem. First, we address a constrained optimization problem, inspired by diversity quotas in some real-world allocation problems, where the items are also partitioned into blocks and there is an upper bound on the number of items from each block that can be assigned to agents in each group. We theoretically analyze the price of diversity, the price of stability, the price of anarchy and report experiments based on real-world data sets. Next, instead of imposing hard constraints, we cast the problem as a variant of fair allocation of indivisible goods – we treat each group of agents as a single entity receiving a bundle of items whose valuation is the maximum total utility of matching agents in that group to items in that bundle; we present algorithms that achieve a standard relaxation of envy-freeness in conjunction with specific efficiency criteria, and we propose a new fairness notion better suited to house allocation problems with disjoint groups (using an axiomatic approach).

3.5 EF1 with Nonmonotone Valuations

Umang Bhaskar (TIFR Mumbai, IN)

For the fair division of items among interested agents, envyfreeness is possibly the most favoured and widely studied formalisation of fairness. For indivisible items, envy-free allocations may not exist in trivial cases, and hence research and practice focus on relaxations, particularly envy-freeness up to one item (EF1), which allows the removal of an item to resolve envy. A significant reason for the popularity of EF1 allocations in theory and practice is its simple fact of existence for many valuation classes. In this talk, we present results on the existence of EF1 for agents with nonmonotone valuations, highlighting both results and challenges.

References
1 Umang Bhaskar, A. R. Sricharan, Rohit Vaish. *On Approximate Envy-Freeness for Indivisible Chores and Mixed Resources.* In Proceedings of APPROX 2021, pp. 1-23, Seattle, Washington, USA, August 16-18, 2021.
2 Umang Bhaskar, Gunjan Kumar, Yeshwant Pandit, Rakshitha. *EF1 Allocations for Identical Trilean and Separable Single-Peaked Valuations.* Currently under submission, 2024.

3.6 Achieving Envy-Freeness through Items Sale

Vittorio Bilò (University of Salento – Lecce, IT)

We consider a fair division setting of allocating indivisible items to a set of agents. In order to cope with the well-known impossibility results related to the non-existence of envy-free allocations, we allow the option of selling some of the items so as to compensate envious agents with monetary rewards. In fact, this approach is not new in practice, as it is applied in some countries in inheritance or divorce cases. A drawback of this approach is that it may create a value loss, since the market value derived by selling an item can be less than the value perceived by the agents. Therefore, given the market values of all items, a natural goal is to identify which items to sell so as to arrive at an envy-free allocation, while at the same time maximizing the overall social welfare. Our work is focused on the algorithmic study of this problem, and we provide both positive and negative results on its approximability. When the agents have a commonly accepted value for each item, our results show a sharp separation between the cases of two or more agents. In particular, we establish a PTAS for two agents, and we complement this with a hardness result, that for three or more agents, the best approximation guarantee is provided by essentially selling all items. This hardness barrier, however, is relieved when the number of distinct item values is constant, as we provide an efficient algorithm for any number of agents. We also explore the generalization to heterogeneous valuations, where the hardness result continues to hold, and where we provide positive results for certain special cases.

3.7 Of Strategyproof and Fair Mechanisms for Multiagent Resource Allocation

Sylvain Bouveret (University of Grenoble, FR)

When allocating indivisible items to agents, it is known that the only strategyproof mechanisms that satisfy a set of rather mild conditions are constrained serial dictatorships: given a fixed order over agents, at each step the designated agent chooses a given number of items

(depending on her position in the sequence). With these rules, agents who come earlier in the sequence have a larger choice of items. However, this advantage can be compensated by a higher number of items received by those who come later. How to balance priority in the sequence and number of items received is a nontrivial question. We use a previous model, parameterized by a mapping from ranks to scores, a social welfare functional, and a distribution over preference profiles. For several meaningful choices of parameters, we show that the optimal sequence can be computed exactly in polynomial time or approximated using sampling. Our results hold for several probabilistic models on preference profiles with a particular emphasis on the Plackett-Luce model. We conclude with experimental results illustrating how the optimal sequence is impacted by various parameters of the problem.

3.8 EFX Allocations on Graphs

Giorgos Christodoulou (Aristotle University of Thessaloniki, GR & Archimedes/RC Athena – Marousi, GR)

We study envy freeness up to any good (EFX) in settings where valuations can be represented via a graph of arbitrary size where vertices correspond to agents and edges to items. An item (edge) has zero marginal value to all agents (vertices) not incident to the edge. Each vertex may have an arbitrary monotone valuation on the set of incident edges. We first consider allocations that correspond to orientations of the edges, where we show that EFX does not always exist, and furthermore that it is NP-complete to decide whether an EFX orientation exists. Our main result is that (EFX) allocations exist for this setting. This is one of the few cases where EFX allocations are known to exist for more than 3 agents.

3.9 Plant and Steal: Truthful Fair Allocations via Predictions

Alon Eden (The Hebrew University of Jerusalem, IL)

Amanatidis, Birmpas, Christodoulou and Markakis showed no truthful mechanism exists with a better approximation to the maximin share (MMS) than $\lfloor m/2 \rfloor$. I describe how to use predictions to devise truthful mechanisms that obtain improved approximation guarantees when the prediction is accurate (consistency), and have near-optimal fallback guarantees when the prediction is off (robustness).

I later discuss with the audience if my mechanism is useful (it is not), and which incentive properties should we go for instead.

3.10 Fairness in Temporal Voting

Edith Elkind (University of Oxford, GB)

We consider the problem of assigning projects to time slots, in the setting where agents have approval preferences over projects that vary over time, and each project can be implemented at most once. While optimising the utilitarian welfare in this setting is easy, optimising egalitarian welfare turns out to be computationally hard, even under strong constraints on agents' preferences. However, we describe a randomised algorithm for this problem that runs in polynomial time if the number of agents is bounded by a constant; our proof uses algebraic techniques and extends to other concepts of welfare and fairness.

3.11 Open Problems in the Complexity of Cake Cutting

Alexandros Hollender (University of Oxford, GB)

In this talk, we will see some of the main open problems in fair cake cutting, namely the problem of partitioning a divisible resource in a fair manner. The focus will be on computing connected envy-free divisions of the cake and we will consider various types of valuation functions and different models of computation.

3.12 Kajibuntan: A House Chore Division App

Ayumi Igarashi (University of Tokyo, JP)

Couples often encounter the challenge of sharing house chores. This raises the fundamental question of how to divide chores. In this paper, we present a new application for a fair division of household chores. Our platform, called Kajibuntan, allows couples to specify the set of chores to be shared, their preferences over them, and the current allocation. Our tool

visualizes the current allocation and makes proposals according to their preferences based on the theory of fair division. The goal of our tool is to provide a systematic and transparent system to divide household chores and help creating harmony in the home.

3.13 Fair Division: Towards a Sandewallian Taxonomy

Jérôme Lang (CNRS – Paris, FR)

This is a manifesto for a collaborative work. The aim is to build a taxonomy (with some carefully chosen set of labels) of fair division problems and papers, similar to Sandewall's famous taxonomy of logical systems for reasoning about action and change [1]. It is based on a set of specialties concerning: item; agents; allocations; preferences; evaluation criteria; protocols; knowledge, behaviour and dynamics; nature of work in the paper. Every speciality can take a value on a feature domain. Every problem of paper can be labelled by a set of features. Benefits of such a taxonomy are: one can see in a few seconds what a paper is about; can retrieve efficiently papers sharing some features even if they don't cite each other or use different terminologies, which will make it easier for a junior research to find their way through the (now huge) bibliography in fair division. The tool is collaborative: features can be added, deleted, renamed, refined, splitted.

References
1 Erik Sandewall, Features and fluents (vol. 1): the representation of knowledge about dynamical systems. Oxford University Press, Inc., 1995.

3.14 Improved EFX Approximation Guarantees under Ordinal-based Assumptions

Evangelos Markakis (Athens University of Economics and Business, GR)

Joint work of Evangelos Markakis, Christodoulos Santorinaios
Main reference Evangelos Markakis, Christodoulos Santorinaios: "Improved EFX Approximation Guarantees under Ordinal-based Assumptions", in Proc. of the 2023 International Conference on Autonomous Agents and Multiagent Systems, AAMAS 2023, London, United Kingdom, 29 May 2023 – 2 June 2023, pp. 591–599, ACM, 2023.
URL https://doi.org/10.5555/3545946.3598689

Our work studies the fair allocation of indivisible items to a set of agents, and falls within the scope of establishing improved approximation guarantees. It is well known by now that the classic solution concepts in fair division, such as envy-freeness and proportionality, fail to exist in the presence of indivisible items. Unfortunately, the lack of existence remains unresolved even for some relaxations of envy-freeness, and most notably for the notion of EFX, which has attracted significant attention in the relevant literature. This in turn has motivated the quest for approximation algorithms, resulting in the currently best known $(\phi - 1)$-approximation guarantee by Amanatidis et al. (2020), where ϕ equals the golden ratio. So far, it has been notoriously hard to obtain any further advancements beyond this factor. Our main contribution is that we achieve better approximations, for certain special

cases, where the agents agree on their perception of some items in terms of their worth. In particular, we first provide an algorithm with a 2/3-approximation, when the agents agree on what are the top n items (but not necessarily on their exact ranking), with n being the number of agents. To do so, we also propose a general framework that can be of independent interest for obtaining further guarantees. Secondly, we establish the existence of exact EFX allocations in a different scenario, where the agents view the items as split into tiers w.r.t. their value, and they agree on which items belong to each tier. Overall, our results provide evidence that improved guarantees can still be possible by exploiting ordinal information of the valuations.

3.15   Explaining the Lack of Locally Envy-Free Allocations

Nicolas Maudet (Sorbonne University – Paris, FR)

In fair division, local envy-freeness is a desirable property which has been thoroughly studied in recent years. We study explanations which can be given to explain that no allocation of items can satisfy this criterion, in the house allocation setting where agents receive a single item. While Minimal Unsatisfiable Subsets (MUSes) are key concepts to extract explanations, they cannot be used as such: (i) they highly depend on the initial encoding of the problem; (ii) they are flat structures which fall short of capturing the dynamics of explanations; (iii) they typically come in large number and exhibit great diversity. In this work we provide two SAT encodings of the problem which allow us to extract MUS when instances are unsatisfiable. We build a dynamic graph structure which allows to follow step-by-step the explanation. Finally, we propose several criteria to select MUSes, some of them being based on the MUS structure, while others rely on this original graphical explanation structure.

3.16   Fairness and Incentives in Federated Learning

Ruta Mehta (University of Illinois – Urbana-Champaign, US)

Federated learning provides an effective paradigm to jointly optimize a model that benefits from rich distributed data while protecting data privacy. Nonetheless, the heterogeneous nature of distributed data makes it challenging to define and ensure fairness among local agents and create incentive issues. For instance, it is intuitively "unfair" for agents with high-quality data to sacrifice their performance due to other agents with low-quality data. Furthermore, on the one hand, agents benefit from the global model trained on shared data, while on the other hand, by participating in federated learning, they may also incur costs (related to privacy and communication) due to data sharing. This may give rise to incentive issues and free riding.

In this talk, I will use social choice theory and game theory to address these fairness and incentive issues. I will show how FL and SCT can inform each other, leading to new insights and avenues for exploration.

References
1 Bhaskar Ray Chaudhury, Linyi Li, Mintong Kang, Bo Li, Ruta Mehta, *Fairness in Federated Learning via Core-Stability*. In Proceedings of the Annual Conference on Neural Information Processing Systems, NeurIPS 2022, New Orleans, LA, USA, November 28 – December 9, 2022.
2 Aniket Murhekar, Zhuowen Yuan, Bhaskar Ray Chaudhury, Bo Li, Ruta Mehta, *Incentives in Federated Learning: Equilibria, Dynamics, and Mechanisms for Welfare Maximization*. In Proceedings of the Annual Conference on Neural Information Processing Systems, NeurIPS 2023, New Orleans, LA, USA, December 10 – 16, 2023.
3 Bhaskar Ray Chaudhury, Aniket Murhekar, Zhuowen Yuan, Bo Li, Ruta Mehta, Ariel D. Procaccia, *Fair Federated Learning via the Proportional Veto Core*. In Proceedings of the 41st International Conference on Machine Learning, ICML 2024, Vienna, Austria, July 21-27, 2024.

3.17 The Versatile Concept of Guarantee in Fair Division

Hervé J. Moulin (University of Glasgow, GB)

The design of welfare guarantees for the division of resources in common property is historically the first modern (i. e., mathematical) principle of fair division and still the focus of a large body of current research. This tutorial defines the concept and corresponding mathematical question in an abstract context-free model of "welfare division" before reviewing instances of the problem where the answer is easy, as well as where it is not: cake cutting, indivisible objects, sharing of a production function and probabilistic voting.

3.18 Theory and Practice of Fair Food Allocation

Alexandros Psomas (Purdue University – West Lafayette, US)

In this talk, I will discuss the ways some non-profits address food insecurity, and what fair division has offered and has to offer to this domain.

3.19   Epistemic EFX Allocations Exist for Monotone Valuations

Nidhi Rathi (MPI für Informatik – Saarbrücken, DE)

The notion of envy-freeness up to any item (EFX) is considered to be the most fascinating fairness concept in fair division of indivisible items. Unfortunately, despite significant efforts, existence of EFX allocations is a major open problem in fair division, thereby making the study of approximations and relaxations of EFX a natural line of research. Recently, Caragiannis et al. introduced a promising relaxation of EFX, called epistemic EFX (EEFX). We say an allocation to be EEFX if, for every agent, it is possible to shuffle the items in the remaining bundles so that she becomes EFX-satisfied". Caragiannis et al. [IJCAI'23] prove existence and polynomial-time computability of EEFX allocations for additive valuations. A natural question asks what happens when we consider valuations more general than additive?

We address this important open question and answer it affirmatively by establishing the existence of EEFX allocations for an arbitrary number of agents with general monotone valuations. To the best of our knowledge, besides EF1, EEFX is the only known relaxation of EFX to have such strong existential guarantees. Furthermore, we complement our existential result by proving computational and information-theoretic lower bounds. We prove that for an arbitrary number of (more than one) agents (even) with identical submodular valuations, it is PLS-hard to compute EEFX allocations and it requires exponentially-many value queries to do so.

The talk is based on a work that will appear in AAAI 2025.

3.20   Degree of Fairness in Efficient House Allocation

Sanjukta Roy (University of Leeds, GB)

The classic house allocation problem is primarily concerned with finding a matching between a set of agents and a set of houses that guarantees some notion of economic efficiency (e.g. utilitarian welfare). While recent works have shifted focus on achieving fairness (e.g. minimizing the number of envious agents), they often come with notable costs on efficiency notions such as utilitarian or egalitarian welfare. We investigate the trade-offs between these welfare measures and several natural fairness measures that rely on the number of envious agents, the total (aggregate) envy of all agents, and maximum total envy of an agent. In particular, by focusing on envy-free allocations, we first show that, should one exist, finding an envy-free allocation with maximum utilitarian or egalitarian welfare is computationally tractable. We highlight a rather stark contrast between utilitarian and egalitarian welfare by showing that finding utilitarian welfare maximizing allocations that minimize the aforementioned fairness measures can be done in polynomial time while their

egalitarian counterparts remain intractable (for the most part) even under binary valuations. We complement our theoretical findings by giving insights into the relationship between the different fairness measures and conducting empirical analysis.

3.21 Pushing the Frontier on Approximate EFX Allocations

Alkmini Sgouritsa (Athens University of Economics and Business, GR)

We study the problem of allocating a set of indivisible goods to a set of agents with additive valuation functions, aiming to achieve approximate envy-freeness up to any good.

The state-of-the-art results on the problem include that (exact) EFX allocations exist when (a) there are at most three agents, or (b) the agents' valuation functions can take at most two values, or (c) the agents' valuation functions can be represented via a graph. In this work, by relaxing the notion of EFX to 2/3-EFX, we obtain existence results for strict generalizations of the settings for which exact EFX allocations are known to exist. We show that 2/3-EFX allocations exist when (a) there are at most seven agents, (b) the agents' valuation functions can take at most three values, or (c) the agents' valuation functions can be represented via a multigraph.

From a different viewpoint, for approximate EFX, it is known that a 0.618-EFX allocation exists for any number of agents with additive valuation functions. Our results beat the barrier of 0.618 and achieve an approximation guarantee of 2/3, but we need to impose restrictions on the setting. Therefore, our results push the frontier of existence and computation of approximate EFX allocations, and provide insights into the challenges of settling the existence of exact EFX allocations.

3.22 Learning Augmented Mechanism Design

Golnoosh Shahkarami (MPI für Informatik – Saarbrücken, DE)

In the strategic facility location problem, a set of agents report their locations in a metric space and the goal is to use these reports to open a new facility, minimizing an aggregate distance measure from the agents to the facility. However, agents are strategic and may misreport their locations to influence the facility's placement in their favor. The aim is to design truthful mechanisms, ensuring agents cannot gain by misreporting. This problem was recently revisited through the learning-augmented framework, aiming to move beyond worst-case analysis and design truthful mechanisms that are augmented with (machine-learned) predictions. The focus of this prior work was on mechanisms that are deterministic and

augmented with a prediction regarding the optimal facility location. In this paper, we provide a deeper understanding of this problem by exploring the power of randomization as well as the impact of different types of predictions on the performance of truthful learning-augmented mechanisms. We study both the single-dimensional and the Euclidean case and provide upper and lower bounds regarding the achievable approximation of the optimal egalitarian social cost.

3.23   Fair Interval Scheduling of Indivisible Chores

Rohit Vaish (Indian Institute of Technology – New Delhi, IN)

We will discuss the problem of fairly assigning a set of discrete tasks, or chores, among a set of agents. Each chore has a designated start and finish time, and each agent can perform at most one chore at any time. We will explore the existence and computation of "fair" (specifically, envy-free up to one chore) and "efficient" (specifically, maximal or Pareto optimal) schedules under various settings. The presentation will cover novel technical ideas, including a color-switching technique and an application of the "cycle-plus-triangles" theorem (originally conjectured by Erdös) for achieving approximate envy-freeness. We will also highlight several open problems and directions for future work.

3.24   Mechanisms that Play a Game, not Toss a Coin

Toby Walsh (UNSW – Sydney, AU)

Randomized mechanisms can have good normative properties compared to their deterministic counterparts. However, randomized mechanisms are problematic in several ways such as in their verifiability. We propose here to de-randomize such mechanisms by having agents play a game instead of tossing a coin. The game is designed so agents play randomly, and this play injects "randomness" into the mechanism. Surprisingly this de-randomization retains many of the good normative properties of the original randomized mechanism but gives a mechanism that is deterministic and easy, for instance, to audit. We consider three general purpose methods to de-randomize mechanisms, and apply these to six different domains: voting, facility location, task allocation, school choice, peer selection, and resource allocation. We propose a number of novel de-randomized mechanisms for these six domains with good normative properties (such as equilibria in which agents sincerely report preferences over the original problem). In one domain, we additionally show that a new and desirable normative property emerges as a result of de-randomization.

References

1 Toby Walsh, *Mechanisms that Play a Game, not Toss a Coin.* Proceedings of 33rd International Joint Conference on Artificial Intelligence, IJCAI 2024, pp. 3005-3013, Jeju, South Korea, August 3-9, 2024.

3.25 From Theory to Practice of Fair Allocation Mechanisms

Yair Zick (University of Massachusetts Amherst, US)

In this talk we provide a brief overview of the course allocation problem. We need to assign a set of class seats to students with subjective preferences over the classes they want to take. In addition, students' assigned classes are constrained; for example, a student may not take two classes with overlapping schedules, nor can they take a seat in the same class twice. Furthermore, students are often constrained in the number of classes they are allowed to take in a semester.

Unlike other instances of the problem of fair allocation of indivisible items, course scheduling often involves thousands of agents (students) and items (course seats). Our goal is to design scalable, computationally efficient mechanisms that compute assignments with both fairness and efficiency guarantees. We argue for the design of simple, sequential allocation protocols prove to be surprisingly effective in these settings.

We show that our approach is effective both in theory and in practice. We offer a comprehensive mathematical analysis of our algorithmic framework, based on the Yankee Swap mechanism proposed by Viswanathan and Zick (2023). In addition, we collect survey responses from over 1000 students at UMass Amherst, at both the undergraduate and graduate level. We show that the Yankee Swap framework outperforms several standard benchmarks on allocative efficiency, fairness, and computational efficiency.

References

1 Vignesh Viswanathan and Yair Zick. *Yankee Swap: a Fast and Simple Fair Allocation Mechanism for Matroid Rank Valuations.* Proceedings of the 22nd International Conference on Autonomous Agents and Multiagent Systems, AAMAS 2023, pp. 179-187, London, UK, 29 May 2023 – 2 June 2023.

4   Open problems

4.1   Fair Division of Chores with Budget Constraints

Edith Elkind (University of Oxford, GB)

This direction concerns chore division with time budgets. The specific problem of interest is to come up with a suitable definition of envy-freeness when each chore may have different time requirements for different agents (e.g., one agent can do the dishes in 20 min, whereas another agent needs 40 min). Several other related questions are also meaningful under this setting.

4.2   Course Allocation Among Faculty Members

Ayumi Igarashi (University of Tokyo, JP)

The problem of assigning courses and committee jobs among faculty members arises in most universities. The challenge lies in the fact that different people have varying levels of experience. While senior members may have already completed a significant number of tasks, junior members may have only taught a few courses. It is an interesting direction to establish fairness criteria in such situations.

4.3   Fairness in Machine Learning versus Fairness in Allocating Goods/Chores

Evangelos Markakis (Athens University of Economics and Business, GR)

In the recent years, there has been a stream of works in machine learning focusing on fairness criteria for learning algorithms. Fairness there refers to eliminating bias of the algorithms or more generally providing equal treatment, especially for algorithms that take important decisions (e.g., deciding whether to give a bank loan to a client). The relevant machine learning community has evolved rather independently from the community of this Dagstuhl Seminar, which focuses mostly on the fair allocation of goods and chores. It would be interesting to foster interactions between these two communities and explore the applicability of concepts from one field to the other.

4.4 Fairness in Federated Learning

Ruta Mehta (University of Illinois – Urbana-Champaign, US)

Federated learning is a powerful paradigm that allows multiple agents to learn a common model by pulling all of their data while not really sharing the data. Due to privacy and processing costs, agents have an incentive to avoid sharing while benefiting from others' data. Several recent works try to address this issue by designing mechanisms, with or without payments. However, most of these mechanisms are centralized, requiring knowledge of agents' data and their loss functions. Designing decentralized federated learning protocols, that work only on gradient-like information, and implementing such mechanisms is an important open question. A more general question is to design trustworthy mechanisms, that are fair, private, and robust to strategic agents, for generativeAI.

4.5 Inter-profile Axioms in the Allocation of Indivisible Goods

Dominik Peters (University Paris-Dauphine, FR)

Recent work on fair allocation of indivisible goods has focussed on proving existence guarantees. Examples of such existence results concern EF1, PO together with EF1, approximations to MMS, and, in special cases, EFX. But for applications, we do not just need existence, but we need a rule that, given valuations, decides on an allocation. To be sure, an existence result implies that a rule satisfying the relevant property exists. But it doesn't guarantee that such a rule is well-behaved. The best existence proofs come equipped with nice rules, like the Maximum Nash Welfare rule that guarantees PO and EF1. But other proofs do not; notably the algorithms powering MMS approximations or EFX in special cases are quite unnatural. I believe that it would be fruitful to focus some research towards turning known existential results into designing good rules (or proving their non-existence). Progress in this direction could be measured through inter-profile axioms, which connect the output allocations for different inputs. An example would be monotonicity properties. For example, if a new item is introduced, one could insist that no agent gets worse off (resource monotonicity). Or if a new agent joins, no existing agent should become better off (population monotonicity). Another monotonicity axiom requires that if an agent increases their valuation for an object that the agent receives, then the agent continues to receive that object. Other examples of inter-profile axioms are strategyproofness and non-bossiness.

4.6 Delegated Fair Division

Alexandros Psomas (Purdue University – West Lafayette, US)

In fair division, we often take the view of the designer who has resources to allocate to individuals who have values/costs for the resources. However, there are many natural scenarios (e.g., Feeding America → small food banks → food insecure individuals, college → departments → faculty, and so on), where the designer interacts with representatives and allocates the resources to them, and then the representatives allocate the resources to the individuals (that get value/incur costs). What should the designer do in these situations? See `https://arxiv.org/abs/2409.19087` for a recent (WINE 2024) paper that studies a specific instantiation of this question. There are many other natural models.

4.7 Fair Division in Metric Spaces

Nidhi Rathi (MPI für Informatik – Saarbrücken, DE)

Assume that the agents and items (can be goods or chores) lie on a metric space, defined by the distance function "d". Let us denote the distance between an agent i and an item j by $d(i, j)$. Now, this distance may represent the agent "i"s utility/disutility for item j. We can now define the valuation function v_i for an agent as follows: $v_i(j) = d(i, j)$ or $1/d(i, j)$ for an item j.

Given the fact that we have structured valuations, can we develop better fair and efficient algorithms? We can even start from agents and items lying on a path and find algorithms that find, say EF1+PO allocations efficiently. What approximations to EFX can be achieved?

4.8 Golden Ticket Mechanisms

Toby Walsh (UNSW – Sydney, AU)

In 2017, the Villum Foundation committed to spend 15% of its annual funds on high risk projects decided by an anonymous process in which every reviewer had one "golden ticket" which guaranteed funding for a project regardless of the other reviews. The Volkswagen Foundation in Hanover, Germany, ran a similar scheme from 2013 to 2021 in which each reviewer was allowed to use a "golden ticket" once per grant-review meeting.

Multiwinner voting and participatory budgeting are two social choice problems where we might consider golden ticket mechanisms. For instance, the golden ticket approval mechanism first selects any candidates given a golden ticket, and then selects the remaining winners (up to the quota) with approval voting. The golden ticket approval mechanism is committee monotone (a candidate continues to win if we increase the number of winners), and candidate monotone (a candidate continues to win if some voter adds the candidate to their approval

set). Unlike normal multiwinner approval voting, it is not support monotone (adding a new voter may prevent a candidate from winning since the new voter may golden ticket another candidate and push them out of the winning set).

It would be interesting to consider strategic behaviours of golden ticket mechanisms. When does a voter use their golden ticket? If a candidate is likely to be popular, then a voter might save their golden ticket for a less popular candidate. So how does a voter compute which candidate to golden ticket?

Participants

- Hannaneh Akrami
  MPI für Informatik –
  Saarbrücken, DE

- Georgios Amanatidis
  University of Essex –
  Colchester, GB

- Haris Aziz
  UNSW – Sydney, AU

- Nawal Benabbou
  Sorbonne University – Paris, FR

- Umang Bhaskar
  TIFR Mumbai, IN

- Vittorio Bilo
  University of Salento – Lecce, IT

- Georgios Birmpas
  University of Liverpool, GB

- Paula Böhm
  TU Clausthal, DE

- Sylvain Bouveret
  University of Grenoble, FR

- Giorgos Christodoulou
  Aristotle University of
  Thessaloniki, GR &
  Archimedes/RC Athena –
  Marousi, GR

- Alon Eden
  The Hebrew University of
  Jerusalem, IL

- Edith Elkind
  University of Oxford, GB

- Amos Fiat
  Tel Aviv University, IL

- Sushmita Gupta
  The Institute of Mathematical
  Sciences – Chennai, IN

- Martin Hoefer
  RWTH Aachen, DE

- Alexandros Hollender
  University of Oxford, GB

- Ayumi Igarashi
  University of Tokyo, JP

- Panagiotis Kanellopoulos
  University of Essex –
  Colchester, GB

- Elias Koutsoupias
  University of Oxford, GB

- Jérôme Lang
  CNRS – Paris, FR

- Evangelos Markakis
  Athens University of Economics
  and Business, GR

- Nicolas Maudet
  Sorbonne University – Paris, FR

- Ruta Mehta
  University of Illinois –
  Urbana-Champaign, US

- Hervé J. Moulin
  University of Glasgow, GB

- Keziah Naggita
  TTIC – Chicago, US

- Paula Navarette Diaz
  University of Massachusetts
  Amherst, US

- Dominik Peters
  University Paris-Dauphine, FR

- Alexandros Psomas
  Purdue University –
  West Lafayette, US

- Nidhi Rathi
  MPI für Informatik –
  Saarbrücken, DE

- Rebecca Reiffenhäuser
  University of Amsterdam, NL

- Sanjukta Roy
  University of Leeds, GB

- Ulrike Schmidt-Kraepelin
  TU Eindhoven, NL

- Alkmini Sgouritsa
  Athens University of Economics
  and Business, GR

- Golnoosh Shahkarami
  MPI für Informatik –
  Saarbrücken, DE

- Max Springer
  University of Maryland, US

- Adaku Uchendu
  MIT Lincoln Laboratory –
  Lexington, US

- Rohit Vaish
  Indian Institute of Technology –
  New Delhi, IN

- Giovanna Varricchio
  University of Calabria –
  Rende, IT

- Toby Walsh
  UNSW – Sydney, AU

- Yair Zick
  University of Massachusetts
  Amherst, US