Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence (IJCAI-16)

Online Fair Division Redux

Martin Aleksandrov
UNSW Australia and Data61 (formerly NICTA)
martin.aleksandrov@nicta.com.au

1 Introduction

Hunger is a major problem worldwide. Food banks around the globe combine forces with various welfare agencies towards alleviating the hunger by assisting people in need. For example, Foodbank Australia cooperates with local charities in order to effectively allocate food as it is donated. In 2014, nearly 10% of these relief organizations could not meet the demand and thus left around $24{,}000$ children with no breakfast in their schools; see [Byrne and Anderson, 2014]. Can we improve the food allocation? Further, the Foodbanking network in Canada has a long-standing tradition in handling customer demands, but in the last year 60% of their sponsorship covered the delivery of the food; see [Carter, 2014]. Can we reduce the transportation costs implied by the food allocation? Finally, the Meal Gap in New York reached 250 millions in 2014; see [Agi, 2015]. How do we allocate food in cities that "never sleep" and in which there are high time and spatial dynamics? Evidently, a food bank needs an allocation mechanism that takes all these features into account. Such a mechanism should be able to (1) allocate resources *online*, (2) be robust to *stochastic* changes in the allocation preferences and (3) inform *dispatching* solutions. I address exactly such complex real-world features in here.

2 Related Work and Research Plan

A greatly investigated topic in resource allocation is offline fair division. Since [Steinhaus, 1948], various mechanisms have been developed that allocate goods offline; see e.g. [Brams and Taylor, 1996]. Today, however, we witness the age of high technologies that enable us to solve complex online problems efficiently. We therefore turn our attention to online fair division. For example, [Walsh, 2011] cut cake online by exploiting offline fair division procedures. Further, we cooperate with Foodbank Australia to improve the food allocation to charities. The food arrives online and is allocated immediately to the charities. [Walsh, 2015] formulated an online model for this setting in which there is a number of agents and indivisible items arrive in rounds. Each agent has some (private) utility for each of the items. As an item arrives, they then bid for the item thus revealing their valuations for it and a mechanism allocates it to one of the agents. [Walsh, 2015] proposed two such mechanisms. LIKE gives uniformly at random an item to an agent that bids positively. The allocations of LIKE can be greatly unfair as an agent can get all items. BALANCED LIKE achieves fairer allocations as it allocates uniformly at random an item to an agent with fewest items among those agents that bid positively. This model provides a simple abstraction of clearly more complex fair division problem that can greatly benefit from the use of a number of existing optimal online matching algorithms; see e.g. [Jaillet and Lu, 2014]. Based on previous work, my thesis addresses a number of more complex features of this *online fair division* problem and focuses on the development of sophisticated and novel mechanisms for it. I next list my research plan that extends the initial one in [Aleksandrov, 2015].

- model with *one item* at a time and mechanisms; see [Aleksandrov *et al.*, 2015]
- model with *unequal entitlements* and mechanisms
- model with *multiple items* at a time and mechanisms;
- model with *divisible items* and *quotas* and mechanisms; see e.g. [Walsh, 2011]
- model with *budget-constrained* agents and mechanisms; see e.g. [Goel *et al.*, 2013]
- model with *coalition formations* and mechanisms
- mixed model with *vehicle-routing constraints* and mechanisms; see e.g. [Aleksandrov *et al.*, 2013]

To analyse mechanisms, I study axioms such as *strategyproofness*, *envy-freeness*, *efficiency* among many others. In addition, I validate their competitiveness against the optimal (offline or online) mechanism using generated and real-world data; see e.g. [Dubey, 1986; Koutsoupias and Papadimitriou, 2009; Mattei and Walsh, 2013]. Moreover, I investigate complexity questions around computing outcomes, optimal strategies and manipulations; see e.g. [Aziz *et al.*, 2015; Bouveret and Lang, 2014].

3 Results

BALANCED LIKE outputs fairer allocations than LIKE, see [Aleksandrov *et al.*, 2015]. However, BALANCED LIKE is not strategyproof. Ideally, a mechanism is strategyproof and fair. In support, the BALANCED QUEUE mechanism computes a subset of agents that have fewest items among those bidding positively and a priority *queue* that orders the agents using "first-in-first-out" principle given their last items or uniformly at random if they have zero items. It then allocates the new item using this queue.

Theorem 1 *The* BALANCED QUEUE *mechanism is strategyproof and bounded envy-free ex post with 0/1 utilities.*

Proof sketch. For strategyproofness, note that to choose an agent uniformly at random is the same as to draw an ordering of the agents uniformly at random. Hence, we can draw an ordering prior round 1 and then use it resolve ties when some agents have zero items. No agent has incentive to bid 1 for items they sincerely value with 0. Each agent gets most items if they bid 1 for items they sincerely value with 1. Otherwise, they are placed later in the queue. For bounded envy-freeness, see Theorem 8 from [Aleksandrov *et al.*, 2015].□

The mechanisms above allocate a single item at each round. But, in practice, we might expect multiple items to arrive simultaneously. At each round, the BALANCED DRAFT mechanism then computes a *balanced* ordering of the agents that gives greater priority to agents with fewer items. Ties are broken uniformly. When a single item arrives per round, BALANCED DRAFT degenerates to BALANCED LIKE. Hence, it is bounded envy-free ex post, but not strategyproof. Next, suppose we give virtual budgets to our agents. We can do this prior the allocation process or at each round. At each round, we now run an auction. ADAPTIVE CLINCHING reveals the item price and an agent that values it not lower competes for it by placing a bid, subject to their budget. The auctioneer then picks uniformly at random an agent with the greatest budget and bid, and charges them the item price. ADAPTIVE CLINCHING is budget-feasible, budget-monotonic, individually-rational, efficient and strategyproof with fixed budgets, see [Goel *et al.*, 2013]. I show that it is no longer strategyproof with increasing budgets.

Theorem 2 *The* ADAPTIVE CLINCHING *mechanism is not strategyproof with increasing budgets.*

Proof. Let us consider the alphabetical division of items a and b between agents *1* and *2*. Suppose a and b cost 0.25$ and 1$. Further, suppose *1* and *2* receive 1$ prior round 1, *1* and *2* receive 1.25$ and 1$ prior round 2, *1* values both items with 1$, and *2* values a with 0.75$ and b with 1.25$. Sincerely, *1* gets a and *2* gets b. Suppose next *1* and *2* bid strategically their total remaining budgets of 2$ at round 2. Then, *1* gets b with probability $\frac{1}{2}$. This is strict improvement.□

I further look into computational questions. BALANCED QUEUE is equivalent to BALANCED LIKE when each agent gets exactly one item in each possible ex post allocation. ADAPTIVE CLINCHING is equivalent to BALANCED LIKE when all budgets, positive bids and item prices are the same. Consequently, computing expected outcomes with BALANCED QUEUE or ADAPTIVE CLINCHING is at least as hard as with BALANCED LIKE. The latter mechanism resembles the popular random priority dictatorship whose exact expected outcomes are hard; see e.g. [Sabán and Sethuraman, 2013]. However, computing ex post outcomes with any of these mechanisms is easy.

To sum up, I presented a strategyproof and bounded envy-free ex post mechanism for the model in [Walsh, 2015], a bounded envy-free ex post mechanism for the extended model with multiple items and a strategyproof repeated auction mechanism for the model with budget-constrained agents when the budgets are fixed.

4 Conclusion

I presented my postgraduate research plan and several new results that extend my previous work. Despite the novelty of my thesis, many interesting questions remain open. For instance, how do we allocate *costs* to the charities? Perhaps some of them are more profitable to the business than others.

References

[Agi, 2015] Liz Agi. The Meal Gap under the Microscope: New York City Families at the Intersection of Food and Financial Insecurity. *Food Bank For NYC*, 2015.

[Aleksandrov *et al.*, 2013] Martin Aleksandrov, Pedro Barahona, Philip Kilby, and Toby Walsh. Heuristics and policies for online pickup and delivery problems. *AAAI Workshops*, 2013.

[Aleksandrov *et al.*, 2015] M. Aleksandrov, H. Aziz, S. Gaspers, and T. Walsh. Online Fair Division: Analysing a Food Bank Problem. In *Pr. of IJCAI*, 2015.

[Aleksandrov, 2015] Martin Aleksandrov. Online fair division. In *Pr. of IJCAI*, 2015.

[Aziz *et al.*, 2015] Haris Aziz, Toby Walsh, and Lirong Xia. Possible and necessary allocations via sequential mechanisms. In *Pr. of IJCAI 2015*, pages 468–474, 2015.

[Bouveret and Lang, 2014] Sylvain Bouveret and Jérôme Lang. Manipulating picking sequences. In *Pr. of ECAI 2014 and PAIS*, pages 141–146, 2014.

[Brams and Taylor, 1996] S.J. Brams and A.D. Taylor. *Fair fivision: from cake-cutting to dispute resolution*. Cambridge University Press, 1996.

[Byrne and Anderson, 2014] M. Byrne and K. Anderson. School Breakfast Program. *2014 Evaluation Report, Edith Cowan University*, 2014.

[Carter, 2014] Cheryl Carter. Creating a Canada where no one goes hungry. *2014 Annual Report, Food Banks Canada*, 2014.

[Dubey, 1986] Pradeep Dubey. Inefficiency of nash equilibria. *Math. Oper. Res.*, 11(1), 1986.

[Goel *et al.*, 2013] Gagan Goel, Vahab Mirrokni, and Renato Paes Leme. Clinching auctions with online supply. In *Pr. of ACM-SIAM Symposium*, SODA '13. SIAM, 2013.

[Jaillet and Lu, 2014] Patrick Jaillet and Xin Lu. Online stochastic matching: New algorithms with better bounds. *Math. Oper. Res.*, 39(3), 2014.

[Koutsoupias and Papadimitriou, 2009] Elias Koutsoupias and Christos H. Papadimitriou. Worst-case equilibria. *Computer Science Review*, 3(2), 2009.

[Mattei and Walsh, 2013] Nicholas Mattei and Toby Walsh. Preflib: A library for preferences http://www.preflib.org. In *Pr. of ADT*, 2013.

[Sabán and Sethuraman, 2013] Daniela Sabán and Jay Sethuraman. The complexity of computing the random priority allocation matrix. In *Pr. of WINE 2013*, page 421, 2013.

[Steinhaus, 1948] Hugo Steinhaus. The problem of fair division. *Econ.*, 1948.

[Walsh, 2011] T. Walsh. Online cake cutting. In *Pr. of ADT*. Springer, 2011.

[Walsh, 2015] T. Walsh. Challenges in resource and cost allocation. In *Pr. of AAAI*, 2015.