Sharing is Caring: Dynamic Mechanism for Shared Resource Ownership

Doctoral Consortium

Ridi Hossain
National University of Singapore
Singapore
ridi-h@comp.nus.edu.sg

ABSTRACT

Shared ownership of computing resources has long been in the practice; here, multiple agents pool their resources together to achieve high utility and low wastefulness. Sharing incentive, non-wastefulness and strategyproofness are three of the most desirable properties for a feasible system. However, Freeman et al. [2018] showed the fact that these three properties are incompatible in a dynamic setting and thus, a trade off must be maintained. In this work, we propose a dynamic allocation mechanism which fairly allocates the shared resources among the agents, and partially satisfies the above desiderata. Our mechanism outperforms the mechanisms proposed by Freeman et al. [2018] in the single resource case in terms of social welfare both in synthetic and real-life data. We also show that in the single resource case, our mechanism allocates the resources in a way that creates a market equilibrium and thus naturally satisfies several additional properties.

KEYWORDS

Algorithmic Game Theory; Resource Allocation; Market Mechanism

ACM Reference Format:

Ridi Hossain. 2019. Sharing is Caring: Dynamic Mechanism for Shared Resource Ownership. 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

The CS department in the University of X has a happy problem: several research labs used their grant funding to purchase servers, where they run their experiments or store data. Each server offers several computing resources (CPU, RAM, cache, storage etc.) in differing amounts. The labs' resource demands vary over time; a lab may require a lot of resources at some point of time (say, before an AI conference deadline), and at others their servers are idle. It may very well be the case that at some point in time, a lab would actually require more resources than it currently has available, and would be happy to use other labs' available server time. This imbalance is naturally undesirable: servers are expensive to maintain and letting them go unused at times when they could be utilized by others is highly wasteful. A *resource pooling mechanism* is a natural solution to this problem.

Consider, for example, a scenario with three research labs, denoted 1, 2, 3, who require a single type of resource. Labs 1 and 2 own 10 units of the resource, and lab 3 owns 40. In order to complete their tasks, they each require 20 units per round; clearly, offering 20 units of the unused resources of lab 3 to labs 1 and 2 would be significantly more efficient. In order to reap the benefits of cooperation the department has to devise a satisfactory allocation mechanism; in this work, we propose and analyze such a mechanism.

1.1 Our Contribution

We consider a dynamic model where agents contribute some fixed amount of resources to a public pool at each round, which must then be allocated to each agent according to their demands. Our mechanism satisfies the *sharing incentive* (SI) property: this baseline demand requires that agents receive at least as much resources as they have available. The mechanism also satisfies the non-wastefulness property: there is no agent with unfulfilled demands when there are unused resources in the system. Our mechanism is strategyproof (SP) for myopic agents: truthful reporting is a dominant strategy each round, assuming that agents only think about their utility in the current round rather than their long-term reward; we call it *turn wise strategyproofness*. However, when agents care about their future rewards, it is entirely possible that they will be incentivized to misreport their current demands in order to achieve high future gains. This is unavoidable, as SI, SP and non-wastefulness are incompatible [8]. That said, our mechanism only incentivizes *under-reporting* of one's demands, and never over-reporting; in other words, our mechanism may reward agents for utilizing less resources in early rounds, in order to reap higher rewards when they anticipate a greater need in later rounds. Unlike prior work, our mechanism allocates resources to agents based on the amount of resources they contributed, rather than the amount of resources they own. Our mechanism ensures some notion of long-term strategyproofness by making agents' allocation at time t is dependent on their contributions up to time $t-1$; thus, agents are incentivized to release unused resources to the common pool, though they may be incentivized to release more resources at some rounds in order to reap future rewards. We also show that our mechanism allocates the shared resources among the demanding agents in a way that it creates a market equilibrium and thus inherits some other desirable fairness and efficiency guarantees like *Pareto Optimality* (PO) and an approximate notion of *Maximin share guarantee* (MMS) [2]. We empirically compare our mechanisms to the mechanisms proposed by Freeman et al. [2018]; our mechanism outperforms the current state of the art in terms of social welfare.

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.

1.2 Related Work

Computational resource allocation is a well-studied problem [3, 9, 19]; the work by Freeman et al. [2018] most closely relates to our own: like this work, [8] study allocation mechanisms for heterogeneous, dynamic demands.

Several other works study resource allocation mechanisms [6, 10–12, 16]; however, most of them consider a static setting. Ghodsi et al. [2011] propose a mechanism based on the maximin fairness policy for multiple types of resources, maximizing the minimum dominant share amongst agents. This mechanism satisfies several desiderata, but does not consider the problem in a dynamic setting. In addition, Ghodsi et al. [2011] do not consider agents who contribute to a joint pool. Parkes et al. [2015] extend the framework of Ghodsi et al. [2011], and consider a notion of personal endowment; however, like [10], they do not consider the dynamic case. Kash et al. [2014] study fair division of multiple types of resources over dynamic setting. Like [10], they also consider agent demands following *Leontief preferences*. They introduce the notion of *dynamic Pareto optimallity* (DPO) and *Dynamic envy freeness* (DEF) and show that their proposed DynamicDRF mechanism satisfies SI, DEF, DOP and SP. However, their mechanism is not appropriate for the setting where agents contribute resources to the system.

Other works compute *competitive equilibria* obtained via artificial currency, in order to arrive at fair allocations [2, 14, 15]. Artificial currency is particularly appealing in domains such as ours, where the mechanism designer has no interest in maximizing revenue, and the use of actual money is prohibited (labs normally purchase resources via grant funding, and charging others for their use is often illegal/highly unethical). While the presence of CE has been proven in almost all cases of divisible goods [1, 18], this is not the case for indivisible goods. *Competitive equilibrium from equal incomes* (CEEI) is useful in producing outcomes where all agents should be treated equally; its market equilibrium properties immediately imply that the resulting allocation is Pareto efficient and envy-free in the case of divisible goods [7]; Budish [2011] studies approximate CEEI, where agents' budgets are perturbed, and some items remain unallocated. Nisan et al. [2007] consider a more general case and prove the existence of equilibria for a small number of items and agents, a result that was subsequently improved upon by Segal-Halevi [2018].

2 PRELIMINARIES

Let $N = \{1, \ldots, n\}$ be a set of *agents* and each of them owns $r_i \in Z^+$ resources. At round t, each agent wishes to complete $task_i^t$ copies of a single task; each instance of the task requires $d_i^t \in Z^+$ resources. At round t, a mechanism takes as input the agents' reported demands per task and the number of task copies; it outputs an *allocation* vector. Let, a_i^t is the amount of resources received by agent i at round t. We assume that tasks could be partially completed. This is captured by the following utility model:

$$u_i^t(d_i^t, a_i^t, task_i^t) = \min\left\{task_i^t, \frac{a_i^t}{d_i^t}\right\} \tag{1}$$

Current algorithmic approaches [8, 16] equate an agent's contribution with the amount of resources they contribute to the server, regardless of their consumption. However, we argue that agents' contribution should be measured in terms of the resources they *actually make available to others*. If at round t, the *shareable units* of resources are e^t, at the end of each round, the contribution parameter c_i^t of each agent i is calculated as follows:

$$c_i^t = \begin{cases} \frac{r_i - a_i^t}{e^t} & \text{if } e^t > 0 \\ 0 & \text{Otherwise.} \end{cases} \tag{2}$$

At the end of each round, the cumulative contribution parameter, C_i^t is updated, i.e. $C_i^t = C_i^{t-1} + c_i^t$.

3 OUR MECHANISM

We propose a dynamic allocation mechanism which allocates shared resources among agents based on their demands and contributions. We begin by allocating each agent an amount of resources that guarantees their welfare is at least what they can achieve on their own. Next, the excess resources, e^t are distributed among the agents who have excess demands. The agents are allocated resources in proportion to their past contributions if these are positive. First, Algorithm allocates resources to agents whose C_i^{t-1} are strictly positive; if such agents exist, they are assigned a share of e^t proportional to C_i^{t-1}; once these agents' demands are fulfilled, Algorithm divides the remaining resources among the agents with $C_i^t \leq 0$. We proof that our mechanism has sharing incentive, resource non-wastefulness and turn-wise strategyproof.

3.1 Market Setting

The concept of competitive equilibrium (CE) has long been used in the literature of fairly allocating resources [2, 5, 15]; in a competitive equilibrium, we set prices to items and allocate bundles to agents so that everyone considers their bundle to be the best possible bundle within their budget. This immediately implies that competitive equilibria satisfy some highly desirable fairness and efficiency properties.

In this work, we show that our mechanism allocates resources in a way that induces a market equilibrium and thus, naturally inherits several properties of competitive market equilibrium like an approximate notion of the *maximin share* (MMS) guarantee [2], and *Pareto Optimality* (PO).

4 CONCLUSION

We propose a novel dynamic resource allocation mechanism, which is based on agents' contribution. Our notion of agent contribution ensures that only agents' actual contributions to the system are measured, as opposed to the amount of resources that they own. We show that our mechanism ensures sharing incentives, resource non-wastefulness and turn-wise strategyproofness. Our mechanism outperforms other mechanisms in the literature in terms of social welfare, on both synthetic and real-life data. We also show that the allocation of our mechanism induces a market equilibrium among the demanding agents and thus, naturally satisfies some other fairness and efficiency guarantees like Pareto optimality, and l-out-of-d maximin share. However, as our model works for both single and multiple types of resources, it would be interesting to see what guarantees does our mechanism have in the multiple types of resource case.

REFERENCES

[1] Kenneth J. Arrow and Gerard Debreu. 1954. Existence of an equilibrium for a competitive economy. *Econometrica: Journal of the Econometric Society* 22, 3 (July 1954), 265–290. https://doi.org/10.2307/1907353

[2] Moshe Babaioff, Noam Nisan, and nbal Talgam-Cohen. 2017. Competitive Equilibria with Indivisible Goods and Generic Budgets. *arXiv:1703.08150* (2017). http://arxiv.org/abs/1703.08150

[3] Eric Boutin, Jaliya Ekanayake, Wei Lin, Bing Shi, Jingren Zhou, Zhengping Qian, Ming Wu, and Lidong Zhou. 2014. Apollo: Scalable and Coordinated Scheduling for Cloud-Scale Computing. In *Proceedings of the 11th USENIX Symposium on Operating Systems Design and Implementation (OSDI)*, Vol. 14. 285–300.

[4] Eric Budish. 2011. The Combinatorial Assignment Problem: Approximate Competitive Equilibrium from Equal Incomes. *Journal of Political Economy* 119, 6 (2011), 1061–1103. https://doi.org/10.1086/664613

[5] Shahar Dobzinski, Michal Feldman, Inbal Talgam-Cohen, and Omri Weinstein. 2015. Welfare and revenue guarantees for competitive bundling equilibrium. In *International Conference on Web and Internet Economics*. Springer, 300–313. https://doi.org/10.1007/978-3-662-48995-6_22

[6] Danny Dolev, Dror G. Feitelson, Joseph Y. Halpern, Raz Kupferman, and Nathan Linial. 2012. No Justified Complaints: On Fair Sharing of Multiple Resources. In *Proceedings of the 3rd Innovations in Theoretical Computer Science Conference (ITCS)*. ACM Press, New York, NY, USA, 68–75. https://doi.org/10.1145/2090236.2090243

[7] Duncan K Foley. 1967. Resource allocation and the public sector. (1967).

[8] Rupert Freeman, Seyed Majid Zahedi, Vincent Conitzer, and Benjamin C. Lee. 2018. Dynamic Proportional Sharing: A Game-Theoretic Approach. *Proceedings of the ACM Conference on Measurement and Analysis of Computing Systems - SIGMETRICS* 2, 1 (April 2018), 3:1–3:36. http://doi.acm.org/10.1145/3179406

[9] Eric Friedman, Christos-Alexandros Psomas, and Shai Vardi. 2017. Controlled Dynamic Fair Division. In *Proceedings of the 2017 ACM Conference on Economics and Computation (EC)*. ACM Press, New York, NY, USA, 461–478. https://doi.org/10.1145/3033274.3085123

[10] Ali Ghodsi, Matei Zaharia, Benjamin Hindman, Andy Konwinski, Scott Shenker, and Ion Stoica. 2011. Dominant Resource Fairness: Fair Allocation of Multiple Resource Types. In *Proceedings of the 8th USENIX Conference on Networked Systems Design and Implementation (NSDI)*, Vol. 11. 24–24.

[11] Ali Ghodsi, Matei Zaharia, Scott Shenker, and Ion Stoica. 2013. Choosy: Max-min Fair Sharing for Datacenter Jobs with Constraints. In *Proceedings of the 8th ACM European conference on Computer systems (EuroSys)*. ACM Press, New York, NY, USA, 365–378. https://doi.org/10.1145/2465351.2465387

[12] Avital Gutman and Noam Nisan. 2012. Fair allocation without trade. In *Proceedings of the 11th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, Vol. 2. Valencia, Spain, 719–728.

[13] Ian Kash, Ariel D Procaccia, and Nisarg Shah. 2014. No agent left behind: Dynamic fair division of multiple resources. *Journal of Artificial Intelligence Research* 51 (2014), 579–603.

[14] Noam Nisan, Tim Roughgarden, Eva Tardos, and Vijay V. Vazirani. 2007. *Algorithmic game theory.*

[15] 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 Multi-Agent Systems (AAMAS)*, Vol. 1. Richland, SC, 873–880. http://dl.acm.org/citation.cfm?id=1838206.1838323

[16] David C. Parkes, Ariel D. Procaccia, and Nisarg Shah. 2015. Beyond Dominant Resource Fairness: Extensions, Limitations, and Indivisibilities. *ACM Transactions on Economics and Computation (TEAC)* 3, 1 (March 2015), 3:1–3:22. https://doi.org/10.1145/2739040

[17] Erel Segal-Halevi. 2018. Competitive Equilibrium For Almost All Incomes. In *Proceedings of the 17th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*. Richland, SC, 1267–1275. http://dl.acm.org/citation.cfm?id=3237383.3237887

[18] Erel Segal-Halevi and Balázs R. Sziklai. 2018. Monotonicity and competitive equilibrium in cake-cutting. *Economic Theory* (2018), 1–39. https://doi.org/10.1007/s00199-018-1128-6

[19] Seyed Majid Zahedi and Benjamin C. Lee. 2014. REF: Resource Elasticity Fairness with Sharing Incentives for Multiprocessors. *SIGARCH Computer Architecture News* 42, 1 (February 2014), 145–160. https://doi.org/10.1145/2654822.2541962