Non-Cooperative Multi-Robot Planning Under Shared Resources

Doctoral Consortium

Anna Gautier
University of Oxford
Oxford, United Kingdom
anna.gautier@eng.ox.ac.uk

ABSTRACT

As more and more single-use robots are introduced to private and public spaces, it will become essential to mediate the interaction between robots. In particular, we consider the problem of resource sharing in non-cooperative multi-robot systems. We discuss the motivation for different types of shared resources and share how we used auctions to address the non-cooperative multi-agent pathfinding problem. We summarize those result, which are presented in full in a separate paper. Finally, we discuss some avenues for future work, including the application of auctions to allocate multi-unit chance-constrained resources under the presence of uncertainty.

CCS CONCEPTS

• **Computing methodologies** → **Multi-agent planning**; • **Theory of computation** → *Algorithmic mechanism design.*

KEYWORDS

Multi-agent systems; Auctions

ACM Reference Format:

Anna Gautier. 2022. Non-Cooperative Multi-Robot Planning Under Shared Resources : Doctoral Consortium. In *Proc. of the 21st International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2022), Online, May 9–13, 2022*, IFAAMAS, 3 pages.

1 INTRODUCTION

As robots are deployed in ever more areas, it is inevitable that systems owned and operated by competing organisations will be required to interact in the same space. Interactions in such settings will be inherently non-cooperative: robots from different organisations will have different goals, and must be assumed to act in pursuit of their goals, with no concern for the goals of robots from other organisations. Just as humans have laws and norms to guide them through daily life, it is important that we, as a society, begin to develop rules to guide how robots interact with each other. Because robots and other automated agents will act rationally, it makes sense to design systems that are inspired by game theory. Rosenschein and Zlotkin [8] termed these systems *rules of encounter* for automated agents. Such systems have long existed for automated online agents. The Contract Net Protocol, designed by Smith [9], is arguably the first such market-based model for automated agents. Today, more sophisticated agents participate in auctions to buy advertising space [4, 6] and markets to relay information through wireless networks [7].

Proc. of the 21st International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2022), P. Faliszewski, V. Mascardi, C. Pelachaud, M.E. Taylor (eds.), May 9–13, 2022, Online.

But there are a few unique challenges that occur in robotic systems over current online non-cooperative systems. Because these systems are physical, competitive robots will always share at least one resource, space. Consider a superstore, with aisles of groceries, clothing, and electronics. The store does not have the technology to build a fleet of robots, so they contract out individual tasks. One company might be hired to deploy a robot to move around the store examining shelves to track inventory. Another company might be hired to deploy cleaning robots. These robots have misaligned incentives; while they both have to interact in the space, they only care about their own contract with the store, and so their only goal is to complete their own task. They may be penalised for failing to complete their tasks either with a built-in monetary penalty or the eventual loss of their contract. As a result, this system is *non-cooperative*. This situation is ill-suited to centralisation, because agents need autonomy over their own decision making and path planning. In other cases, robots share physical resources like electricity, human controllers, or use of physical objects. One of the main challenges that occurs in this setting is the presence of uncertainty. In order to allocate resources optimally, it is important to distribute and manage scarce resources offline and ahead of time. But often, the environment and the results of robots' actions are uncertain. For example, an indoor package delivery robot may not know if a door is open, a chair is in its path, or which room its recipient is in. All those factors may affect the WiFi networks it needs to use, the amount of times it needs guidance from a teleoperator, the time it takes to complete a delivery, or the amount of battery it uses up in its operations. So while agents may need a certain set of resources at allocation time, this may change at execution time, and it is important for any model to consider this. By combining tools from mechanism design and multi-robot planning, we can design sensible rules of interaction for non-cooperative multi-robot resource allocation problems, like the two presented above.

2 MECHANISMS FOR MULTI-AGENT PATH FINDING

The first vector of research that we have explored focused on a specific shared resource: floor space shared by mobile robots. When multiple independently operated robots interact in the same space, their mobility can be significantly hindered by congestion, particularly in terrains with bottlenecks [10]. A warehouse setting displays this clearly. If two robots create optimised long term plans that take them into the same narrow isle at the same time, they waste time proceeding all the way out of the aisle and replanning. In settings with three or more robots, it is possible to 'trap' other robots. Another setting that illustrates this point is retail stores, which often contract different companies to provided autonomous robots for

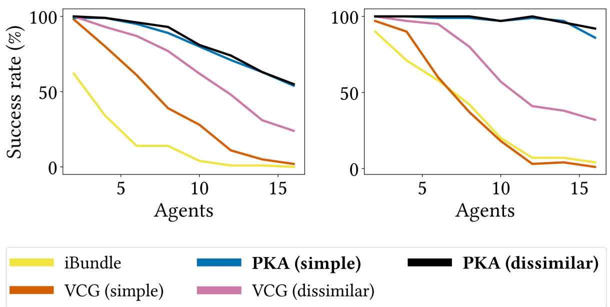

Figure 1: Success probabilities for *Dragon Age: Origins* maps lak108d (left) lak110d (right).

different tasks. Some robots, like cleaning robots, would be given tasks on short notice. Other robots, owned and operated by a different company, may have tasks which requires long term planning, like tracking inventory. If the robots do not communicate with each other, inefficiencies are inevitable. This type of non-cooperative scenario can be modelled as a multi-agent pathfinding (MAPF) problem. Agents have unique start locatations and unique goals, and the system objective is to find a series of paths through the map that are never in the same location at the same time. The cooperative analog is well studied, and the non-cooperative problem can be solved through combinatorial auctions as shown in [1]. But this approach can be prohibitively time consuming when off-the-shelf auction tools are used, particularly in large open spaces.

Instead, we developed an offline non-cooperative planning mechanism to deconflict agents paths [5]. Our Privileged Knowledge Auction (PKA) consists of a modified combinatorial Vickrey-Clarke-Groves auction. Each agents submits a set of bids, where each bid contains a timed path through the map and the respective cost of that path. Our approach limits the initial number of bids in the Vickrey-Clarke-Groves auction, then uses the privileged knowledge of the auctioneer to identify and solve path conflicts that may occur in the submitted paths. In order to maintain agent autonomy in the non-cooperative system, individual agents are provided with final say over paths. The mechanism provides a heuristic method to maximise social welfare whilst remaining computationally efficient. Prices based on the traditional Vickrey-Clarke-Groves auction deter agents from manipulating the system and lying about their path costs. We also consider the problem of single-agent bid generation and propose a similarity metric to use in dissimilar shortest path generation, with the goal of providing a diverse set of path options to the auctioneer to mitigate conflicts.

Our experiments with synthetic data outperform existing work on the non-cooperative problem. One simulation we conducted was on maps from *Dragon Age: Origins*, lak108d and lak110d. Results in Figure 1 show how our algorithm (PKA) performs against iBundle, which is the solution proposed in [1]. Our algorithm consistently found more successful solutions than iBundle, when given a timeout of 5 minutes. We also demonstrate that our dissimilar path generation mitigates conflicts in the first stage of the auction, providing a speed up. Full details on the simulations can be found in [5].

3 CURRENT WORK

The current vector of research we are exploring deals with allocating many identical resources (referred to as multi-unit) under the presence of uncertainty. There is a wealth of literature on a similar problem in the multi-agent planning community, where it is assumed that robots are willing to cooperate. In the cooperative shared resource allocation literature, it is assumed each agent has a Markov Decision Process (MDP) that governs its interactions with the world. Actions take a certain amount of resources, and uncertainty is built into the transition model. de Nijs et al. [2] provides a comprehensive review of these cooperative approaches. When agents request a certain number of resources to be pre-allocated, they are likely uncertain about the true number of resources they consume when they execute a given plan because of the stochastic environment. Say an agent is allocated a fixed number of resources. If they plans for the worst-case (i.e., assuming that uncertainty is resolved with maximum resource usage), resources may be underutilised as agents are unable to choose policies that use too many resources with a very low probability. On the other hand, if they plan only for expected resource usage, it could result in agents counting on resources without any guarantee that they are actually available. So instead, we consider allocating resources that are chance constrained, as in [3]. This allows for agents to choose less conservative plans where resources violations are unlikely, but may occur some fixed percentage of the time. Our goal is to consider the multi-robot planning problem under chance-constrained resources in a non-cooperative setting. We plan to address this problem with an auction where agents are asked to bid on resources, reporting both how much they value each resource and how likely they are to exceed a given resource amount. Agents could then be allocated resources though solving a constrained optimisation problem that ensures a chance constraint over the probability of exceeding the resource limit is met over all agents.

We would also like to consider richer understandings of uncertainty. While this planning under a chance constraint represents a good trade off between expected resource use and worst-case resource use, it only provides information on how often a resource violation can happen, without considering how bad that violation is when it does happen. Instead we hope to consider *risk-aware* systems, which consider tail cases in their entirety. One risk metric we are considering is Conditional Value at Risk, which asks "In the worst x% of cases, how bad will the outcome be on average."

Another future avenue concerns incorporating continuous or logic-based time into our existing models. Both the multi-agent pathfinding problem described in Section 2 and the resource allocation under uncertainty problem described above rely on discrete time, which requires synchronised clocks, perfect communication between agents, and uniform action duration. We hope to incorporate interval time logic into these systems to better account for real world robotics capabilities.

ACKNOWLEDGMENTS

This work was supported by a Studentship from Amazon Web Services and from the AIMS Centre for Doctoral Training.

REFERENCES

[1] Ofra Amir, Guni Sharon, and Roni Stern. 2015. Multi-Agent Pathfinding as a Combinatorial Auction. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence* (Austin, Texas) *(AAAI'15)*. AAAI Press, Palo Alto, California, 2003–2009. https://doi.org/10.5555/2886521.2886599

[2] Frits de Nijs, Erwin Walraven, Mathijs De Weerdt, and Matthijs Spaan. 2021. Constrained multiagent Markov decision processes: A taxonomy of problems and algorithms. *Journal of Artificial Intelligence Research* 70 (2021), 955–1001. https://doi.org/10.1613/jair.1.12233

[3] Frits de Nijs, Erwin Walraven, Mathijs M. de Weerdt, and Matthijs T. J. Spaan. 2017. Bounding the Probability of Resource Constraint Violations in Multi-Agent MDPs. In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence* (San Francisco, California, USA) *(AAAI'17)*. AAAI Press, 3562–3568. https://doi.org/10.5555/3298023.3298086

[4] Benjamin Edelman, Michael Ostrovsky, and Michael Schwarz. 2007. Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords. *American Economic Review* 97, 1 (2007), 242–259. https://doi.org/10.1257/aer.97.1.242

[5] Anna Gautier, Bruno Lacerda, Nick Hawes, and Michael Wooldridge. 2022. Negotiated Path Planning for Non-Cooperative Multi-Robot Systems. In *Proceedings of the 21st International Conference on Autonomous Agents & Multiagent Systems* (Online) *(AAMAS 2022)*. IFAAMAS.

[6] Cade Metz. 2017. Facebook Doesn't Make as Much Money as It Could-On Purpose. https://www.wired.com/2015/09/facebook-doesnt-make-much-money-couldon-purpose/

[7] S. Ren and M. van der Schaar. 2011. Pricing and Distributed Power Control in Wireless Relay Networks. *IEEE Transactions on Signal Processing* 59, 6 (2011), 2913–2926. https://doi.org/10.1109/TSP.2011.2121064

[8] Jeffrey S. Rosenschein and Gilad Zlotkin. 1994. *Rules of Encounter: Designing Conventions for Automated Negotiation among Computers.* MIT Press, Cambridge, MA, USA. https://doi.org/10.5555/180311

[9] Reid G Smith. 1980. The contract net protocol: High-level communication and control in a distributed problem solver. *IEEE Trans. Comput.* 12 (1980), 1104–1113. https://doi.org/10.1109/TC.1980.1675516

[10] Charlie Street, Sebastian Pütz, Manuel Mühlig, Nick Hawes, and Bruno Lacerda. 2021. Congestion-Aware Policy Synthesis for Multirobot Systems. *IEEE Transactions on Robotics* (2021), 1–19. https://doi.org/0.1109/TRO.2021.3071618