

Decision-Making under Uncertainty for Multi-Robot Systems

Bruno Lacerda, Anna Gautier, Alex Rutherford, Alex Stephens, Charlie Street and Nick Hawes

Oxford Robotics Institute, University of Oxford, United Kingdom

E-mails: bruno@robots.ox.ac.uk, alexander.rutherford@pmb.ox.ac.uk, stephens@robots.ox.ac.uk, anna.gautier@eng.ox.ac.uk, cstreet@robots.ox.ac.uk, nickh@robots.ox.ac.uk

Abstract. In this overview paper, we present the work of the Goal-Oriented Long-Lived Systems Lab on multi-robot systems. We address multi-robot systems from a decision-making under uncertainty perspective, proposing approaches that explicitly reason about the inherent uncertainty of action execution, and how such stochasticity affects multi-robot coordination. To develop effective decision-making approaches, we take a special focus on (i) temporal uncertainty, in particular of action execution; (ii) the ability to provide rich guarantees of performance, both at a local (robot) level and at a global (team) level; and (iii) scaling up to systems with real-world impact. We summarise several pieces of work and highlight how they address the challenges above, and also hint at future research directions.

Keywords: Multi-robot systems, Decision-making under uncertainty, Markov models, Formal methods, Asynchronous execution

1. Introduction

The Goal-Oriented Long-Lived Systems (GOALS) Lab is a research group based at the Oxford Robotics Institute, University of Oxford, UK. Established in 2017, GOALS focuses on *decision-making under uncertainty for autonomous systems*. In particular, we focus on long-term autonomy and task and mission planning for mobile robots which must operate for extended periods (days, weeks or months) in dynamic, uncertain environments. To create long-term autonomous behaviour, we explore the application of artificial intelligence and formal methods to robots. We focus on the use of decision-making under uncertainty and machine learning, such that the longer robots act in an environment, the better they perform. We specialise in optimising and providing formal performance guarantees given rich behaviour specifications, such as temporal logics, risk-aware planning and multi-objective scenarios.

A particularly relevant strand of research within GOALS is the field of multi-robot systems (MRS). Our MRS research focuses on exploring how robots interact with each other over time, and how modelling these interactions can improve planning. These physical interactions occur between embodied agents with limited sensory and actuation capabilities, and in environments that are often dynamic and not fully known a priori. This leads to inherently uncertain interactions, and we believe it is key to consider such uncertainty in the modelling formalisms that are employed. To do so, we typically use variants of Markov decision processes (MDPs) [26], often extending them to consider particular aspects of the problem being tackled. We also adopt techniques developed in the broader multi-agent systems (MAS) field, specialising them to the MRS setting. In this paper, we aim to provide our perspective on the main challenges to be addressed in this context, describe approaches we have developed to address them, and give an overview of our current research directions and interests.

1.1. Challenges

Teams of mobile robots evolve in *space and time*, and appropriately modelling their dynamics is key for performant multi-robot planning approaches. Furthermore, an overarching challenge in all MRS research is tackling the difficulties of scaling up to realistic applications. This leads to several interconnected challenges:

- *Modelling asynchronous behaviour.* Most modelling approaches for MAS under uncertainty, such as multi-agent MDPs (MMDPs) [4] or decentralised partially observable MDPs (Dec-POMDPs) [25], assume the agents act synchronously, in lock-step. For many MAS problems, this assumption is not overly restrictive. However, in MRS where robots are navigating between locations, enforcing synchronous actions can lead to sub-optimal behaviour and communication overhead [23]. Moreover, the *duration* of actions are also uncertain. Therefore, we require models that not only model the uncertainty over action outcome, but also explicitly consider the uncertainty over action duration.
- *Interplay between local (single-robot) and global (team) views.* When modelling MRS, one can consider two perspectives. A global, team-level, perspective related to modelling and specifying objectives that must be achieved by the team; and a local, single robot perspective related to individual models of action execution and individual robot goals. Models like MMDPs and Dec-POMDPs consider both perspectives simultaneously, but at the cost of joint state spaces composed of the Cartesian product of each robot’s internal state features (e.g. location, battery level), along with possible environmental features (e.g. presence of humans, location of objects). This leads to poor scalability. We believe that different problems require different levels of fidelity of the local and global perspectives, and scaling up requires principled approaches that make use of probabilistic guarantees of single-robot behaviour to provide guarantees at the global level. This interplay between the local and global levels is key to achieving performant solutions that can scale to realistic problems.
- *Rich behaviours.* Decision-making under uncertainty approaches typically specify goals as a reward to maximise or a cost to minimise. Whilst these are suitable in many situations, many problems require considering other, richer, objectives. We believe intended system behaviour should be specified using not only a reward or cost, but also should consider other objectives such as temporal logic specifications, safety, or risk. Furthermore, conjunctions of these objectives should often be considered simultaneously, through the use of techniques from multi-objective reasoning or constrained optimisation. Finally, they should also be considered at local or global levels, depending on the problem.

1.2. Applications

Our work typically considers cases where the goals given to the robots (either individually or as a team) involve visiting a sequence of spatial locations, sometimes within a bound on the time or energy used to complete this sequence. As such our application focus falls largely into two areas, (intra-)logistics and large-scale inspection and monitoring. Within the logistics setting our multi-robot methods optimise the flow of goods in a single, shared environment. We are working in this space with industrial partners including Accenture Labs and the Honda Research Institute Europe. Accenture Labs have used our methods to explore the automatic generation and evaluation of new warehouse layouts and human-robot team compositions. We have also taken inspiration from applications in agriculture, particularly fruit picking [7, 16], where a team of robots dynamically collect goods from a team of humans. In the area of inspection, we have been working to develop methods that could scale to multi-robot inspection of large areas such as solar farms and nuclear sites.

1.3. Other Multi-Agent Systems Research in GOALS

MAS are present in other aspects of GOALS research. In particular, we have considered shared autonomy scenarios [6, 28], which can be viewed as a MAS composed of an autonomous artificial agent and a human; and have also considered stochastic game formulations for robust and risk-averse planning [27, 29, 30], modelling the problem as a game between the robot and the environment. For this paper, we chose to focus on our research on MRS as this is where concepts and approaches from MAS are more prevalent.




Fig. 1. A mobile robot in its environment, and the corresponding map and navigation graph. Blue (bi-directional) edges represent possible navigation actions between states. Reprinted from [20].

2. Main Approaches

Our approaches consider teams of spatially distributed mobile robots that must navigate in an environment to achieve a certain goal. Depending on the setting, these might be problems where team members have individual goals, problems where a formal specification of desired behaviour is provided at the team level, or both. To abstract from the continuous dynamics of robot navigation and focus on high-level aspects of multi-robot coordination, we consider robots that evolve on a *navigation graph* $G = \langle V, E \rangle$, where vertices in V represent relevant locations in the environment and edges in E represent the ability to navigate between two locations without visiting another location. Fig. 1 depicts a robot navigating, alongside a navigation graph overlaying a metric map representation of an environment. In the following approaches, we assume there is a single navigation graph shared by all robots.

2.1. Multi-Robot Generalised Stochastic Petri Nets

We start by describing work that considers a global model of the team’s execution. This was our first work tackling asynchronous action execution under uncertain navigation times. We model a multi-robot team navigating on a topological map using a generalised stochastic Petri net (GSPN). A GSPN is formed of a set of places, immediate transitions and exponential transitions. It models the flow of tokens, with the firing of transitions moving tokens across places according to the model’s structure. In a GSPN for multi-robot navigation (MR-GSPN), we model the triggering of actions as immediate transitions, and their duration using exponential transitions. Places represent nodes and edges in the navigation graph. Crucially, the *robots are modelled as tokens*, i.e. the tokens in each place represent the number of robots in the corresponding edge or node of the environment. This allows for a *compact representation* of a global perspective of the multi-robot team that scales better than traditional multi-agent models such as MMDPs or Dec-POMDPs, which represent the state as a Cartesian product of local state features. Furthermore, the event-based semantics of MR-GSPN inherently considers asynchronous execution. Fig. 2 depicts a fragment of a MR-GSPN.

This approach relies on the anonymity and homogeneity of robots to allow for a compact representation of the team. As such, it is suitable for specifications where the goal is *global*, i.e. can be optimised without reasoning explicitly about which robot does what. To optimise specifications, we exploit the interpretation of GSPNs as *Markov automata* (MA) [10]. MA are an extension of MDPs that explicitly model immediate and exponentially timed transitions. MA can be solved for several classes of specifications, typically by reducing the optimisation problems over MA to optimisations over MDPs, which might be discrete- or continuous-timed depending on the specification [14, 15]. In our previous work, we investigated two classes




Fig. 2. Example of construction of a MR-GSPN. Left: two nodes v and v' connected by an edge in a navigation graph; Right: The fragment of the GSPN representing a team of robots navigating between v and v' . Places are represented by circles, immediate transitions by filled rectangles, exponential transitions by empty rectangles and tokens as filled circles. The depicted marking represents a state where one robot is at v and two robots are navigating from v to v' . Reprinted from [22].

of specifications. In [22], we considered the problem of maximising the total reward the system can obtain until a failure occurs. However, considering failures to be fatal, and not allowing the system to gather more reward when a failure occurs, led to policies that were often overly conservative. This, along with our interest in long-term autonomy, led to the work in [2], which focused on long-run average (LRA) reward optimisation. This allowed us to additionally quantify the cost of recovering from a failure. Empirical experiments in a wildfire monitoring scenario showed that optimising for LRA leads to better results than the more typical approach of optimising for infinite-horizon discounted reward.

2.2. Context-Aware Multi-Agent Simulation

The MR-GSPN model’s compactness allows for scaling to larger numbers of robots than the more traditional multi-agent joint state representation employed in MMDPs and Dec-POMDPs [22]. However, the model of each robot is extremely impoverished, with all robots being represented as tokens which are indistinguishable. Thus, in [34], we proposed a methodology that models MRS using a joint state, composed of the Cartesian product of each robot’s local state features and a set of global state features, but still allows the modelling of asynchronous execution in continuous time. We achieve this by extending MAs to the multi-robot case. Further, the proposed multi-robot MA (MRMA) considers *context* when computing distributions over action outcome and duration. Context can be viewed as an abstraction of the global state that summarises the *influence* [24] the environment and other robots have on a specific robot’s action execution. An example context for MRS is *congestion*, which we define as the number of robots close to the robot executing a navigation action. Contexts allow us to distinguish between different situations when executing an action, without considering the full joint state. This enables the construction of MRMA models from empirical data, and in [34] we propose data gathering policies. Finally, we introduce the context-aware multi-agent simulator (CAMAS), a discrete event simulator based on the MRMA model. We validate an MRMA obtained from data gathered in the Gazebo physics simulator by comparing statistics of CAMAS executions with Gazebo.

The MRMA is a centralised multi-agent joint model, which makes it intractable for planning. However, due to the accuracy of its underlying MRMA, CAMAS can be used to quickly evaluate policies. This enables fast prototyping and evaluation of planning solutions.

2.3. Congestion-Aware Planning

The MRMA model presented in [34] is an accurate representation of the global state of the MRS and can be used to evaluate both global and local properties of a set of policies. However, it is not suitable for planning due to its joint state representation. To plan using context-aware models, we have taken a more local perspective where each robot has its own goal and the team objective is to find a makespan-minimising solution (i.e. a solution that minimises the cost for all robots to achieve their goals) [33, 35]. The congestion-aware planner (CAP) focuses on congestion as a form of context and the impact of this context on navigation duration. To facilitate planning, we consider single-robot models augmented with congestion information obtained from a *probabilistic reservation table*, which contains information regarding the plans of other robots. These yield (single-robot) MA which are time-varying due to the congestion happening at specific times and locations. We propose a sequential planning mechanism, where each robot considers

only the plans of robots that have planned previously and ignores the behaviour of subsequent robots. These plans are modelled as continuous-time Markov chains induced by the synthesised policies over the MA and, when planning for subsequent robots, are used to compute congestion probabilities via analysis of transient probabilities [19]. Whilst suboptimal, this simple planning mechanism allows for good scalability and has been empirically proven to yield good solutions in multi-robot navigation problems [31].

2.4. Non-Cooperative Multi-Agent Path Finding

The congestion-aware planner described above deals with a variant of a multi-agent pathfinding (MAPF) problem [32] that explicitly considers continuous time and uncertainty over action durations. In MAPF, given a set of robots and a goal location for each robot, one aims to find a non-conflicting goal-reaching path for each robot. In this context, the navigation graph typically represents a grid map and a conflict occurs when two robots are in the same node or try to swap nodes. Furthermore, in its standard form, MAPF considers deterministic models.

MAPF is a canonical multi-robot coordination problem where robots must coordinate and share resources (a certain location at a certain time) and, as such, we considered it a fitting problem to establish our research on *non-cooperative* systems. Therefore, in [13], we considered a non-cooperative MAPF scenario, i.e. a scenario where we assume the agents are self-interested. In such a scenario, the global guarantees relate to mechanism design notions, specifically *incentive compatibility* and *individual rationality*. Broadly speaking, incentive compatibility means an agent cannot benefit from lying and individual rationality means each agent is better off participating in the auction than they would be otherwise.

Our approach is based on a modified combinatorial Vickrey-Clarke-Groves (VCG) auction [18]. For scalability, we limit the initial number of bids in the VCG auction, and then use the knowledge the auctioneer obtains from the bids to identify and solve path conflicts. We also consider single-agent bid generation and propose a similarity metric to use for dissimilar shortest-path generation. We show empirically that our method enables the solution of larger problems than previous work on non-cooperative MAPF [1], and that computing bids using the proposed dissimilar paths generation method increases the success likelihood of the auction.

2.5. Simultaneous Task Allocation and Planning

In all the work described above, we did not reason explicitly about the allocation of tasks to robots. However, multi-robot task allocation (MRTA) [17] is a key problem in multi-robot coordination. Many of the methods proposed for MRTA, e.g. auctioning approaches [21], decouple the task allocation from the (single-robot) planning process. However, this separation typically means that the task allocation process cannot be informed by the plans of the individual robots, which prevents allocation mechanisms from exploiting opportunities or avoiding hindrances, that are only evident once planning has been performed. Thus, in [11], we propose a framework for simultaneous task allocation and planning under uncertainty (STAPU). The approach considers an MDP model for each robot, and a mission specified as a set of co-safe linear temporal logic (LTL) tasks that must be completed, along with a safe LTL constraint that must hold throughout the MRS execution. Instead of building a full joint model of the problem, we exploit task independence assumptions to build a model which considers each robot sequentially. The obtained policies are then executed jointly. This allows the approach to scale linearly in the number of robots, in contrast with the exponential explosion associated with the joint multi-agent model.

Approach	Global View	Local View	Planning	Uncertainty	Asynchronous Execution	Task Allocation	Scalability
MR-GSPN [2, 22]	GSPN	Robots as tokens	Yes	Action duration	Yes	Implicit	10 robots
MRMA [34]	MRMA	Context-aware MA	No	Action duration and outcomes	Yes	A priori	20 robots
CAP [33, 35]	Team makespan	Context-aware MA	Yes	Action duration and outcomes	Yes	A priori	10 robots
Non-cooperative MAPF [13]	Social welfare; Individual rationality; Incentive compatibility	Grid map	Yes	None	No	A priori	15 robots
STAPU [11]	Sequential team MDP	MDP	Yes	Action outcome	No	Explicit	8 robots

Table 1

Summary of main features and assumptions of the presented methods. Scalability refers to the maximum number of robots the method was experimented on in the corresponding publication(s).

3. Discussion and Future Directions

3.1. Methods Features and Assumptions

The presented works address different aspects of planning for MRS and do so under different assumptions and model fidelity, both at the local (single-robot) and the global (team) levels. Table 3.1 summarises the features and assumptions of each method. We summarise the main points:

- The MR-GSPNs provide a global model of team behaviour which naturally considers concurrent action execution and uncertainty over action duration. Furthermore, task allocation is achieved implicitly as part of solving the global model. However, the single-robot aspect is greatly simplified, by considering robots only as tokens. This introduces limits over the kind of problem that can be tackled, as the model only counts the number of robots that are in each location, but does not consider anything else regarding their internal state. The MR-GSPN could be extended to do so, but at the cost of its compactness. In the limit, if all the internal features of each individual robot were taken into account, the MR-GSPN would become equivalent to the MRMA.
- The MRMA is an accurate and fine-grained model of both the individual robots and the team. It inherently considers uncertainty over action outcome and duration and is the highest fidelity model we have proposed. However, the joint state representation scales prohibitively and does not allow for planning for realistic problems. It can, however, be used to provide accurate statistics on the performance of policies obtained using other approaches, by using CAMAS to simulate the execution of such policies.
- The CAP takes a single-robot perspective of the MRMA discussed above. This allows for coordination to arise from planning on the single-robot models, since they include the presence of other robots, in the form of contexts of action execution. However, the local perspective only allows for planning in the presence of other robots, rather than considering more general global specifications, such as those that include task allocation.
- The non-cooperative MAPF solver considers a different high-level problem of self-interested agents, removing the assumption of fully cooperative systems that underlies our other work. It can provide global guarantees on individual rationality and is designed to prevent exploitation by the participating robots. However, it does so under assumptions of deterministic and fully synchronised robot behaviour.

- The STAPU approach explicitly considers the task allocation aspect of multi-robot coordination, and does so under rich temporal logic specifications. Independence assumptions over tasks allow us to consider the robots sequentially rather than jointly, which brings substantial scalability benefits. However, these assumptions disallow cooperative tasks. Furthermore, the approach assumes synchronous policy execution which, as mentioned previously, leads to suboptimal behaviour and communication overhead.

3.2. Research Directions

The main objective of the MRS research carried out in GOALS is to push our methods to improve along the main axis we consider: rich local models and specifications; rich global models and specifications; and scalability. These objectives are often conflicting and our main goal is to find ways to improve on all of them. To do so, we will consider what relevant aspects of each problem require precise modelling, and at which level. Furthermore, we intend to more explicitly consider the *interplay between local and global guarantees*. This general research direction instantiates into more concrete pieces of work that we are currently pursuing:

- In recent years, the field of *multi-agent reinforcement learning* (MARL) has seen substantial developments, stemming from the advances in deep reinforcement learning approaches [12]. However, its use in MRS applications is under-explored. We believe there are two main challenges related to this (i) the asynchronous action execution of teams of mobile robots is not considered in state-of-the-art MARL approaches; and (ii) there is a lack of simulators of multi-robot behaviour that can quickly and accurately provide large amounts of data for training. We believe CAMAS [34] is appropriate to tackle point (ii), and are currently researching adaptations of MARL algorithms to consider asynchronous action execution, intending to develop a MARL algorithm that enables context-aware multi-robot coordination. The use of MARL approaches will allow us to scale up to larger problems and relax observability and communication assumptions during execution, since robots will condition their actions only on (local) contexts.
- We are continuing our work on non-cooperative MRS and are extending it to consider uncertainty. Specifically, we are investigating approaches that ensure the fair allocation of a multi-unit resource to a set of robots. This problem has been considered in the cooperative setting as *resource-constrained MAS* [8], and we are extending into the non-cooperative setting by developing auction-based approaches with a price structure that ensures that robot bids are truthful. This requires a non-straightforward adaptation of traditional pricing structures because bids also include the probability of a certain resource use. The auctioning approach can also be applied to the cooperative setting, by not considering the pricing structure and simply using the auction to allocate resources without any kind of payment. We believe this work has the potential to advance the state-of-the-art on cooperative resource-constrained MAS whilst simultaneously addressing the unstudied non-cooperative aspect.
- We are continuing our work on continuous-time models of multi-robot execution and extending them to also consider *spatiotemporal models of task announcement*. We aim to bring these fine-grained local models of robot execution and task announcement to the global level, and develop *proactive* task allocation approaches that utilise the task announcement models to improve the effectiveness of the team by assigning robots to areas where tasks are likely to occur.
- Finally, we intend to extend approaches that address the online learning of environment dynamics from single-robot to multi-robot settings. In particular, we will consider our work on safe exploration and planning using Gaussian process models of a priori unknown environment features [3, 5, 9] in the multi-robot setting. We aim to develop approaches that can quickly learn *a priori* unknown environments by utilising the MRS to cover the environment efficiently.

4. Conclusion

In this paper, we described the work developed in the GOALS group related to MAS, in particular the coordination of MRS. The basis of our work relies on model-based reasoning under uncertainty, and the search for model representations of both single robots and of the overall team behaviour that enable the specification and synthesis of rich behaviour. We believe that bringing techniques and models developed in the context of more general MAS to the MRS domain has the potential to yield substantial advancements to the development and application of MRS in the real world. However, there are several challenges to be tackled and one must consider the specificities of teams of embodied agents that evolve in space and time. The MRS research in GOALS strives to develop principled techniques that achieve this in a range of scenarios.

Acknowledgements

We would like to thank our collaborators for their invaluable contributions to our work, in particular the co-authors of the papers discussed here: Fatma Faruq, David Parker, Masoumeh Mansouri, Federico Pecora, Manuel Mühlig, Carlos Azevedo, Pedro Lima, Sebastian Pütz, and Michael Wooldridge.

This work was supported by the EPSRC Programme Grant ‘From Sensing to Collaboration’ [EP/V000748/1], Honda Research Institute Europe GmbH, the AIMS Centre for Doctoral Training, and a gift from Amazon Web Services.

References

- [1] Ofra Amir, Guni Sharon, and Roni Stern. Multi-agent pathfinding as a combinatorial auction. In *Proceedings of the 29th AAAI Conference on Artificial Intelligence (AAAI)*, Palo Alto, California, 2015.
- [2] Carlos Azevedo, Bruno Lacerda, Nick Hawes, and Pedro Lima. Long-run multi-robot planning under uncertain action durations for persistent tasks. In *Proceedings of the 2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2020.
- [3] Fernando S. Barbosa, Bruno Lacerda, Paul Duckworth, Jana Tumova, and Nick Hawes. Risk-aware motion planning in partially known environments. In *Proceedings of the 2021 IEEE Conference on Decision and Control (CDC)*, 2021.
- [4] Craig Boutilier. Planning, learning and coordination in multiagent decision processes. In *Proceedings of the 6th Conference on Theoretical Aspects of Rationality and Knowledge (TARK)*, 1996.
- [5] Matthew Budd, Bruno Lacerda, Paul Duckworth, Andrew West, Barry Lennox, and Nick Hawes. Markov decision processes with unknown state feature values for safe exploration using gaussian processes. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2020.
- [6] Clarissa Costen, Marc Rigter, Bruno Lacerda, and Nick Hawes. Shared autonomy systems with stochastic operator models. In *Proceedings of the 31st International Joint Conference on Artificial Intelligence (IJCAI)*, 2022.
- [7] Gautham Das, Grzegorz Cielniak, Pal From, and Marc Hanheide. Discrete event simulations for scalability analysis of robotic in-field logistics in agriculture—a case study. In *Proceedings of the IEEE International Conference on Robotics and Automation, Workshop on Robotic Vision and Action in Agriculture (WRVAA)*, 2018.
- [8] Frits De Nijs, Erwin Walraven, Mathijs De Weerd, and Matthijs Spaan. Constrained multiagent Markov decision processes: A taxonomy of problems and algorithms. *Journal of Artificial Intelligence Research*, 70:955–1001, 2021.
- [9] Paul Duckworth, Bruno Lacerda, and Nick Hawes. Time-bounded mission planning in time-varying domains with semi-MDPs and Gaussian processes. In *Proceedings of the 2020 Conference on Robot Learning (CoRL)*, 2020.
- [10] Christian Eisentraut, Holger Hermanns, Joost-Pieter Katoen, and Lijun Zhang. A semantics for every GSPN. In *Proceedings of the 34th International Conference on Applications and Theory of Petri Nets and Concurrency (Petri Nets)*. Springer, 2013.
- [11] Fatma Faruq, Bruno Lacerda, Nick Hawes, and David Parker. Simultaneous task allocation and planning under uncertainty. In *Proceedings of the 2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Madrid, Spain, 2018.
- [12] Jakob Foerster, Ioannis Alexandros Assael, Nando De Freitas, and Shimon Whiteson. Learning to communicate with deep multi-agent reinforcement learning. In *Proceedings of Advances in Neural Information Processing Systems 29 (NeurIPS)*, 2016.

- [13] Anna Gautier, Alex Stephens, Bruno Lacerda, Nick Hawes, and Michael Wooldridge. Negotiated path planning for non-cooperative multi-robot systems. In *Proceedings of the 21st International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2022.
- [14] Dennis Guck, Hassan Hatefi, Holger Hermanns, Joost-Pieter Katoen, and Mark Timmer. Analysis of timed and long-run objectives for markov automata. *Logical Methods in Computer Science*, 10(3), 2014.
- [15] Hassan Hatefi and Holger Hermanns. Model checking algorithms for Markov automata. *Electronic Communications of the EASST*, 53, 2012.
- [16] Muhammad W Khan, Gautham P Das, Marc Hanheide, and Grzegorz Cielniak. Incorporating spatial constraints into a bayesian tracking framework for improved localisation in agricultural environments. In *Proceedings of the 2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 2440–2445. IEEE, 2020.
- [17] G Ayorkor Korsah, Anthony Stentz, and M Bernardine Dias. A comprehensive taxonomy for multi-robot task allocation. *The International Journal of Robotics Research*, 32(12):1495–1512, 2013.
- [18] Vijay Krishna. *Auction theory*. Academic press, 2009.
- [19] Marta Kwiatkowska, Gethin Norman, and David Parker. Stochastic model checking. In *Proceedings of the International School on Formal Methods for the Design of Computer, Communication and Software Systems: Performance Evaluation (SFM)*, pages 220–270. Springer, 2007.
- [20] Bruno Lacerda, David Parker, and Nick Hawes. Optimal and dynamic planning for Markov decision processes with co-safe LTL specifications. In *Proceedings of the 2014 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Chicago, IL, USA, 2014.
- [21] M.G. Lagoudakis, M. Berhault, S. Koenig, P. Keskinocak, and A.J. Kleywegt. Simple auctions with performance guarantees for multi-robot task allocation. In *Proceedings of the 2004 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 698–705, 2004.
- [22] Masoumeh Mansouri, Bruno Lacerda, Nick Hawes, and Federico Pecora. Multi-robot planning under uncertain travel times and safety constraints. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI)*, Macau, China, 2019.
- [23] João Vicente Messias, Matthijs T. J. Spaan, and Pedro U. Lima. GSMDPs for multi-robot sequential decision-making. In *Proceedings of the 27th AAAI Conference on Artificial Intelligence (AAAI)*, 2013.
- [24] Frans Oliehoek, Stefan Witwicki, and Leslie Kaelbling. Influence-based abstraction for multiagent systems. In *Proceedings of the 26th AAAI Conference on Artificial Intelligence (AAAI)*, 2012.
- [25] Frans A Oliehoek and Christopher Amato. *A concise introduction to decentralized POMDPs*. Springer, 2016.
- [26] Martin L Puterman. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. John Wiley & Sons, Inc., 1994.
- [27] Marc Rigter, Paul Duckworth, Bruno Lacerda, and Nick Hawes. Planning for risk-aversion and expected value in MDPs. In *Proceedings of the 32nd International Conference on Automated Planning and Scheduling (ICAPS)*, 2022.
- [28] Marc Rigter, Bruno Lacerda, and Nick Hawes. A framework for learning from demonstration with minimal human effort. *Robotics and Automation Letters (RA-L)*, 5(2), 2020.
- [29] Marc Rigter, Bruno Lacerda, and Nick Hawes. Minimax regret optimisation for robust planning in uncertain Markov decision processes. In *Proceedings of the 35th AAAI Conference on Artificial Intelligence (AAAI)*, 2021.
- [30] Marc Rigter, Bruno Lacerda, and Nick Hawes. Risk-averse bayes-adaptive reinforcement learning. In *Proceedings of Advances in Neural Information Processing Systems 34 (NeurIPS)*, 2021.
- [31] David Silver. Cooperative pathfinding. In *Proceedings of the first AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment (AIIDE)*, pages 117–122. AAAI Press, 2005.
- [32] Roni Stern, Nathan R. Sturtevant, Ariel Felner, Sven Koenig, Hang Ma, Thayne T. Walker, Jiaoyang Li, Dor Atzmon, Liron Cohen, T. K. Satish Kumar, Eli Boyarski, and Roman Bartak. Multi-agent pathfinding: Definitions, variants, and benchmarks. In *Proceedings of the 12th International Symposium on Combinatorial Search (SoCS)*, pages 151–158, 2019.
- [33] Charlie Street, Bruno Lacerda, Manuel Mühlig, and Nick Hawes. Multi-robot planning under uncertainty with congestion-aware models. In *Proceedings of the 19th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2020.
- [34] Charlie Street, Bruno Lacerda, Michal Staniaszek, Manuel Mühlig, and Nick Hawes. Context-aware modelling for multi-robot systems under uncertainty. In *Proceedings of the 21st International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2022.
- [35] Charlie Street, Sebastian Pütz, Manuel Mühlig, Nick Hawes, and Bruno Lacerda. Congestion-aware policy synthesis for multirobot systems. *IEEE Transactions on Robotics*, 38(1), 2022.