

# EgoPlan: A Framework for Multi-Agent Planning Using Single Agent Planners

Mark McArthur, Yashar Moshfeghi, and Michael Cashmore

University of Strathclyde, Glasgow, Scotland

{mark.mcarthur.2016,yashar.moshfeghi,michael.cashmore}@strath.ac.uk

## Abstract

Planning problems are, in general, PSPACE-complete; large problems, especially multi-agent problems with required coordination, can be intractable or impractical to solve. Factored planning and multi-agent planning both address this by separating multi-agent problems into tractable sub-problems, but there are limitations in the expressivity of existing planners and in the ability to handle tightly coupled multi-agent problems. This paper presents EGOPLAN, a framework which factors a multi-agent problem into related sub-problems which are solved by iteratively calling on a single agent planner. EGOPLAN is evaluated on a multi-robot test domain with durative actions, required coordination, and temporal constraints, comparing the performance of a temporal planner, OPTIC-CPLEX, with and without EGOPLAN. Our results show that for our test domain, using EGOPLAN allows OPTIC-CPLEX to solve problems that are twice as complex as it can solve without EGOPLAN, and to solve complex problems significantly faster.

## 1 Introduction

Many planning problems naturally model situations with multiple acting agents, i.e. trucks in a logistics domain or teams in an emergency response problem, that have an inherent factoring into smaller, connected sub-problems. Scenarios that can be solved by a single complex agent can often be tackled more cheaply and efficiently with a team of simpler heterogeneous agents (Carreno et al. 2020). Multi-agent planning techniques and factored planning techniques can take advantage of factoring to reduce the complexity of a planning problem. This allows these techniques to solve larger problems than unfactored single agent planners and to solve large problems faster.

Both multi-agent and factored planning solve factored versions of the problem to generate a valid plan for the original problem. The difference between these techniques is in the nature and purpose of the factorisation. Typically, multi-agent planners must use an imposed factorisation based on privacy or limited communication. Most multi-agent planners are designed to be suitable for one type of imposed factorisation (Torreño et al. 2018), such as privacy-preserving planners (Torreño, Onaindia, and Óscar Sapena 2014). On

the other hand, factored planners typically analyse the problem to generate an efficient factorisation, and are not restricted by any imposed factoring (Brafman and Domshlak 2006). When factorisation cannot completely decouple sub-problems the use of specialised multi-agent or factored planners is required. These planners often cannot handle the same combinations of expressive features as single agent planners, e.g. uncertainty and time (Torreño et al. 2018).

In this paper we introduce EGOPLAN<sup>1</sup>, a novel factored planning technique that enables the use of a single agent planner (for any expressive planning language) in multi-agent problems where interaction between agents is required. Our initial approach restricts agents to one of two roles: those that achieve goals (*Egobots*) or support other agents (a single *Sidekick*). The two groups iteratively converge upon a solution. Removing the restriction of these two roles is discussed in Section 7.

EGOPLAN assumes a factorisation imposed by separate agents in a real-world problem. However, EGOPLAN does not handle privacy constraints or limited communication. As such, it is somewhere between factored planners and multi-agent planners, and is best compared to factored planners in terms of performance.

Our approach is agnostic to the choice of planner, meaning the approach can tackle more expressive problems than most multi-agent or factored planners, even when sub-tasks require interaction between agents. Our implementation is based on temporal problems in the Planning Domain Definition Language (PDDL2.2) (Edelkamp and Hoffmann 2004).

This paper describes related work in factored and multi-agent planning, then defines a multi-agent problem and the EGOPLAN framework. An evaluation of EGOPLAN is given in Section 6.

## 2 Related Work

In this section, we give an overview of multi-agent and factored planning and highlight similar approaches to EGOPLAN. A more detailed survey of the state-of-the-art in multi-agent planning is provided by Torreño et al. (2018).

<sup>1</sup>An implementation of EGOPLAN and benchmark domains are available online: <https://github.com/strathclyde-artificial-intelligence/egobots-and-sidekicks>

Multi-agent planners can be distinguished by how they combine planning and coordination. Some planners plan cooperatively, with high levels of communication between agents to allow them to function like a threaded single agent planner (Torreño, Onaindia, and Óscar Sapena 2014). In other planners, each agent plans independently, and then constraint satisfaction or replanning techniques are used to combine those plans (Nissim, Brafman, and Domshlak 2010).

A key distinction between factored planners and multi-agent planners is that factored planners search for a decomposition of the problem that minimises the search space. Decomposition can be done by considering the possible orderings between *serializable subgoals* (Korf 1987; Yu et al. 2004) or causal graph analysis (Wang and Williams 2015). Other approaches depend on domain knowledge, for example Buksz et al. (2019) decompose problems based on the *locality* of subgoals. Carreno et al. (2020) build upon this to include agent capabilities. More similar to our approach, REALPLAN (Srivastava 2000) decomposes tasks based on available resources, such as independent robots.

These approaches allow the direct use of single agent planners with potentially much more expressive planning languages, but are limited by the ways in which the factored problems may interact with each other (Crosby, Rovatsos, and Petrick 2013). For example, A# proposed by Jezequel and Fabre (2012) models interactions between sub-problems by biasing each planning agent’s search towards plans which are mutually compatible with other planning agents, rather than directly modelling requests from one agent to another.

The current implementation of EGOPLAN handles temporal and numeric problems in PDDL2.2 (Edelkamp and Hoffmann 2004), which is used in our evaluation. While most factored planners are limited in expressivity, temporal factored planners do exist such as tBurton (Wang and Williams 2015) and TFPOP (Kvarnström 2011). tBurton uses a single agent planner embedded in a larger framework, similar to the EGOPLAN framework, and factors the problem based on sub-goals. The partial plans which can achieve these sub-goals are then unified using a backtracking search. More similar to EGOPLAN, TFPOP factors by acting agents, assuming that each acting agent will have long strings of actions which require no interaction with other agents. This assumption describes loosely coupled problems, and when it does not hold TFPOP is still able to solve problems but with less factoring.

Crosby et al. (2014) proposed a framework that transforms problems between multi-agent and single agent. Their method transforms a multi-agent problem with required concurrent actions into a simpler single agent problem, where concurrent actions are modelled as single actions containing predicates for each involved agent. While also using a single agent planner to solve a multi-agent problem, our approach builds a framework around the single agent planner and solves many small sub-problems whereas Crosby et al.’s method combines the sub-problems into a single problem the single agent planner can solve.

The previous work most similar to EGOPLAN is A-SHOP (Dix et al. 2003), which implemented a SHOP planner as an agent in the IMPACT multi-agent environment.

This allowed the use of SHOP’s planning capabilities to solve multi-agent problems and allowed SHOP to use arbitrary inputs from other IMPACT agents. The system could also handle multiple separate SHOP planning agents in a multi-agent team. However, there was no specific functionality for multi-agent planning, as the planners were agents that would take inputs and give outputs in a non-planner environment. The depth of interaction between planning agents is unclear.

### 3 Problem Definition

In EGOPLAN, problems are solved by a team of planning agents that can be separated into two roles: *Egobots* and *Sidekicks*. Egobots have goals to achieve and may request help from Sidekicks. Sidekicks do not have their own goals but receive requests from Egobots and try to satisfy as many of those requests as they can. EGOPLAN is a first step towards a more generally applicable system in which teams of planning agents can both send and receive requests. We define the multi-agent problem as follows.

**Definition 1 (Multi-Agent Problem).** The Multi-Agent Planning Problem  $\Pi$  is a 6-tuple  $\langle \Phi, \Psi, P, I, G, A \rangle$  where:  $\Phi$  is a set of agents. Each agent  $\varphi_i$  is associated with a set of agents to which it may make requests, denoted as  $\Phi_i \subseteq \Phi$ . These mappings are represented as the set  $\Psi : \{\Phi_i, \forall \varphi_i\}$ .  $P$  is a finite set of propositions and fluents. Each agent  $\varphi_i$  can observe only a subset of propositions and fluents  $P_i \subseteq P$ . These subsets are not disjoint.  $I \subseteq P$  is the initial state.  $G \subseteq P$  is the goal condition.  $G_i \subseteq G$  is the goal condition observable by agent  $\varphi_i$ . The goal condition observable by each agent is disjoint, so that  $G_i \cap G_j = \emptyset$  for all  $\varphi_i, \varphi_j \in \Phi$ .  $A$  is a set of durative actions. For each action, there is a single planning agent  $\varphi_i$  that is executive over that action, and the action contains only one acting agent.  $A_i$  is the set of actions for which  $\varphi_i$  is the executive.

Factoring the multi-agent problem results in a set of factored problems,  $\Pi_i$  for each agent  $\varphi_i$ , defined as follows.

**Definition 2 (Factored Agent Problem).** The problem  $\Pi_i$  for agent  $\varphi_i$  is a 5-tuple  $\Pi_i = \langle P_i, I_i, G_i, \hat{A}, R_i \rangle$  where:  $P_i$  and  $G_i$  are the propositions, fluents, and goal conditions known to the agent as defined above.  $I_i = P_i \cap I$  is the initial state known to agent  $\varphi_i$ .  $\hat{A} : A_i \cup \bigcup_{\varphi_j \in \Phi_i} A_j$  is the set of durative actions over which agent  $\varphi_i$  is executive, plus actions belonging to agents to which  $\varphi_i$  can send requests.  $R_i$  is the set of requests  $\varphi_i$  receives from other agents, described in more detail in Section 5.

### 4 Framework

The planning process of EGOPLAN is illustrated by Figure 1. This is an iterative process in which each iteration generates the next part of the Sidekick’s final plan. The Egobots re-plan completely in each iteration. Once there are no more requests for the Sidekick to complete, the process concludes.

First a set of agents  $\Phi$  must be chosen for the problem  $\langle P, I, G, A \rangle$ , giving the multiagent problem  $\langle \Phi, \Psi, P, I, G, A \rangle$ . Each agent will either be an Egobot or Sidekick.

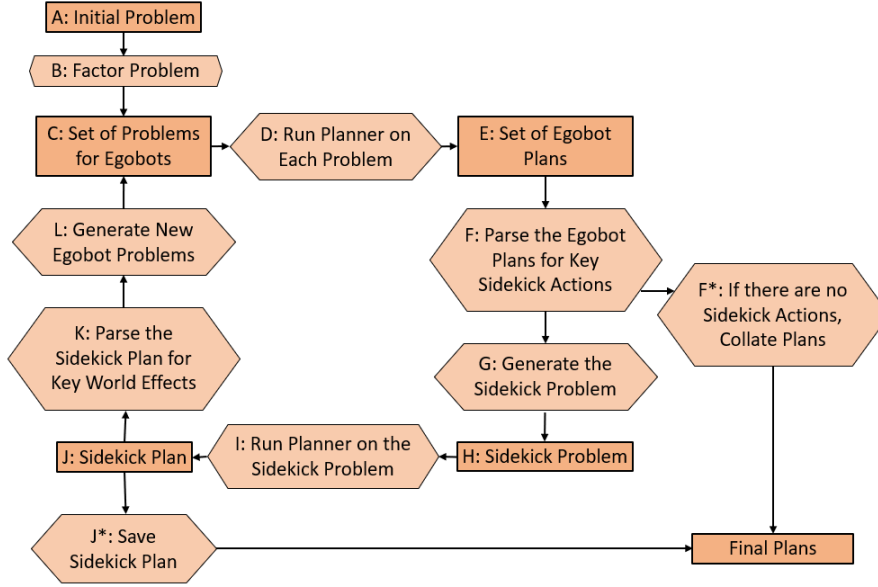


Figure 1: Execution flow of the EGOPLAN framework.

An Egobot  $\varphi_i$  has a limited knowledge of the world ( $P_i$ ), a goal condition ( $G_i$ ), executive control over some actions ( $A_i$ ), and knowledge of the Sidekick’s executive actions ( $\hat{A}_i = A_i \cup A_s$ ). A Sidekick  $\varphi_s$  has full knowledge of the world ( $P_s = P$ ), an empty goal condition ( $G_s = \emptyset$ ), executive control over some actions  $A_s$ , and no knowledge of other actions ( $\hat{A}_s = A_s$ ). The planning agents are split into these two roles so that no agent will be expected to both receive and send requests.

Because Egobots can only ask for help and Sidekicks can only give help, there is no capacity for interaction between Egobots or interaction between Sidekicks. This limits EGOPLAN to multi-agent problems in which no agents need to both receive and give help. Similarly, Egobot planning agents cannot handle interference from other Egobots, and Sidekick planning agents cannot handle interference from other Sidekicks. This limits EGOPLAN to domains in which agents of the same type are kept separated. Future work to generalise EGOPLAN to different settings without these limitations is discussed in Section 7.

The multi-agent problem  $\Pi$  is factored into separate problem files for each Egobot (fig. 1 A-C). The planner is then called once for each Egobot problem generating plans (fig. 1 D, E) in which the executive actions of that Egobot and of the Sidekick are used to meet that Egobot’s goals.

Each Sidekick action in an Egobot’s plan whose effects achieve part of the Egobot’s goal or supports a later Egobot action is transformed into a *request* for the Sidekick (fig. 1 F). The details of transforming a Sidekick action into a request are described in Section 5.

A problem is generated for the Sidekick with soft goals for each request. The planner generates a partially optimised plan in which the Sidekick completes as many requests as it is able (fig. 1 G-J). This plan will be part of the final plan

(fig. 1 J\*). Each part of the Sidekick plan that modifies part of the world state known to an Egobot (some  $p \in P_i$ ) is transformed into a timed effect (fig. 1 K). In PDDL2.2 we use Timed Initial Literals (TIL). This includes an effect that adds the final location of the Sidekick at the end of its plan.

A set of Egobot problems is generated (fig. 1 L, C). These problems include the effects of Sidekick actions as TILs including the TIL that re-enables the Sidekick. A plan for these Egobot problems cannot include actions for the Sidekick before the TIL that re-enables the Sidekick, and so new requests are only generated later than that time.

Once an Egobot generates a plan that does not include any Sidekick actions (fig. 1 F\*), this is one of the final Egobot plans and that Egobot does not replan in later iterations. Once all Egobots have generated a plan that does not include any Sidekick actions, the final Egobot plans and the set of Sidekick plans are combined into a single final plan for the original problem.

## 5 Requests

In this paper we define a *request*  $r = \langle G_r, S_r \rangle$  as a communication from one planning agent to another containing a temporally extended goal and the value of the goal. In EGOPLAN, useful requests are identified by giving the Egobot planning agents knowledge of Sidekick actions, so that their plans include helpful, feasible Sidekick actions.

Each precondition of an Egobot action that is supported by the effect of a Sidekick action becomes a request goal  $G_r$ . The condition is first generalised by “lifting” symmetric objects (Fox and Long 1999). For example, an Egobot might request that a welder is delivered to a location, represented by the condition  $(\text{at loc12 welder4})$ . However, the Egobot plan is agnostic to the specific welder object, so the goal becomes:  $(\text{at loc12 ?w - welder})$ . The goal is

then temporally extended with a deadline at or soon after the end of the relevant Sidekick action in the Egobot plan.

The value of each request is defined by a *score* fluent  $S_r$  added to the Sidekick problem during factorisation. The values were set using domain specific information about the importance of different requests. The Sidekick problem is an optimisation problem to maximise *score* by fulfilling requests. Actions that fulfill requests were modelled such that they can only be applied if an Egobot has requested it, can only be completed once per request, can only be completed before the request’s deadline, and increase the score by  $S_r$ .

## 6 Evaluation

A domain for robotic maintenance of a space station was used to test EGOPLAN. Problems using this domain have goals to inspect or patch panels. Inspect actions require no tools while patch actions require and consume a patch held by the robot. When patches are separated from the areas with patch goals, the requirement for a Sidekick to deliver patches means that the patch goals can only be achieved by coordination between two agents.

In this domain there are two types of acting agent representing bulky robots which must be supported by a lightweight drone. The bulky acting agents (Egobots) are limited to a set of locations such that no two Egobots can interfere with each other. The lightweight drone (Sidekick) cannot perform the patch panel action, but it can travel to any location to inspect panels and deliver patches.




Figure 2: Example Problem Layout

All problems were generated with one Sidekick agent (fig. 2 “sid”) starting in a central location (fig. 2 “loc00”), which initially contains all patch objects. Each problem had some number of Egobot agents and some number of accessible locations per Egobot. Figure 2 shows five Egobots each with three accessible locations. These locations were arranged in lines, one line per Egobot, with each line connecting to the central location where the Sidekick starts. The Egobots were each given 4 goals per location they could access, giving the

problem shown in Figure 2 a total of 60 goals. 1 in 5 of the goals were patch goals, the rest were inspect goals.

Problems were generated with  $E = 3 \dots 20$  Egobots and  $L = 3 \dots 13$  locations per Egobot. The problem was solved by EGOPLAN with the temporal planner OPTIC-CPLEX (Benton, Coles, and Coles 2012). The same planner was applied to the unfactored problem.

The two approaches were compared based on planning success, planning time, and plan length. The single agent planner was given a timeout of 30 minutes. The time reported for the single agent planner is the time taken to find the first plan, which in all cases was also the best plan.

The iterative process of EGOPLAN makes it harder to apply timeouts, as timeouts must be set for each agent in each iteration. EGOPLAN cannot backtrack to optimise, so EGOPLAN cannot be directly asked to optimise for 30 minutes.

EGOPLAN was given internal timeouts for each iteration: Egobot problems were solved with a timeout of 30 minutes, returning the first plan found without optimisation; Sidekick problems were given  $5 * E * L$  seconds and returned the best plan found in that interval. This Sidekick timeout was chosen through trial and error as it generally avoided the Sidekick failing to find a plan. The time reported for EGOPLAN was the sum of the times taken for the planner to find the plans used in each iteration, not including time used to optimise that did not result in a better plan being found. This was selected as the closest equivalent to the single agent reporting the time taken to find the best plan, but not reporting the entire 30 minutes of planning.

In total 547 problems were attempted by each approach (large problems were not attempted if preliminary testing showed problems with significantly lower  $N$  and  $L$  could not be solved by either approach). The single agent planner was able to solve 12% of the problems while EGOPLAN was able to solve 59%. There were no problems solved by the single agent that were not also solved by EGOPLAN.

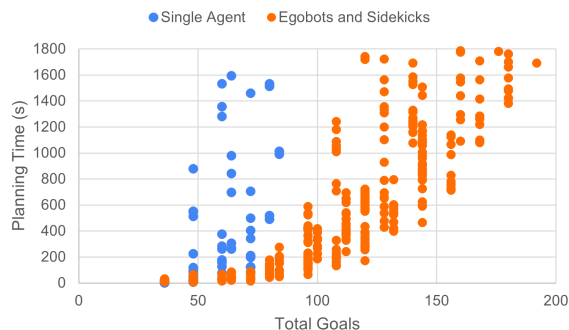


Figure 3: Planning Time by Complexity

The time taken to solve problems is shown in Figure 3, with the total number of goals in the problem used as an approximation of the problem’s complexity. The single agent approach couldn’t solve any problems with more than 96 goals, while the most complex problem solved by EGOPLAN has 192 goals. The wide range of times taken to solve problems with the same number of goals is due to factors

other than total goals, such as the distribution of patch goals, impacting the complexity of the problem.

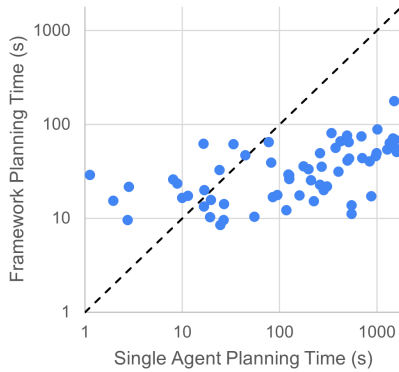


Figure 4: Comparison of Planning Time

Figure 4 shows the 63 problems that the single agent approach was able to solve and compares the time taken for both approaches to solve these problems. The dots above the dashed line show the problems which the single agent approach was able to solve faster than EGOPLAN. There were 14 problems solved faster by the single agent approach, these are problems with up to 60 goals that didn't take long for either approach to solve. The other 49 problems that were solved by the single agent approach were solved faster by EGOPLAN, and a further 259 problems were not solved by the single agent approach but were solved by EGOPLAN.

Figure 3 shows that for both approaches more complex problems took more time to solve, but that planning time increased with complexity much faster for the single agent than for EGOPLAN. This was expected because factorisation is more effective for larger problems, but typically comes at a cost to plan quality as the single agent approach considers more interactions than the factored approach.

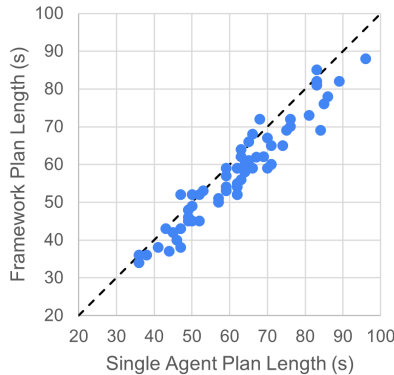


Figure 5: Comparison of Plan Length

Plan quality is measured as total plan duration, shorter being better. Quality was similar for both approaches for the 63 problems the single agent approach was able to solve, as shown in Figure 5. Despite solving a factored problem, EGOPLAN was able to find slightly higher quality plans in

general. The solution found by EGOPLAN is valid with respect to the unfactored problems, and better solutions might also exist, but were not found within the time limit due to the complexity of the problem.

A problem factored by agents can reduce plan quality when interaction between those agents is restricted by the factorisation. EGOPLAN does not restrict one-way binary interaction between agents, and the high quality plans found by EGOPLAN imply that this level of interaction was sufficient for planning in this domain. Applying EGOPLAN to domains which require more complex interaction is discussed in Section 7.




Figure 6: Most Complex Problems Solved by Each Approach

The two forms of complexity in the test problems were the number of Egobots and the number of locations (there were 4 times as many goals as locations) per Egobot. Figure 6 shows how many Egobots and how many locations per Egobot there were in the most complex problems solved by each approach. It was expected that EGOPLAN would be able to solve problems with more Egobots than the single agent approach could, because the advantage of EGOPLAN is the factoring of the problem. However, EGOPLAN also outperformed the single agent approach for low numbers of Egobots with complex problems. With as few as three agents, the single temporal planning problem was not able to scale to large numbers of goals, and benefited from factorisation.

## 7 Conclusion

In this paper we presented EGOPLAN, a framework that allows a single agent planner such as OPTIC-CPLEX to solve factored multi-agent problems with coordination. Our evaluation showed that EGOPLAN allowed OPTIC-CPLEX to scale to larger problems and to solve large problems more quickly than when OPTIC-CPLEX is used without EGOPLAN. As EGOPLAN is a framework that extends a single agent planner, and generates requests through planning, it is able to cope with both expressive problem features (i.e. action concurrency) and required coordination between agents.

In our current implementation and evaluation, the Egobot agents plan in sequence. In reality, multiple physical agents would necessarily perform this step in parallel, potentially reducing the real time required for EGOPLAN to find a solution by a significant factor. As EGOPLAN handles coordina-

tion through the sending and receiving of requests, it is well suited to decentralised planning and as a next step we will integrate EGOPLAN with existing task-allocation systems for multi-robot planning (Carreno et al. 2020). In addition, we will compare to existing centralised factored planning approaches, such as TFPOP, investigating the differences and limitations in factoring and problem features. Beyond this, there are several opportunities for generalising the framework which we discuss below.

**Two-way communication.** The framework will be extended to handle scenarios in which all agents are capable of both sending and receiving requests. This more advanced interaction between agents would allow a planner to solve more general problems than EGOPLAN, which is limited to problems in which the agent mappings ( $\Psi$ ) can be represented as a tree. Additionally, two-way communication would remove the planning bottleneck of the single Sidekick in EGOPLAN. Implementing two-way communication requires addressing negotiation “deadlock”, in which agents cannot reach a consensus on how to act.

**Generating requests efficiently.** EGOPLAN identifies useful requests and confirms the request is feasible through plan generation, but this can require significant planning effort. This effort would increase significantly if two-way communication allowed each agent to request help from any other agent. This can be mitigated in two ways. First, limiting which other agents each agent can send requests to using some agent mapping  $\Psi$  that ensures each agent only models itself and nearby agents. Second, by using an abstracted model of the supporting agents, approximating the feasibility of requests instead of directly planning to confirm.

**Action commitment.** EGOPLAN iteratively builds a final plan, committing to actions included in each Sidekick plan. Understanding when it is safe to commit to an action during planning or execution is very challenging (Cushing, Benton, and Kambhampati 2008; Shperberg et al. 2020; Cashmore et al. 2021). We intend to investigate past work in action commitment and apply it to the context of a multi-agent environment, where a planning agent may not be aware of the full consequences of its actions.

## References

- Benton, J.; Coles, A.; and Coles, A. 2012. Temporal planning with preferences and time-dependent continuous costs. In *International Conference on Automated Planning and Scheduling*.
- Brafman, R. I., and Domshlak, C. 2006. Factored planning: How, when, and when not. In *AAAI*.
- Buksz, D.; Cashmore, M.; Krarup, B.; Magazzeni, D.; and Ridder, B. 2019. Strategic-tactical planning for autonomous underwater vehicles over long horizons. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) 2018*, 3565–3572.
- Carreno, Y.; Pairet, E.; Petillot, Y.; and Petrick, R. P. A. 2020. A decentralised strategy for heterogeneous auv missions via goal distribution and temporal planning. In *Proceedings of the International Conference on Automated Planning and Scheduling*, 431–439.
- Cashmore, M.; Coles, A.; Cserna, B.; Karpas, E.; Magazzeni, D.; and Ruml, W. 2021. Replanning for situated robots. In *Proceedings of the International Conference on Automated Planning and Scheduling*, 665–673.
- Crosby, M.; Jonsson, A.; and Rovatsos, M. 2014. A single-agent approach to multiagent planning. In *ECAI*.
- Crosby, M.; Rovatsos, M.; and Petrick, R. 2013. Automated agent decomposition for classical planning. *ICAPS 2013 - Proceedings of the 23rd International Conference on Automated Planning and Scheduling* 46–54.
- Cushing, W.; Benton, J.; and Kambhampati, S. 2008. Replanning as a deliberative re-selection of objectives. Technical report, Arizona State University CSE Department.
- Dix, J.; Muñoz-avila, H.; Nau, D. S.; and Zhang, L. 2003. Impacting shop: Putting an ai planner into a multi-agent environment. *Annals of Mathematics and AI* 37:381–407.
- Edelkamp, S., and Hoffmann, J. 2004. Pddl2. 2: The language for the classical part of the 4th international planning competition. Technical report, Technical Report 195, University of Freiburg.
- Fox, M., and Long, D. 1999. The detection and exploitation of symmetry in planning problems. In *IJCAI*, volume 99, 956–961.
- Jezequel, L., and Fabre, E. 2012. A-sharp: A distributed version of a\* for factored planning. In *2012 IEEE 51st IEEE Conference on Decision and Control (CDC)*, 7377–7382.
- Korf, R. E. 1987. Planning as search: A quantitative approach. *Artificial Intelligence* 33(1):65 – 88.
- Kvarnström, J. 2011. Planning for loosely coupled agents using partial order forward-chaining. In *Proceedings of the 21st International Conference on Automated Planning and Scheduling (ICAPS)*.
- Nissim, R.; Brafman, R.; and Domshlak, C. 2010. A general, fully distributed multi-agent planning algorithm. In *Proceedings of the International Joint Conference on Autonomous Agents and Multiagent Systems*, 1323–1330.
- Shperberg, S.; Coles, A.; Karpas, E.; Shimony, E.; and Ruml, W. 2020. Trading plan cost for timeliness in situated temporal planning. In *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI-20*, 4176–4182.
- Srivastava, B. 2000. Realplan: Decoupling causal and resource reasoning in planning. In *AAAI/IAAI*, 812–818.
- Torreño, A.; Onaindia, E.; Komenda, A.; and Stolba, M. 2018. Cooperative multi-agent planning: A survey. *ACM Computing Surveys* 50(6):1–32.
- Torreño, A.; Onaindia, E.; and Óscar Sapena. 2014. Fmap: Distributed cooperative multi-agent planning. *Applied Intelligence* 41(2):606–626.
- Wang, D., and Williams, B. 2015. Tburton: A divide and conquer temporal planner. In *AAAI*, 3409–3417.
- Yu, H.; Marinescu, D.; Wu, A.; and Siegel, H. J. 2004. Planning with recursive subgoals. In *Knowledge-Based Intelligent Information and Engineering Systems*, 7–27.