



Using Pre-Computed Knowledge for Goal Allocation in Multi-Agent Planning

Nerea Luis¹ · Tiago Pereira^{2,3,4} · Susana Fernández¹ · António Moreira^{2,4} · Daniel Borrajo¹ · Manuela Veloso³

Received: 24 June 2018 / Accepted: 11 April 2019 / Published online: 15 May 2019
© The Author(s) 2019

Abstract

Many real-world robotic scenarios require performing task planning to decide courses of actions to be executed by (possibly heterogeneous) robots. A classical centralized planning approach has to find a solution inside a search space that contains every possible combination of robots and goals. This leads to inefficient solutions that do not scale well. Multi-Agent Planning (MAP) provides a new way to solve this kind of tasks efficiently. Previous works on MAP have proposed to factorize the problem to decrease the planning effort i.e. dividing the goals among the agents (robots). However, these techniques do not scale when the number of agents and goals grow. Also, in most real world scenarios with big maps, goals might not be reached by every robot so it has a computational cost associated. In this paper we propose a combination of robotics and planning techniques to alleviate and boost the computation of the goal assignment process. We use Actuation Maps (AMs). Given a map, AMs can determine the regions each agent can actuate on. Thus, specific information can be extracted to know which goals can be tackled by each agent, as well as cheaply estimating the cost of using each agent to achieve every goal. Experiments show that when information extracted from AMs is provided to a multi-agent planning algorithm, the goal assignment is significantly faster, speeding-up the planning process considerably. Experiments also show that this approach greatly outperforms classical centralized planning.

Keywords Multi agent planning · Actuation maps · Goal allocation · Robotics · Distributed planning · Path planning

1 Introduction

Real-world robotic scenarios, in which a set of robots need to solve a certain amount of tasks, usually require the combination of path-planning and motion-planning techniques. An example of this type of scenarios is the coverage problem, which consists of distributing the space among the set of robots, so that each one explores a certain region of the environment. The coverage problem planning task is to find a route for each robot so that all the feasible space is covered by the robots' actuators, while minimizing the execution

time. Vacuum cleaning robots can be potential candidates for this problem. We assume that we have a team of heterogeneous robots with different sizes. While the smallest robot can reach more areas, a bigger robot cleans a wider area while traveling a smaller distance. Nevertheless, other similar problems can also be solved with our contributed technique e.g. heterogeneous robots executing surveillance tasks, cooperative mapping of the environment or search and rescue tasks. As long as there exist (1) some navigation graph where we can extract information to help the planner; and (2) agents with similar or different capabilities, it will be a potential domain to solve with our approach. We have encoded our problem as a Multi-Agent Planning (MAP) task. Automated planning is the field of Artificial Intelligence which deals with the computation of plans. A plan is a sequence of actions that, if executed in order from the initial state, reaches another state where all the feasible goals are achieved.

The problem is modeled with the standard PDDL language [14]. For that purpose, we use a discrete representation of the map, e.g., a 2D grid of waypoints, which is 8-connected regarding robot motion. This means robots can

✉ Nerea Luis
nluis@inf.uc3m.es

¹ Universidad Carlos III de Madrid, Avda. Universidad 30, Madrid, Spain

² Faculty of Engineering, University of Porto, Porto, Portugal

³ Carnegie Mellon University, Pittsburgh, PA, USA

⁴ INESC-TEC, Porto, Portugal

move horizontally, vertically and diagonally in the grid of waypoints. Robots can move from one waypoint to another as long as they are grid neighbors and do not collide with obstacles. Moreover, robots can actuate other waypoints if their distance to the robot's current position is less than their actuation radius. Here we consider actuation as a robot performing an operation that results in some change to the environment. However, the actuation capabilities could be modeled not only as the operations where a robot changes its surrounding, but also as any perception-like operation. The planning problem to solve would be equivalent if the overall goal is, for instance, to clean all reachable space or to measure the temperature everywhere. Our framework allows to solve planning problems where sensing operations need to be executed at specific waypoint locations, such as mapping the Wi-Fi signal strength in buildings, taking measurements of temperature and humidity on a set of pre-defined locations such as a computer cluster and server sites, inspecting some regions or even executing some surveillance tasks. For vacuum cleaning robots, the actuation is performed through the robot's actuator that cleans the floor. Cleaning a specific waypoint location can be seen as the robot accomplishing a specific cleaning operation, which is an actuation goal. Thus the coverage problem would represent a set of robots moving through a map and accomplishing multiple actuation goals. To accomplish those goals, the robots would have to execute multiple cleaning operations, i.e. actuating on all available waypoints on the environment with the objective of cleaning all reachable regions of the environment. In the cleaning case, the actuation capabilities of each robot depend on the specific actuator each robot uses. In general, for each robot and domain a different actuation model can be considered. For example, the actuation radius can either be smaller than the robot's footprint, equal to the robot's footprint, or extend further than the robot's footprint. An example of the last case is a mobile manipulator, where an arm can be extended and actuate on regions beyond the robot's range in terms of its shape and footprint. In this paper, for simplicity we only focus on two cases. First, circular robots where the actuation range is smaller or equal to the robot footprint's radius, and second, any-shape robots whose actuation range is the same as its footprint.

We expect the planner to output a plan that accomplishes all the feasible actuation goals, by moving the robots to all the reachable locations from where they can actuate the goal waypoints. Given the robot heterogeneity, some goals might only be feasible for a subset of the robots.

From a MAP point of view, the multi-robot problem we propose forces us to deal with two issues regarding the performance of the planning process: (1) the size of the search space grows with the number of waypoints and goals; and (2) some goals are not feasible for some robots. On one hand, real-world scenarios are big enough to make

almost impossible for a planner to solve this problem in a reasonable amount of time by just assigning all goals to all agents (following a centralized planning approach). On the other hand, some Multi-Agent planners invoke a goal-allocation phase before starting to plan to decrease the effort of computing individual plans [4, 24]. During goal allocation, a relaxed plan is computed per goal and robot to either return an estimated cost or to identify unfeasibility. This process would be repeated multiple times, concretely $|Agents| \times |Goals|$ resulting in a huge loss of computation time, especially when exploring all of the search space to identify unfeasible goals.

Therefore, we contribute a methodology that uses Actuation Maps (AMs) to extract path-planning related information. That information is used to later boost the performance of a multi-agent planner. In concrete, we use these maps as a preprocessing step to speed-up the goal assignment phase. AMs are built only once before the planning process, one per robot, at a very low cost in comparison to the impact on time savings observed later in goal assignment. The AMs, from the robot-dependent reachability maps [27], not only determine the feasibility of each pair robot-goal but also allow us to efficiently compute an estimated cost of achieving that goal. As a result, the planner receives the estimated cost information as input, and saves time by simplifying the goal allocation computation, directly assigning goals to robots.

In our previous work, we showed that combining a preprocessing operation with MAP could bring huge savings on planning time [29]. More specifically, we used AMs to estimate the cost of each robot when actuating each goal waypoint. Then, we introduced those estimations into the goal assignment phase to distribute the goals among the robots and plan individual paths that together solved the coverage problem.

In this paper, we extend that approach with various contributions:

- Introduction of collision avoidance, with a replanning phase solving the conflicts between the individual paths of each robot.
- Generalization of robot shape models, from circular only to any-shape, updating both the AMs and the information extraction to generate the PDDL problem.
- Evaluation of four different configurations of our approach against two centralized planners and three multi-agent planners.
- Description of the overall planning architecture, now organized into four modules, and detailed explanation of the preprocessing step.
- Description of the type of planning instances in the coverage problem where our approach has the greatest impact.

- Description of the Goal Allocation process and the deletion of unfeasible goals.
- Description of the MAP algorithm - including a new version that detects and solves interactions among agents.

This paper is structured as follows: in Section 2 we describe the kind of problems for which our approach has been designed, and we include some planning formalizations. Then, on Section 3 we describe the coverage problem we want to solve. On Section 4, we present our contributed planning approach with pre-computed knowledge. Section 5 explains in detail the preprocessing step, which is our main contribution. After that, Section 6 contains the description of the MAP algorithm. On Section 7 we briefly extend our approach to MAP problems that involve interactions between robots. Then, on Section 8, we extend the formalization of circular robots to any-shape robots. On Section 9 we show illustrative experimental results of our algorithm on different scenarios. Finally, we discuss the related work and we present our conclusions and directions for future work.

2 General Problem Formulation

Our approach can be easily applied to any robotic problem that involves at least the following elements:

- a map of the environment;
- a set of potential goals, such as cleaning actuation goals, to be executed by a set of agents over the environment;
- a way to model that scenario into a PDDL domain and problem.

The map can be modeled in different ways. In this work, we start with an image representing the world's floor plan. From that image, we extract a navigation graph over a grid of waypoints, which is used to generate the PDDL problem. However, from the floor plan we can also extract additional information related to the environment, which is accomplished through the process of building the AMs.

The set of potential goals can vary depending on the problem to solve. In this paper we are focusing on the coverage problem and as a result it is enough for the robots to move through the environment. On alternative problems, the goals could be: looking for objects, opening doors or achieving some clients' orders through the environment.

The potential of our approach relies on the ability to extract information from the map related to the goals to be achieved. The aim is to transform that information into a set of estimation costs that can speed up the planning process.

For the coverage problem, we compute the estimated cost as the distance from the robots to each of the waypoints. Again, on alternative problems, the estimated cost could be related to the distance to a required object, the dangerousness or reliability of a path, or the features of a robot, such as maximum velocity or the existence of manipulation capabilities.

In order to transform this kind of problems into PDDL we have to model (1) a domain; (2) a problem; and compute (3) a set of estimated costs. The domain and problem are a lifted representation in predicate logic of the planning task. However, most of the planners always perform a grounding transformation from the domain and problem to generate the planning task. The aim of solving a planning task is to find a plan that reaches the goals specified on the task while minimizing some cost.

Definition 1 (Planning Task (Single Agent)) A single-agent classical planning task [12] is a tuple $\Pi = \langle F, A, I, G \rangle$, where F is a set of propositions, A is a set of instantiated actions, $I \subseteq F$ is an initial state, and $G \subseteq F$ is a set of goals.

Each action $a \in A$ is described by (1) a set of preconditions ($\text{pre}(a)$) that represent literals that must be true in a state to execute the action; (2) and a set of effects ($\text{eff}(a)$), which are literals that are expected to be added ($\text{add}(a)$ effects) or removed ($\text{del}(a)$ effects) from the state after the execution of the action. The definition of each action might also include a cost function, $C(a)$ (the default cost is one). The solution of a planning task is a plan, which is a sequence of actions $\pi = (a_1, \dots, a_n)$ that, if executed in order from the initial state, reaches a state where all the goals in G are satisfied.

As we are working with multiple agents, we consider a MAP formalization where a set of m agents, $\Phi = \{\phi_1, \dots, \phi_m\}$, has to solve the given task.

Definition 2 Multi-Agent Planning task. The MAP task is formed by a set of planning subtasks, one for each agent, $M = \{\Pi_1, \dots, \Pi_m\}$ where M refers to the MAP task. Each planning subtask Π_i includes only the facts, actions, goals and initial state related to the agent ϕ_i .

Last, we define the estimated cost per agent and goal.

Definition 3 Estimated cost per agent and goal (EC). $EC = \{(g, \phi_i, c) \mid g \in G, \phi_i \in \Phi, c = C(g, \phi_i)\}$ such that the C function represents the cost for an agent ϕ_i to reach the goal position g from the agent's initial state. If a goal cannot be reached by an agent, c will be ∞ .

Usually, these estimated costs are computed to divide the goals among the agents before the planning process starts. In

Listing 1 Action Navigate in PDDL

```
(:action navigate
:parameters (?r – robot ?y – waypoint ?z – waypoint)
:precondition (and (connected ?r ?y ?z) (at ?r ?y))
:effect (and (not (at ?r ?y)) (at ?r ?z))
)
```

planning, c is obtained with a heuristic function. However, in our approach, the C function is given by the AMs. Even though MAP is domain-independent, our function to compute the estimated cost is domain-dependent and should be set up differently on each domain, though it would be very similar to the one we propose here for most of the related robotic domains.

3 The Coverage Problem Description

In this Section we consider heterogeneous teams of robots that actuate in a 2D environment, where the world is represented by a 2D image that can be downsampled to a 2D grid of waypoints. The Actuation Map gives information about the actuation capabilities of each robot, as a function of robot size and initial position [27].

At first, we assume that robots are circular and thus the only robot feature is its size, with 2D grid positions being rotation-invariant. Other shapes can also be trivially considered in our approach by extending the PDDL domain file to take into consideration robot orientation as well. This aspect is further explained in Section 8.

As it was previously said, we modeled the domain and problem using PDDL. The domain has two types of objects: *robots*, which act as agents; and *waypoints*, which represent positions in the discretized world.

The planning task of the coverage problem consists on finding a sequence of navigation and actuation actions for each robot so that all actuation goals are covered by any of the robots' actuators, while minimizing the execution time. An actuation goal refers to the atomic action of marking a waypoint as "actuated". Going back to the vacuum cleaning example, the actuation goal represents the objective of cleaning a certain waypoint position. For the coverage problem, the feasible actuation goals are waypoints that are reachable to some agent (in terms of actuation reachability). In alternative problems, the actuation goal of the coverage problem can be replaced by other types of goals.

Therefore, the set G is a list of waypoints to actuate on (positions that need to be covered). The PDDL domain we created has four predicates:

- `At (robot, waypoint)`: defines the robot position;
- `Connected (robot, waypoint, waypoint)`: establishes the navigation connectivity between waypoints, specified for each robot, and given the robot

heterogeneity, some connections might be traversable by some robots and not by others;

- `Actuated (waypoint)`: indicates which waypoints were already actuated; this predicate is used to specify goals;
- `Actuable (robot, waypoint, waypoint)`: shows which waypoints can be actuated by a robot when located on a different waypoint location.

For the coverage problem, robots have to actuate every waypoint in G (as long as the goal is feasible). The waypoints, when connected, generate a navigation graph for a certain robot. The three actions that are defined in the domain are called `navigate` (Listing 1), `actuate-on` (Listing 2) and `actuate-other` (Listing 3). The action `actuate-on` is used to actuate the current position of the robot. The third action is employed to mark a waypoint as actuated if the waypoint is identified as *actuable* from the robot's current location i.e. the waypoint is located inside the robot's actuation radius on the real environment. Actions `navigate` and `actuate-on/other` can be executed by an agent when it is placed on a waypoint. Both `actuate-on` and `actuate-other` have as effect the predicate `actuated`.

In order to generate a PDDL problem, the waypoints' grid resolution is defined in advance using a discretization step. After that, a navigation graph and a set of actuable waypoints are defined for each robot, taking into account their physical characteristics. All this information is generated on the preprocessing step, further explained in Section 5.

4 Architecture

As it was previously said, this work combines Actuation Maps (AMs) with Multi-Agent Planning (MAP). The contributed architecture can be seen in Fig. 1. It has been divided into four modules and receives as input the map of the environment, the general knowledge related to the task

```
(:action actuate-on
:parameters (?r – robot ?y – waypoint)
:precondition (at ?r ?y)
:effect (actuated ?y)
)
```

Listing 2 Action Actuate-on in PDDL

Listing 3 Action Actuate-other in PDDL

```
(:action actuate-other
:parameters (?r – robot ?y – waypoint ?z – waypoint)
:precondition (and (at ?r ?y) (actuatable ?r ?y ?z))
:effect (actuated ?z)
)
```

to solve and the features of the set of robots. The aim of each module is described below.

1. Actuation Maps module: it is in charge of generating the AMs for each given robot. It also extracts the map features that can potentially alleviate the planning process (e.g. estimation costs), and it generates the planning problem in PDDL. It is explained on Section 5.2.
2. Multi-Agent Planning Task Generation module: once the outputs from the prior module and the domain are received as input, the goal assignment process is launched. This module is in charge of dividing the goals among the agents following some goal-strategy. Then, a set of domain and problem is generated for each agent, which is known as factorization. It is explained on Section 5.3.
3. Multi-Agent Planning Algorithm module: the individual planning process and the merging phase are run on this module. It is explained on Section 6.
4. Conflicts solver module: if any interactions need to be solved, this module employs a plan-reuse-planner to fix them. It is explained on Section 7.

The following subsections explain the essential information regarding both components of our contribution: AMs

and MAP. We introduce here the base work needed to later explain the preprocessing step and the MAP algorithm.

4.1 Actuation Maps

Our system receives as input the *Environment map* which represents a 2D environment (e.g. building floor plan) and *m Robot* models with the agents' features. There is a third input provided by the user that refers to the *General knowledge* of the environment (i.e. tasks to solve). These three inputs represent the input information described in Fig. 1.

We briefly summarize here the process of building the AMs [27]. We assume there is an occupancy grid map, i.e., a gray-scale image representing the environment. In this image each pixel has a value with the probability of the corresponding world position being occupied by an obstacle. This occupancy grid map is first transformed into a binary image of free and obstacle pixels, using a fixed threshold. An example of the resulting black and white image is shown in Fig. 2a.

We define \mathcal{G} as the set with all pixel positions from the input binary image. The input image is a visual representation of \mathcal{M} , the set with the obstacle pixel positions, where white represents free space and black represents obstacles.

Fig. 1 Complete architecture that combines actuation maps and multi-agent planning

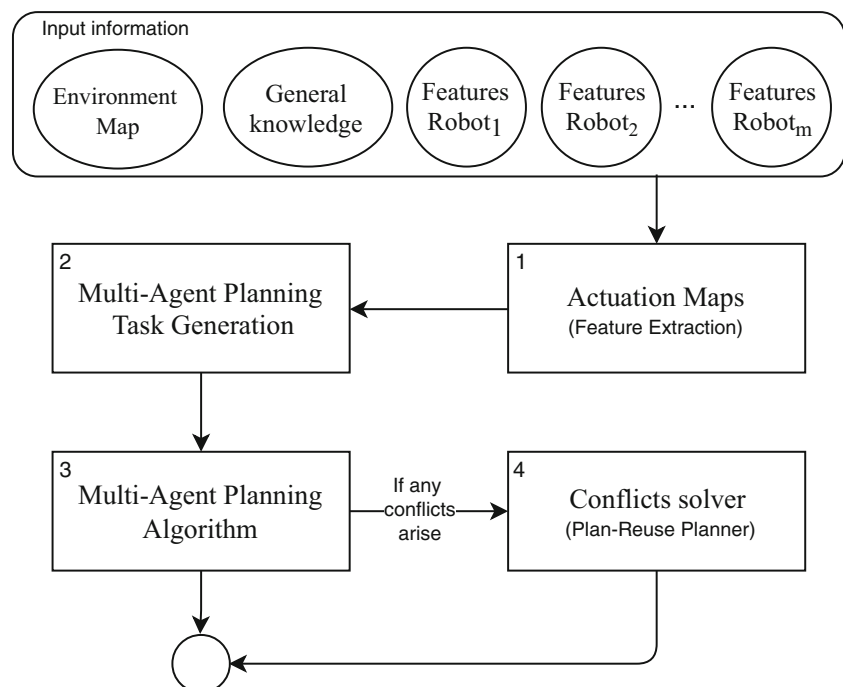



Fig. 2 Simulated map and two heterogeneous robots with different sizes in (a); colored regions represent the navigable space, $\mathcal{L}_i(\mathbf{r}_i^0)$, for 2 robots with different sizes, depending on size and initial position of robots



(a) Map



(b) Robot1



(c) Robot2

We define the structuring element as an image that represents the robot shape. Assuming first the circular robot case, we represent the robot with \mathcal{R}_i , the set with pixel positions from a circle with radius equal to the robot size. The morphological operation dilation on the obstacle set \mathcal{M} by \mathcal{R}_i is:

$$\mathcal{M} \oplus \mathcal{R}_i = \bigcup_{\mathbf{r} \in \mathcal{R}_i} \mathcal{M}_{\mathbf{r}} \quad (1)$$

where $\mathcal{M}_{\mathbf{r}}$ is the translation of \mathcal{M} by vector \mathbf{r} .

The visual output of applying this dilation operation to a map of obstacles is the inflation of obstacles by the robot size. The free configuration space, \mathcal{C}^{free} , is then defined as:

$$\mathcal{C}_i^{free} = \{\mathbf{p} \in \mathcal{G} \mid \mathbf{p} \notin \mathcal{M} \oplus \mathcal{R}_i\} \quad (2)$$

where \mathcal{G} is the grid set with all the pixel positions. The free configuration space represents the feasible positions for the robot center, but does not give any information about the regions that can be actuated by the robot.

In order to determine which regions of the environment are actuable, the partial morphological closing operation is used. Morphological closing is a dilation operation followed by a morphological erosion. Because dilation and erosion are dual operations, the morphological closing of obstacles (erosion applied to the image with inflated obstacles) is equivalent to the dilation of free configuration space.

However, the morphological closing cannot be used to determine the actuation capabilities, because it does not consider which points are reachable from the initial robot position, and different initial positions do change the overall reachability in terms of actuation. The partial morphological closing was introduced in order to consider the initial robot position when determining the actuation capabilities. In order to use the partial morphological closing, the algorithm needs to find the navigable regions first. The set of navigable regions from a starting robot point \mathbf{r}_i^0 is always a subset of \mathcal{C}_i^{free} .

$$\mathcal{L}_i(\mathbf{r}_i^0) = \{\mathbf{p} \in \mathcal{G} \mid \mathbf{p} \text{ connected to } \mathbf{r}_i^0 \wedge \mathbf{p} \in \mathcal{C}_i^{free}\} \quad (3)$$

The navigable set $\mathcal{L}_i(\mathbf{r}_i^0)$ is the set of points that are connected to the initial position \mathbf{r}_i^0 through a path of adjacent cells in the free configuration space.

In Fig. 2 we show a simulated map with 2 robots with different sizes, and the respective navigable space.

Finally, by applying the the second dilation operation of the morphological closing to the navigable set (subset of \mathcal{C}^{free}) instead of applying it to the free configuration space, we obtain the partial morphological closing operation. The *actuation space* is thus the dilation of the navigable space. The structuring element for this second operation is the one that models the actuation capabilities, \mathcal{T} , which dilates the navigable space according with the actuation model. If instead the structuring element \mathcal{R} is used again, that would be equivalent to assuming an actuating ability completely coincident with the entire robot footprint.

$$\mathcal{A}_i(\mathbf{r}_i^0) = \mathcal{L}_i(\mathbf{r}_i^0) \oplus \mathcal{T}_i \quad (4)$$


While the actuation space, $\mathcal{A}_i(\mathbf{r}_i^0)$, is the set containing all the waypoints that can be actuated by a robot, its representation as an image, as shown in Fig. 3, is the *Actuation Map*. We also use the term *Actuation Map* to refer to the overall technique to determine actuation capabilities of robots. In the figures below we show the actuation map, obtained after applying the partial morphological closing operation to the original map. The actuation map can be used as a visual representation to show what the robot can actuate from any point reachable from its initial position, and we use them to visually represent in figures the sets corresponding to each robot's actuation space.

As an example, we can consider again the vacuum cleaning robot case. The configuration space represents the possible center positions for the robot; the actuation space \mathcal{A} represents the regions the robot can clean; and the non-actuable regions are positions that the robot cannot clean. For circular vacuum cleaning robots, corners of the environment are non-actuable regions they cannot clean due to the robot's circular shape.


4.2 Multi-Agent Planning

When multiple agents are involved (e.g. robots, workers, drivers) we talk about Multi-Agent Planning (MAP) [34]. MAP computes a plan for/by a set of agents that jointly


Fig. 3 Colored regions in (b) and (c) represent actuation spaces for respective robots, i.e. the points in the environment that each robot can actuate, depending on their size and initial position shown in (a); the actuation capability in this example is completely coincident with the entire robot footprint, i.e., the actuation range is equal to the robot radius



(a) Map with 2 Robots



(b) Actuation Map 1



(c) Actuation Map 2

solve a planning problem. Usually, MAP tasks have to deal with some coordination issues among agents or the sharing of resources. From the perspective of MAP, planning domains exhibit a coupling level that ranges from loosely-coupled to tightly-coupled, depending on the degree of interaction between agents plans [7]. Agents on loosely-coupled domains barely interact with each other. Tightly-coupled domains have a considerable level of interaction among agents. This implies solving interactions while planning (agents' communication) or afterwards (conflict solving). Our domain is loosely-coupled, as agents barely interact with each other. However, there might be collisions that are not detected by the planner. A different approach involving collision-detection is further explained on Section 7.

In Multi-Agent Planning two main approaches have been commonly used: centralized and distributed. The centralized approach involves a master agent which knows everything about the agents and the environment. The master agent sees the rest of agents as resources and is also responsible for coordinating and solving the interactions that might arise during the planning process. On the distributed approach each agent builds its own plan synchronously with the rest of the agents. Depending on the amount of communication allowed among the agents, they will need to share their information during the planning process or they have to later merge their plans and solve the conflicts that might have arisen. Our MAP algorithm follows the latter, using plan merging [13, 25] to build the solution plan after the individual planning process. Thus, no cost of communication is involved and there is no implementation inside the algorithm regarding communication.

5 Preprocessing

The contributed preprocessing step is shown in Fig. 4. This is the point where both techniques, AMs and MAP, are combined and complement each other. Section 5.1 describes the discretization process, downsampling the waypoint density from the original pixel grid. Section 5.2 describes the

generation of goals, the detection of unfeasible regions and the computation of estimated costs. Section 5.3 describes the process of generating the MAP task and how the task is factorized (divided) in subtasks.

5.1 Discretization

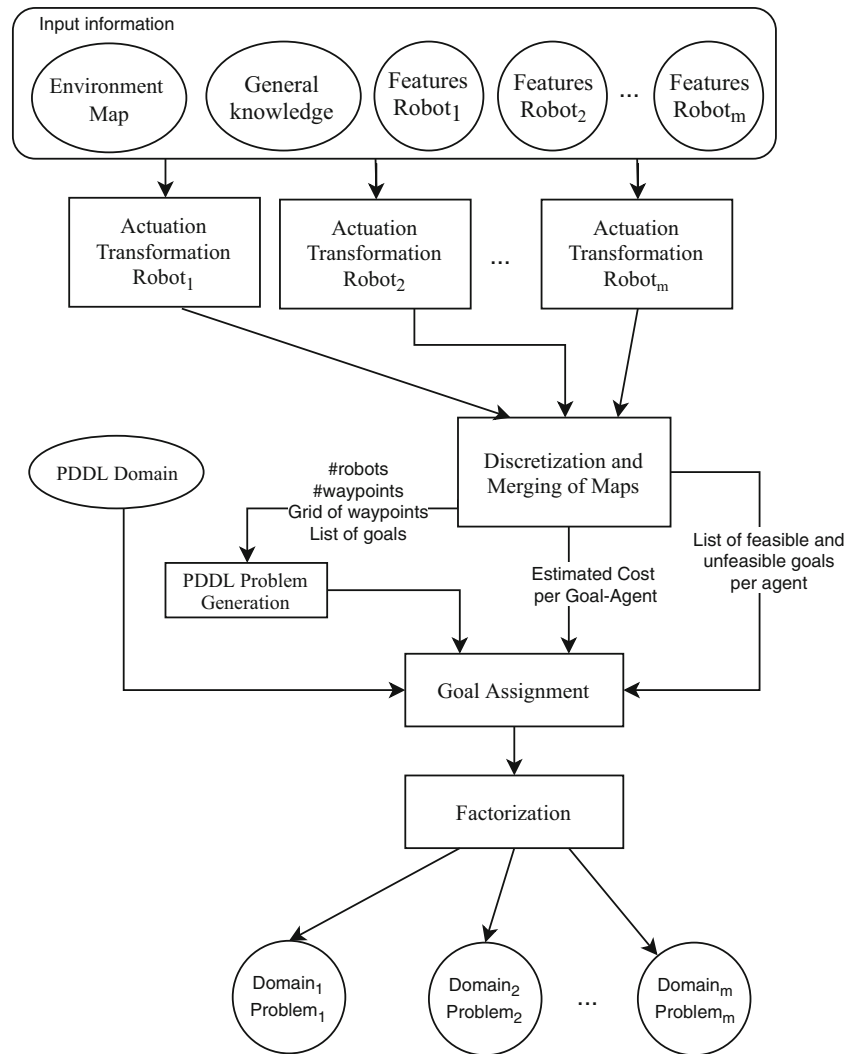
For the planning problem, it is possible to consider each individual pixel as a waypoint. However, that approach results in a high density of points that would make the planning problem excessively complex. Moreover, there is some redundancy in having points that are too close to each other, as their difference is not significant in terms of the environment size and localization accuracy.

Therefore, we reduced the set of locations from all pixels to a smaller set of locations. We considered again waypoints distributed into a grid, but now the grid-size is greater than one pixel. The reduced grid of waypoints is obtained using a downsampling rate s_r , such as the number of waypoints in each direction in the new grid is s_r times less than the number of waypoints in the original grid of pixels.

After downsampling, we can find the connectivity between points to construct the navigation graph of each robot, shown in Fig. 5a. It is also possible to find which waypoints can be actuated from other waypoints using the distance between them, as shown in Fig. 5b, by considering the maximum actuation radius.

The problem of such discretization is the change in the actuation space topology. Adjusting the position of waypoints could allow a better representation of the topology of the environment, but the multi-robot nature of the problem compromises that solution. In order to deal with multiple robots with different configuration spaces, for each agent, we independently adjust the waypoint position -temporarily- in a hidden manner invisible to the other agents. When discretizing each robot's configuration space, we might consider a waypoint as belonging to the free configuration space even if it is strictly outside it, as we assume an error margin to compensate for the discretization error. Nevertheless, we still maintain the original waypoint position in further steps, such as determining actuation

Fig. 4 Preprocessing stage before the planning process starts. First, inputs are processed in order to generate the AM for each agent. Then, a discretization is applied to generate all the required information for planning such as the navigation graph for the PDDL problem, and the list of estimated costs. Once this information has been generated, GA starts and the MAP problem is divided into subproblems with specific goals assigned to each individual problem



feasibility of that waypoint, and for visualization purposes as well. When determining the navigation graph of each robot, an unreachable waypoint position might be moved

to the closest point in the configuration space, if the adjustment is under a given margin δ . As stated previously, the adjustment is always temporary to the construction of

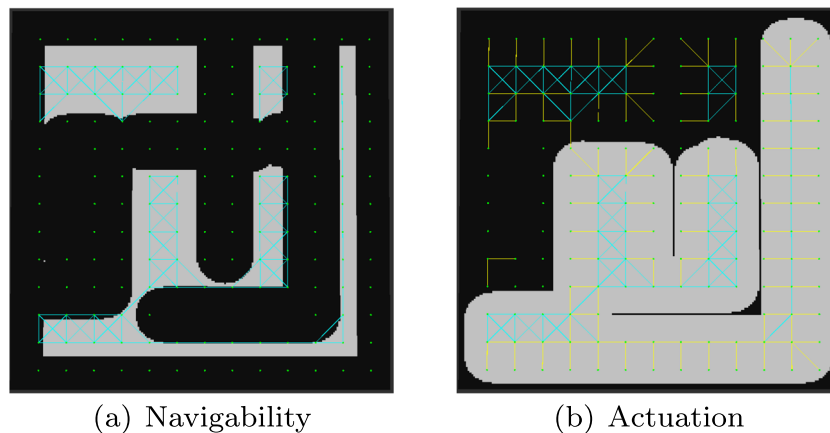


Fig. 5 In **a** we have an example of the free configuration space, with the discretization waypoints shown as green dots. The blue lines represent the connectivity between waypoints in the navigation graph of the robot. Using parameters δ and α it is possible to maintain the topology

of the free configuration space by allowing points in the navigation graph that were originally unfeasible for the robot. In **b**, the actuation map of the same robot, and the respective actuation graph represented with yellow lines

the connectivity graph of each robot. After the navigation connectivity is tested, the waypoint position resets to its default grid position for the next steps, such as determining the actuation feasibility, and the navigation and actuation graphs of other robots.

Moreover, when determining the connectivity of waypoints for the navigation graph, only the eight grid neighbors are considered. A^* is then used to determine the real distance between waypoints (e.g., around obstacles), and connectivity is only considered if the real distance is at most a factor of $\alpha = 1.2$ the straight line distance between them.

Finally, all waypoints that belong to the robot actuation map should be connected to some waypoint of its navigable graph. If that is not the case after the previous steps, we connect the isolated waypoints to the closest navigable vertex in line of sight, even if their distance is greater than the maximum actuation distance, again to compensate for the discretization error. Therefore, while the planner may return an `actuate` action to cover waypoint A from the navigable waypoint B in the discretized world, a real robot would have to move closer from the waypoint B to waypoint A in order to actuate the latter.

The grid density is chosen manually in order to adjust the level of discretization. As for the α and δ parameters, they were tuned empirically such as the free space topology is still maintained even while using lower density discretization of the environment. By trial and error, we found empirically that $\alpha = 1.2$ works for all the tested scenarios. As for the δ parameter, we set it to always start with a value of 3 pixels, then build the discretized model and verify if it is valid, i.e., if all the waypoints belonging to the actuation map become feasible for the respective robot in terms of the discretized representation. If not, we increment the parameter until a topologically consistent representation is found (number of feasible goals equals number of waypoints inside actuation map). Even though this fine-tuning methodology seems sensitive to the robot heterogeneity, the truth is that the final δ value depends on the size of the bigger robot, because the correct discretization of the configuration space is more sensitive to the δ parameter for bigger robots. Through experimentation, we found out that if a certain value of the δ parameter works well for the biggest robot, it always produces the correct discretization for smaller robots. Moreover, we also observed that $\delta = 4$ pixels worked well for all the different and very diverse maps we tested in our experiments with circular robots, only failing for the any-shape experiments where the configuration space discretization is more sensitive to the possible robot orientation. For the any-shape robot experiments, we found that $\delta = 6$ pixels was enough to obtain good a discretization for all the environments tested. The consistency of the δ parameter over different environment maps shows that these

parameters can be map-independent to a certain extent, with most of the work being easily automated.

5.2 Extracting Cost Information from Actuation Maps

The downsampling rate s_r was set manually, depending on the resolution one wants for the grid of waypoints. If the original pixel resolution would be used, the resulting grid of waypoints \mathcal{G}' would contain all pixels and it would be equivalent to \mathcal{G} . Otherwise the set \mathcal{G}' represents the grid waypoint positions after downsampling.

Using the concept of AMs it is possible to very easily find UG , the list of *unfeasible goals per agent*:

$$UG = \{g \in \mathcal{G}', \phi_i \in \Phi \mid g \notin \mathcal{A}_i(\mathbf{r}_i^0)\} \quad (5)$$

The positions in the actuation space $\mathcal{A}_i(\mathbf{r}_i^0)$ are feasible actuation goals for agent ϕ_i . The information from the UG list can speed-up goal assignment by avoiding computation related to unfeasible goals, but it does not provide any information about the cost for each robot to accomplish a feasible actuation goal.

For that purpose, we present the following extension. We build the navigable space \mathcal{L}_i in an iterative procedure, from the starting position \mathbf{r}_i^0 . In the first iteration we have $\mathcal{L}_i^0(\mathbf{r}_i^0) \leftarrow \{\mathbf{r}_i^0\}$, and then the following rule applies:

$$\mathcal{L}_i^j(\mathbf{r}_i^0) = \{\mathbf{p} \in \mathcal{G} \mid \exists \mathbf{q} \in \mathcal{L}_i^{j-1}(\mathbf{r}_i^0) : \mathbf{p} \text{ neighbor of } \mathbf{q} \wedge \mathbf{p} \in \mathcal{C}_i^{free} \wedge \mathbf{p} \notin \mathcal{L}_i^a(\mathbf{r}_i^0) \quad \forall a < j\} \quad (6)$$

When using this recursive rule to build the navigable space, we guarantee that any point in the set $\mathcal{L}_i^j(\mathbf{r}_i^0)$ is exactly at distance j from the initial position \mathbf{r}_i^0 . Because we want to build the complete actuation space, the iterative computation of Eq. 6 stops when all pixels have been expanded at least once.

Furthermore, if we build the actuation space sets with the intermediate navigable sets $\mathcal{L}_i^j(\mathbf{r}_i^0)$,

$$\mathcal{A}_i^j(\mathbf{r}_i^0) = \mathcal{L}_i^j(\mathbf{r}_i^0) \oplus \mathcal{T}_i \quad (7)$$

then the intermediate actuation set $\mathcal{A}_i^j(\mathbf{r}_i^0)$ represents the points that can be actuated by the robot from positions whose distance to \mathbf{r}_i^0 is j . The actuation space defined in the previous Section can also be alternatively defined as

$$\mathcal{A}_i(\mathbf{r}_i^0) = \{\mathbf{p} \in \mathcal{G} \mid \exists a : \mathbf{p} \in \mathcal{A}_i^a(\mathbf{r}_i^0)\} \quad (8)$$

The actuation cost is defined for $g \in \mathcal{A}_i(\mathbf{r}_i^0)$:

$$\mathcal{AC}_i(\mathbf{r}_i^0, g) = \min\{j \mid g \in \mathcal{A}_i^j(\mathbf{r}_i^0)\} + 1 \quad (9)$$

The actuation cost $\mathcal{AC}_i(\mathbf{r}_i^0, g)$ represents, for each $g \in \mathcal{A}_i(\mathbf{r}_i^0)$, the minimum number of actions needed for the robot to actuate the grid waypoint g if starting from the initial position \mathbf{r}_i^0 , measured in the pixel-based grid \mathcal{G} . In

Eq. 9, the minimum j^* represents the minimum distance (i.e., minimum number of navigate actions) needed to travel from \mathbf{r}_i^0 to some point from where g can be actuated. The added one in Eq. 9 accounts for the one actuate action needed to actuate g , after the j^* navigate actions needed to reach a place from where the robot can actuate g .

Thus, the cost function C presented previously in Section 2 is defined in Eq. 10, where s_r is the downsampling rate. The division by s_r transforms the estimated cost of actions measured in the pixel-based grid \mathcal{G} , $\mathcal{AC}_i(\mathbf{r}_i^0, g)$, to the respective cost value in the downsampled grid of waypoints \mathcal{G}' . The *ceil* function rounds up the result of the division to the smallest integral value that is not less than $\mathcal{AC}_i(\mathbf{r}_i^0, g)/s_r$. The cost function C is domain-dependent and works for the coverage problem. If a different problem is given as input, the cost function should be redefined.

$$C(g, \phi_i) = \text{ceil} \left(\mathcal{AC}_i(\mathbf{r}_i^0, g) / s_r \right) \quad (10)$$

Finally, the *Estimated Cost per Goal-Agent* list EC is defined in Eq. 11.

$$EC = \{ \langle g, \phi_i, c \rangle \mid g \in \mathcal{G}' \wedge \phi_i \in \Phi \wedge g \in \mathcal{A}_i(\mathbf{r}_i^0) \wedge c = C(g, \phi_i) \} \quad (11)$$

5.3 Multi-Agent Planning Task Generation

Once the discretization of maps has been performed, we have all the information needed to generate the MAP task M , which is formed by a domain (received as input) and a problem (generated through the discretization). The inputs to the Goal Assignment (GA) phase are (1) the PDDL domain; (2) the PDDL problem; (3) the list of estimated costs where c is computed as the number of steps for an agent to reach the goal position g from its initial position; and (4) the list of unfeasible goals $UG = \{g \in G, \phi_i \in \Phi \mid C(g, \phi_i) = \infty\}$. The cost of navigating between two neighbor grid waypoints is 1 unit. As long as EC is provided, UG is not used inside the MAP algorithm. The case when EC is not provided is later explained in this Section.

In Multi-Agent systems, in order to perform task allocation [8, 18] some strategy has to be determined or implemented, as the aim it is to divide the MAP task in subtasks to alleviate the planning process afterwards. In addition, a goal-assignment strategy (GAS) needs to be chosen to define the way goals are assigned to agents by the system.

In our approach we took the Load-Balance (LB) strategy previously defined in [2], that first calculates $k = \lceil \frac{|G|}{|\Phi|} \rceil$, which represents the average number of goals per agent. Then, it assigns each goal $g \in G$ to the agent $\phi_i \in \Phi$ that achieves g with the least cost. This strategy avoids, if

possible, assigning more goals than k to each agent. The LB assignment strategy is used when minimizing the maximum number of actions per agent (makespan). As a second option, we also took the Best-Cost (BC) strategy also defined before in [2], which simply assigns each goal to the agent that can achieve it with the least cost. The BC strategy is used when minimizing the total number of actions over all robots (plan length).

As in [2], only when the information about estimated cost per pair robot-goal is not available for some reason, our MAP algorithm would perform GA computing a relaxed plan using the FF heuristic [19]. This is not a contribution of the paper itself, as it was already in [24]. However, in that work, when a goal was unfeasible for every agent, it was assigned to all of them. In our approach, when the goal has been identified as unfeasible by all agents, the relaxed plan is not computed for that pair robot-goal and the goal is not included into the M task. This behavior is not common in classical deterministic Automated Planning, as planners expect that the problem does not contain any unfeasible goal. As our approach separates goal allocation from planning, we can easily deal with unfeasible goals. This small contribution gives us more flexibility when working for real environments, as it is better to obtain a plan that solves 95% of the goals than just failing during planning. To plan using soft-goals [23] or working on oversubscription planning [17, 31] would have been other ways to deal with unfeasibility, but they are out of the scope of this work. In summary, there are two contributions to the GA process: (1) the detection and deletion of unfeasible goals is a contribution that helps not only on skipping the computation of those relaxed plans but also avoids the planning process to fail; and (2) to use information from AMs, as the algorithm receives and processes the estimated costs from the AMs to skip the computation of the relaxed plans.

The first step is to allocate the feasible goals to the agents. This step uses the information of estimated costs received from AMs. Goal assignment phase (GA) returns as output (1) a subset of Φ' agents, $\Phi' = \{\phi'_1, \dots, \phi'_n\}$, that will be the only ones who will plan to solve the problem; and (2) a new MAP task $M' = \{\Pi_1, \dots, \Pi_n\}$. As a result, a specific PDDL domain and problem will be generated for each ϕ'_i agent which only includes the goals each agent has to achieve. As it previously mentioned, if a goal is unfeasible for all the Φ' agents, the MAP algorithm will discard it from the new MAP task M' , so that the process of planning will not fail. The output of this process is a pair (domain, problem) for each agent as in Fig. 4 after the *Factorization* step and in Eq. 12.

$$M' = \langle \Pi_1, \dots, \Pi_n \rangle = \left\{ \langle D_1, P_1 \rangle, \langle D_2, P_2 \rangle, \dots, \langle D_n, P_n \rangle \right\} \quad (12)$$

Fig. 6 High level description of the MAP algorithm. Inputs: MAP task (M'), Single-Agent Planner (SAP). Output: resulting plan (π) or no solution

The Multi-Agent Planning Algorithm

Inputs: M' , SAP
Output: π

```

1 Forall  $\phi'_i \in \Phi'$  do  $\pi_i = \text{plan}(\Pi_i, SAP)$ 
2  $\pi_{seq} = \text{concatenate}(\pi_1, \dots, \pi_n)$  /* where  $n = |\Phi'|$  */
3 If  $\text{valid}(\pi_{seq})$ 
4 Then return  $\pi = \text{parallelize}(\pi_{seq})$ 
5 Else return no-solution

```

6 The Multi-Agent Planning Algorithm

In order to solve the MAP task we use the distributed phase of a previous work of the authors [24]. Our algorithm receives as input a MAP task, which consists of a PDDL domain and problem files (Fig. 6) for each agent in $\phi'_i \in \Phi'$. On the first step, each agent builds its plan individually (line 1). Then the agents' plans are merged by a simple concatenation of plans in one resulting plan (line 2). If the merged plan is valid, we parallelize the plan (line 4) generated by the merging step.

Parallelization is performed in two steps: converting the input total-order plan into a partial-order one by a similar algorithm to [35]; and parallelizing this partial-order plan by ordering actions in the first time step that satisfies all ordering constraints in the partial-order plan. The main advantage of parallelization is to be able to return a plan in which on each time step, more than one action can be executed at the same time, taking advantage of having multiple robots in the environment as well. As the domain does not have any interactions, there is no need to fix the plan regarding interactions.

To set up the algorithm, an off-the-shelf planner had to be chosen. The advantage of this MAP algorithm is that any state-of-the-art planner can be easily included on the distributed planning phase without further code modification. Our configuration employs LAMA-UNIT-COST as SAP . LAMA-UNIT-COST corresponds to the first search that LAMA performs, using greedy-best-first with unit costs for actions [30]. The merged plan is validated using VAL [20], the validator from the International Planning Competition (IPC).¹ On Fig. 7 we show a solution example obtained from the MAP algorithm. It corresponds to the scenario called Corridor-High later on the experiments.

7 Dealing with Interactions

Real-world robotics environments might imply to deal with potential interactions among robots (at the very least)

e.g. collisions, sharing a resource, or cooperation. On the previous description of the coverage problem we did not explicitly consider any kind of interactions. Our idea was (1) to test first the scalability of the MAP algorithm; and (2) to generate, as fast as possible, a valid solution. Robots might occasionally collide at some specific step of the solution plan. However, that collision could be easily resolved during execution by forcing one of the robots to wait until the other robot has left the conflict zone. Then, the stopped robot will continue executing the rest of its plan. On the other hand, there is a subarea of Automated Planning called Planning by Reuse that has been widely employed in areas such as Case-Based Planning [6], or replanning when plan execution fails [15]. Usually, planners that perform plan repair receive three inputs: a domain, a problem and a plan to be fixed. Examples of this kind of planners are LPG-ADAPT [15] or ERRT-PLAN [5]. Therefore, an improvement of our approach is to detect and fix potential collisions right after the individual planning process using an off-the-shelf plan reuse planner. This new feature makes our architecture more robust when executing the solution plan in a real environment. Thus, we slightly changed our PDDL domain to track the collisions by adding a new predicate called *occupied*.

- *Occupied (waypoint)*: indicates that there is a robot on that waypoint.

That predicate is set as a new precondition of the *navigate* action described in Listing 3. This allows the agent to only traverse a connection if the destination waypoint is not occupied by a robot. The predicate *occupied* in combination with the parallelization algorithm avoids two situations: (1) Two robots cannot be on the same waypoint at the same plan step. (2) As we build a parallel plan, neither of them can swap positions during the same plan step. The reason is that in order to move *robot1* to y from z , y should be not occupied first. Same happens to z for *robot2*. Thus, both actions are mutually exclusive and the parallelization algorithm does not allow both actions to be performed together. One of the robots would need to move somewhere else first.

The new version of our MAP algorithm is described in Fig. 8. As the MAP algorithm starts with the individual

¹<http://icaps-conference.org/index.php/main/competitions>




Fig. 7 These figures represent the Corridor scenario used in the experiments. The waypoint discretization is shown in (a). The resulting path for each robot is shown in (b) to (e), after solving the planning problem

using load balance as goal-strategy. Path 1 belongs to the smallest robot. Path 4, to the biggest one

planning phase, no collisions will be detected at that point (line 1). After concatenation, the solution plan is validated by VAL (line 3). The validator will detect, if any, mutex actions related to occupied positions as explained above. If so, the plan will be invalid. As a result, the M' task and the invalid plan are sent to the plan reuse planner (line 6). When the plan is fixed, the parallelization step is applied (line 7). Finally, the MAP algorithm runs VAL again (line 8). If the plan is valid, it is returned as the solution. The configuration of our algorithm is the same as the previous version - the Single-Agent Planner is LAMA-UNIT-COST. The plan reuse planner is LPG-ADAPT.

8 Extending the Approach to Any-Shape Robots

For the case of non-circular robot footprints (Fig. 9), given the robot model is not rotation invariant, we need to discretize orientation as well. We use a world representation that is composed of multiple layers, where each layer

represents one orientation. We then using the partial morphological closing operation on each layer, and with that approach we determine individually for each orientation the corresponding actuation space [28].

First, the algorithm needs images to model both the robot and its actuation capabilities. Both are parametrized by images that can be rotated and scaled to represent any robot. As input, it is also necessary to give the center of the robot and actuation in terms of their model images, and their relative position.

Here we assume a quantization of the orientation given by n_θ layers. Using the input image for the robot model, we rotate it by $\theta_j = 2j\pi/n_\theta$, where $0 \leq j < n_\theta$, in order to build a model of the robot for each possible orientation, as shown in Fig. 10. The variable j is used as a layer index.

In terms of morphological operations, we consider two structuring elements, \mathcal{R} and \mathcal{T} , to represent the robot and actuation models respectively. After rotating them, we get $\mathcal{R}(\theta_j)$ and $\mathcal{T}(\theta_j)$, with $0 \leq j < n_\theta$.

Using the structuring element for the robot model, we again apply morphological operations to determine the

Listing 4 Action Navigate that now checks occupied positions

```
(:action navigate
:parameters (?r - robot ?y - waypoint ?z - waypoint)
:precondition
  (and (connected ?r ?y ?z) (at ?r ?y)
        (occupied ?y) (not (occupied ?z))
  )
:effect
  (and (not (at ?r ?y)) (not (occupied ?y))
        (at ?r ?z) (occupied ?z)
  )
)
```

Fig. 8 MAP algorithm that also includes the plan-reuse phase. Inputs: MAP task (M'), Single-Agent Planner (SAP), Plan-reuse planner (R). Output: resulting plan (π) or no solution

The Multi-Agent Planning and Reuse Algorithm

Inputs: M' , SAP , R
Output: π

```

1 Forall  $\phi'_i \in \Phi'$  do  $\pi_i = \text{plan}(\Pi_i, SAP)$ 
2  $\pi_{seq} = \text{concatenate}(\pi_1, \dots, \pi_n)$  /* where  $n = |\Phi'|$  */
3 If  $\text{valid}(\pi_{seq})$ 
4 Then  $\pi = \text{parallelize}(\pi_{seq})$ 
5 Else
6    $\pi_{reuse} = \text{plan-reuse}(M', \pi_{seq}, R)$ 
7    $\pi = \text{parallelize}(\pi_{reuse})$ 
8 If  $\text{valid}(\pi)$  return  $\pi$ 
9 Else return no-solution

```

free configuration space, now for each possible quantized orientation, dilating the map using a different robot shape model for each layer. We use a circular representation for the layered orientation, where the next layer after layer $j = n_\theta - 1$ is layer with index 0.

$$\mathcal{C}^{free}(\theta_j) = \{\mathbf{p} \in \mathcal{G} \mid \mathbf{p} \notin \mathcal{M} \oplus \mathcal{R}(\theta_j)\} \quad \forall 0 \leq j < n_\theta \quad (13)$$

In order to model a robot that navigates through waypoints, we need to establish the type of connectivity between points in different layers, such as it is equivalent to the type of motion the robot actually has. As an example, using the connectivity graph from Fig. 11, where one point is connected to all its neighbors in the same layer, and the respective positions in adjacent layers, is equivalent to considering an omnidirectional model of navigation.

Given the connectivity model, it is then possible to find all points in any layer of the configuration space that connect with the starting robot location \mathbf{r}_i^0 , obtaining the navigable sets $\mathcal{L}_i(\mathbf{r}_i^0, \theta_j)$, for $0 \leq j \leq n_\theta$ layers. The initial orientation is still coded on \mathbf{r}_i^0 , and the navigable set




Fig. 9 Environment and robot models used to test the extended approach to any-shape robots

$\mathcal{L}_i(\mathbf{r}_i^0, \theta_j)$, represents layer j of the navigation space, which has orientation θ_j . We show on Fig. 12 different layers of the navigation space for two different robots.

For the any-shape robots, a multi-layer representation is used to determine the Actuation Map, representing different orientations. However, in terms of accomplishing goals, we assume it is irrelevant the orientation from which a robot actuates on a waypoint position.

Therefore, while on the rotation-invariant scenario the domain was discretized in a series of 2D waypoints, for the any-shape case there are two types of waypoints: the 3D waypoints representing (x, y, θ) position, and the 2D waypoints representing (x, y) positions invariant to orientation.

The navigability graph now becomes a graph of 3D waypoints connected to each other, modeling the motion capabilities of robots in the world in terms of both rotation and translation, individually or combined, as exemplified for different orientation layers on Fig. 13.


We then use a second dilation operation to the navigable space in each layer to obtain the actuation space for each orientation. The structuring element for this second operation is the one that models the actuation capabilities, \mathcal{T} , which dilates the space according with the actuation model. If instead the structuring element \mathcal{R} is used again, that would be equivalent to assuming an actuating ability completely coincident with the entire footprint. Thus the actuation space for each layer is given by

$$\mathcal{A}(\mathbf{r}_i^0, \theta_j) = \mathcal{L}_i(\mathbf{r}_i^0, \theta_j) \oplus \mathcal{T}(\theta_j) \quad (14)$$

The actuation space gives the actuation capabilities for each orientation for a given robot shape and starting position. So, if a point belongs to $\mathcal{A}(\mathbf{r}_i^0, \theta_j)$, then it can be actuated by the robot. We show in Fig. 12 the navigable and actuation spaces for different layers, given the robots and map shown in Fig. 9.

After determining the actuation space for each layer, we can obtain the overall actuation map in a rotation-invariant

Fig. 10 Example of an image representing the robot footprint, rotated for three different angles, and used as structuring element in the morphological operations applied to the respective orientation layers; robot center shown in red



representation by projecting the multiple layers into one single 2D image.

$$\mathcal{P}(\mathbf{r}_i^0) = \bigcup_{\theta_j} \mathcal{A}(\mathbf{r}_i^0, \theta_j) \quad (15)$$

$\mathcal{P}(\mathbf{r}_i^0)$ has the same kind of representation we had with the circular robot, where the actuation map is a single 2D image not depending on the orientation.

The actuation graph is now a graph of 3D waypoints connected to 2D waypoints, representing the actuation of a rotation-independent position in the projected 2D actuation map, from a 3D robot waypoint location, also shown in Fig. 13. The predicates on the PDDL problem are represented as follows:

- Connected (robot, 3Dwaypoint, 3Dwaypoint)
- Actuable (robot, 3Dwaypoint, 2Dwaypoint)

For each 2D waypoint in the circular robot scenario, there is now n_θ 3D waypoints in the same (x, y) position, representing the different orientations a robot can have on the same 2D waypoint. As we show in Fig. 14, the two graphs are constructed independently of the initial position, allowing very easily to change the starting location of any robot and solve a different instance of the same problem. Thus, there were no modifications in the modeling of the

PDDL problem. The 3D to 2D representation is transparent to the planning process.

The navigate action moves through 3D waypoints, and the actuate action makes 2D waypoints have the actuated predicate. The list of goals to solve the problem is still given by a list of 2D waypoints that cover all the space. Thus for the same map, the coverage problem is still the same in terms of goal waypoints and its modelization, but now we plan for robots to move through the environment and actuate goal positions from some planned orientation. The PDDL domain did not need any further modifications.

If we project the multiple layers of the graphs in a 2D image, we can analyze which waypoints are navigable in terms of the robot motion, and which ones are only feasible through an actuation action. As we show in Fig. 15, some of the waypoints are not feasible by any of the robots, and all the feasible waypoints lie inside the Actuation Space (grey region of the images).

9 Experiments and Results

In this Section we show the results of the experiments that were designed to test the impact of the preprocessing on two different versions of our algorithm, MAP. First, on the

Fig. 11 Three adjacent layers of the discretized orientation, showing in blue the neighbor points of a central orange dot, representing the connectivity/motion model


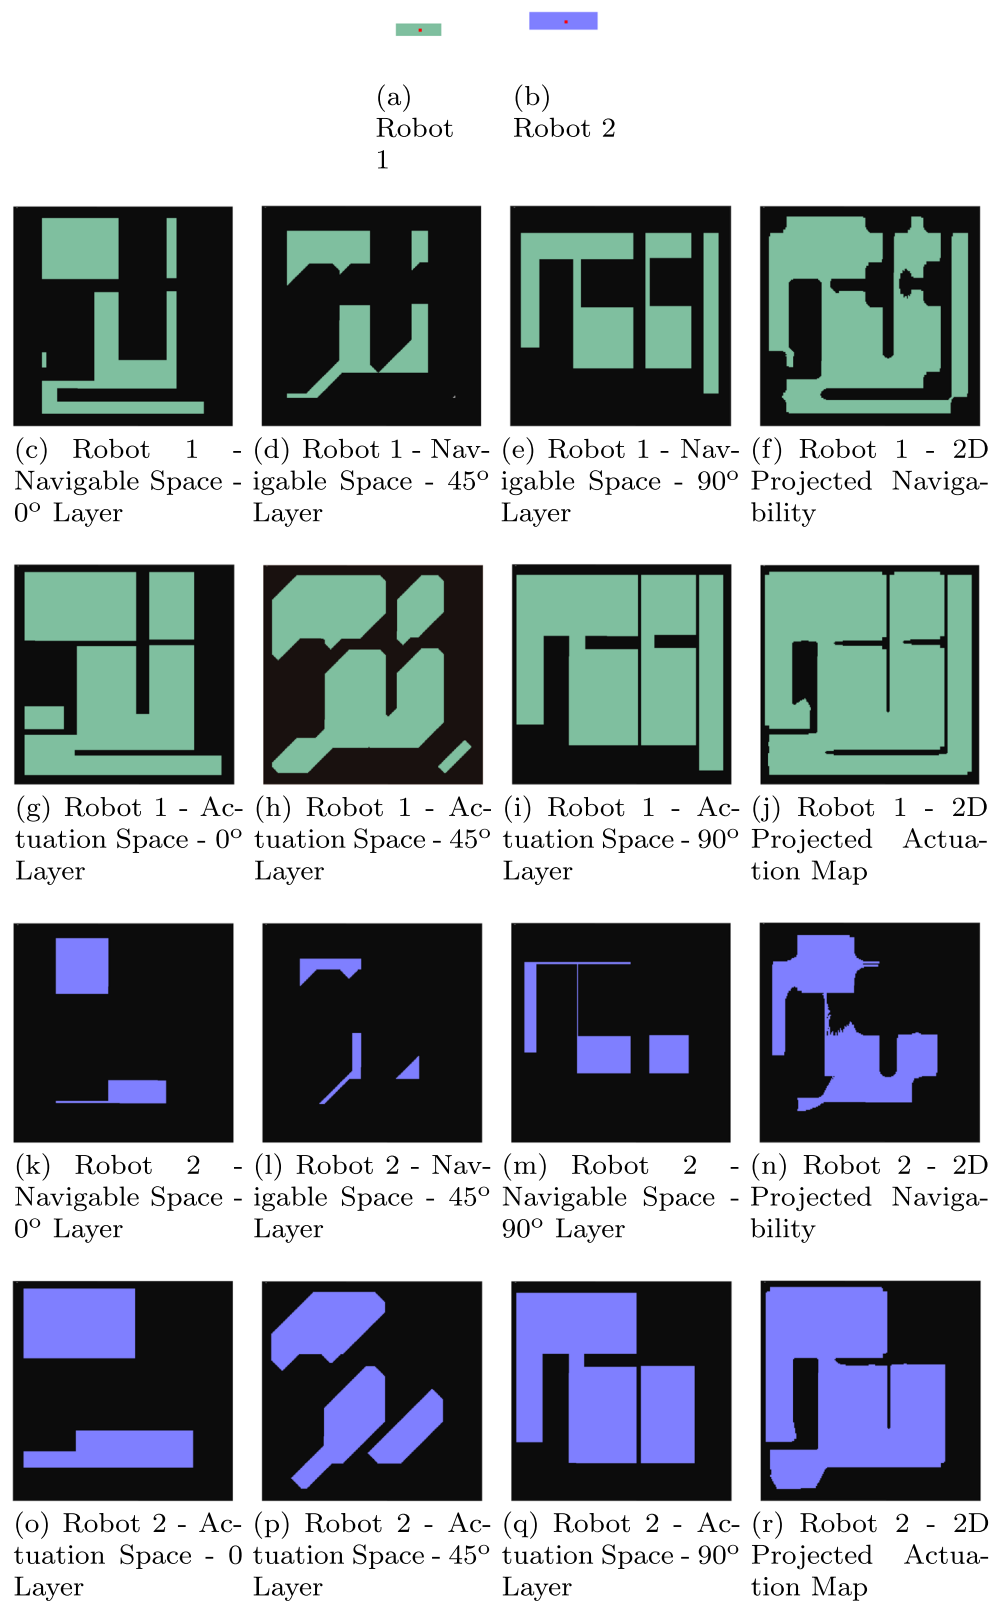


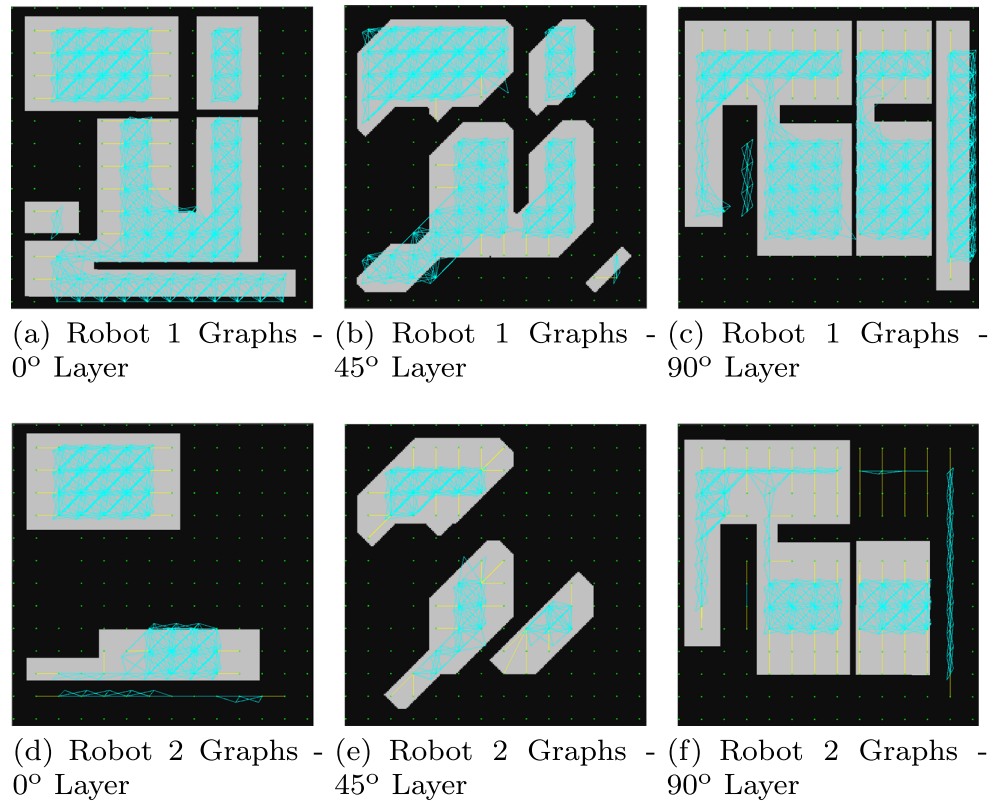
Fig. 12 Navigable and Actuation Space for 2 non-circular robots with different sizes, for the scenario shown in Fig. 9



following Section we describe the five scenarios designed to run the experiments. Then, on Section 9.2 the experiments on the coverage problem are analyzed. These results were

partially included on the previous version of the paper [29]. Finally, on Section 9.3 we show the results on the coverage problem including collision detection.

Fig. 13 The connected and actuable graphs shown in blue and yellow, respectively; as shown for each layer, the yellow actuation graph connects 3D waypoints to the original 2D green waypoints, and the blue connectivity graph connects 3D waypoints not only to neighbors in the same layer, but also in adjacent layers



9.1 Simulation Description

Here we describe in detail the scenarios used for running the experiments. We designed five different scenarios, shown in Fig. 16, each one with two levels of waypoint density (H, the higher, and L, the lower density). The scenarios are designed for circular robots except for the last one (called *Rooms*) that is designed for any-shape robots.

- Mutual Exclusive: three wide parallel horizontal halls, connected between them by two narrow vertical halls;

3 robots move within the horizontal sections, one in each, and their actuation reachabilities are mutually exclusive.

- Maze: maze-like scenario with narrow halls and passages with different sizes, resulting in bigger robots not reaching some parts of the maze, or needing to traverse bigger paths to arrive to the same locations as smaller robots.
- Corridor: four wide sections with openings of different sizes connecting them; the opening decreases from the top to the bottom, with all 4 robots being able to actuate

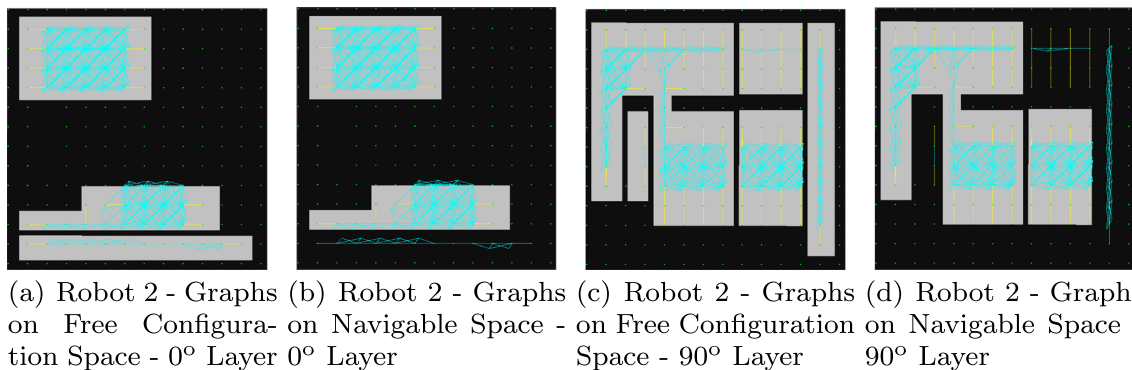



Fig. 14 The discretized graphs constructed are independent of the initial robot positions, allowing to run the problem from different initial positions; the white regions (navigable space, dependent on initial

position) are covered by the graphs, but some black regions (if grey in the respective configuration space, independent of initial position) are also covered by the constructed graphs

Fig. 15 All goal waypoints are shown as spheres on top of the Actuation Map: green represent unfeasible waypoints, in red the ones covered by the connected graph, and in blue the ones only covered by the actuatable graph; for the smaller robot 1, the two graphs are the same



- in the top region, to only 1 being able to reach the bottom.
- Extremities: wide open section with three halls departing to different directions, where all 4 robots actuate; at the end of each hall there is a room that can be accessed through an opening, with only one robot reaching the extremity connected with the smallest opening, to three reaching the one connected with the biggest opening.

- Rooms: simple floorplan environment with some room-like spaces connected through passages of different sizes as well, used to test the non-circular robot case where they can traverse the passages using only certain orientations.

Furthermore, in Table 1 we present the size of each map image, and the number agents and feasible and unfeasible goals for each scenario. We present in Table 1 the grid size

Fig. 16 Maps of the five scenarios used in the experiments. Grey regions represent out-of-reach regions which cannot contain goal waypoints. They are unfeasible for all the robots. Robots are represented with blue circles positioned in the region of their starting position

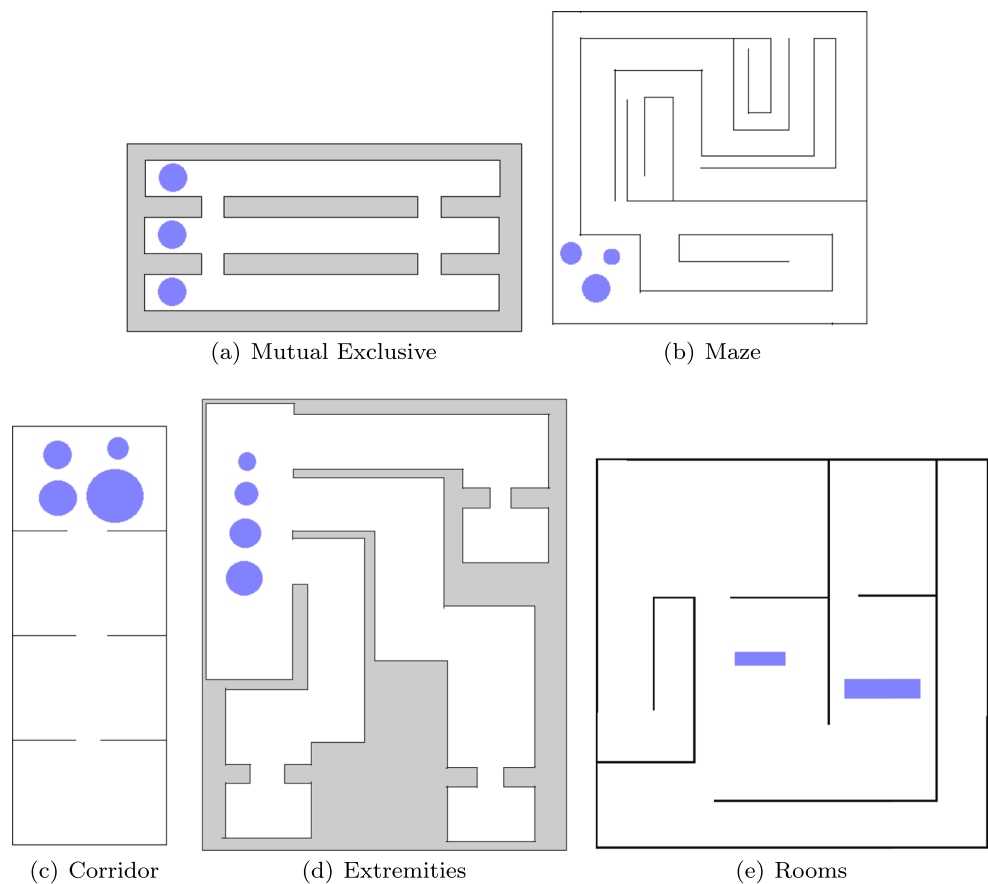


Table 1 Number of agents, feasible and unfeasible goals and respective grid size for each problem

Scenario	Agents	Feasible	Unfeasible	Grid size
CorridorH	4	819	118	49x19
CorridorL	4	384	92	33x13
ExtremeH	4	1993	1325	51x63
ExtremeL	4	896	589	34x42
MutExH	3	499	513	45x21
MutExL	3	223	242	30x14
MazeH	3	1389	154	38x38
MazeL	3	672	100	25x25
Rooms2r	2	192	52	13x13
Rooms10r	10	192	52	13x13
Extreme10r	10	1442	589	34x42
Maze10r	10	672	100	25x25
RoomsH*	2	835	182	28x28
RoomsL*	2	131	61	13x13

Scenarios are designed for circular robots except for those marked with (*), where robots are any-shape

in terms of the down-sampled grid of waypoints. The original image had a pixel size approximately 10 times bigger, with a pixel resolution corresponding to 10cm. Therefore, the maps we tested represent environments with a size always bigger than 300 square meters.

We have generated two problems per scenario, one of them with low density of waypoints (which we identify as L in tables) and the other one with a higher density of waypoints (H). We have also designed versions of *Maze*, *Extremities* and *Rooms* for 10 robots in order to test the behavior of the planners in crowded scenarios. *Rooms2r* is a similar version of *RoomsL* but for circular robots.

For the experiments on this Section, the actuation model is always considered to be equal to the robot footprint. The Actuation Map determination was developed in C++.

9.2 Experiments on the Coverage Problem

In this Section, we show some experiments that test the impact of the preprocessing in our approach. As it was previously said, we have modeled five different scenarios that include up to four agents with different sizes, and thus different actuation capabilities. Planning results are shown using as metrics the time in seconds, the length of the resulting plan and the makespan. In non-temporal domains, we refer as makespan the length of the parallel plan (number of execution steps, where several actions can be executed at the same execution step). Given that we are dealing with MAP tasks that almost have no interactions, it is expected that agents can execute their actions in parallel whenever possible.

Four different configurations of our approach have been set up:

- MAP-LB-EC with estimated-cost information (EC). EC refers to the configuration that combines AMs and MAP using the Load-Balance strategy (LB).
- MAP-BC-EC with estimated-cost information (EC), also combining AMs and MAP but using instead the Best-cost strategy (BC).
- MAP-LB, same as before but without EC information.
- MAP-BC same as before but without EC information.

As it was mentioned in Section 5.3, the LB strategy helps to minimize the *makespan* metric. The BC strategy focuses on minimizing the *plan length* metric. We also run the problems without the preprocessing stage in order to evaluate our impact in terms of computation time and plan quality.

Furthermore, the following state-of-the-art planners have been chosen as a comparison baseline:

- LAMA [30], centralized planer and winner of IPC 2011.
- YAHSP [36], a greedy centralized planner and winner of the Agile track of IPC 2014.
- ADP [9], a multi-agent planner that performs a decomposition of the domain by following three indicators: (1) dependencies are reduced, (2) goals can be achieved independently and (3) coordination between agents is minimized. The aim is to find the most suitable *agentification* i.e. identify the agents, in order to divide the problem among the agents.
- SIW [26], a multi-agent planner that factorizes the problem into subproblems solving one atomic goal at a time until all atomic goals are achieved jointly.
- CMAP [3], a multi-agent planner that employs a centralized approach to solve the problem.

The three multi-agent planners that have been chosen participated on the 1st Competition of Distributed and Multi-agent Planners (CoDMAP²) and obtained good results on the final classification.

Neither of these five planners perform a goal allocation phase separated from the planning process. Thus, we had to test them using equivalent PDDL problems that do not contain unfeasible goals. Also, in order to fairly compare the results of the makespan metric, we had to apply our parallelization algorithm to the resulting plans of ADP and SIW, as they only return the sequential plan.

Before discussing the results on the tables we need to clarify that a maximum of two hours was given to each planner to solve each scenario. YAHSP results do not appear in the tables because it could not solve any of the scenarios.

²<http://agents.fel.cvut.cz/codmap/>

The maximum time spent on the preprocessing for any scenario was 170 milliseconds, for the Extremities problem with 4 robots. We included the preprocessing times (to generate the AMs) in the GA column of Table 2, and in the total time in Tables 3 and 6. Hardware used for running the planner was IntelXeon 3,4GHz QuadCore 32GB RAM. AMs were computed using a 2.5GHz DualCore 6GB RAM. Table 2 is shown to prove the remarkable impact that information from Actuation Maps (AMs) has in combination with MAP. Goal assignment (GA) times in Table 2 are minimal (MAP-LB-EC) in comparison with the ones of MAP-LB, where it needs to compute the relaxed plans for every goal-agent pair. Even though the individual planning and parallelization time for MAP-LB-EC is similar to MAP-LB, the time gains in GA completely dominate the overall planning time. Before running any problem, MAP performs a MAP compilation of the original problem to generate each agent's individual problem after goals are assigned (M' task). Usually this transformation takes seconds and it is included in GA time. However, we observed that given the size and complexity of any-shape scenarios (*RoomsH* and *RoomsL*), the compilation time increases considerably and becomes more than half of the time spent on solving the task. This phenomenon is marked with + in column Total time from Table 2.

Regarding time results in Table 3, MAP-LB-EC is generally faster if all total times are summed up except in Maze. Also, the impact of combining information from AMs with MAP can be easily appreciated if columns from MAP-LB-EC and MAP-LB are compared. The same

happens with BC configurations. Our two configurations MAP-LB-EC and MAP-BC-EC solved every problem. There is an exception in *RoomsH*, in which the parallel plan of both solutions could not be obtained in the remaining time before the 7200 seconds were reached. In general, the easiest scenario to be solved using planning is the Mutual Exclusive (*MutExH*, *MutExL*) because it is designed for each robot to traverse a mutual exclusive subset of waypoints. This is the reason why time results are very similar among all planners except for MAP-LB and MAP-BC where the planner needs to compute the relaxed plans for each pair robot-goal. However, CMAP had some trouble during planning in the high density scenario. The circular robot version of *Rooms* (*Rooms2r*) is also very easy to solve, even though the number of goals is higher than the any-shape version (*RoomsL*). If times from *Rooms2r* and *RoomsL* are compared, the complexity of just changing from circular to any-shape robots can be empirically appreciated. ADP reached the memory limit in *Maze* when planning the solutions before the two hours limit. Even though ADP is a multi-agent planner, the effort of computing plans in big-size environments when all goals are assigned to all agents is very big. LAMA has the same issue as ADP because of its centralized approach (*Maze*, *Extremities*, *RoomsH*). From the set of planners that we chose to compare our approach, SIW is the one that obtains the best results.

Table 4 shows the results regarding the plans' length and Table 5 the results regarding makespan. We have used the words *timeout* to indicate that a planner consumed the allotted time and could not return a solution, *memlimit* to

Table 2 Detailed time results in seconds for the MAP algorithm using the Load Balance strategy with and without estimated cost information

Scenario	Time (s)							
	MAP-LB-EC				MAP-LB			
	Total	GA	Planning	Parallel	Total	GA	Planning	Parallel
CorridorH	88.07	0.90	68.34	18.83	1748.38	1672.46	58.83	17.09
CorridorL	13.37	0.29	10.54	2.54	484.48	179.19	303.28	2.01
ExtremH	639.54	4.53	427.91	207.10	timeout			
ExtremL	86.78	1.22	64.11	21.45	5491.00	5377.80	84.64	28.56
MutExH	10.59	0.44	7.36	2.79	1276.18	1265.62	7.78	2.78
MutExL	2.10	0.11	1.64	0.35	103.40	100.72	2.31	0.37
MazeH	1200.37	0.69	944.21	255.47	timeout			
MazeL	1179.21	0.24	1152.71	26.26	882.93	603.78	240.3	38.85
Rooms2r	0.19	0.05	0.05	0.09	8.31	5.70	1.85	0.76
Rooms10r	4.70	0.19	4.17	0.34	40.54	36.34	3.84	0.36
Extreme10r	93.79	2.97	68.89	21.93	timeout			
Maze10r	36.37	0.81	28.74	6.82	2121.11	2070.49	39.6	11.02
RoomsH	6180.37 ⁺	9.85	1022.19	1922.87	timeout			
RoomsL	448.63 ⁺	2.31	2.54	95.05	timeout			

From left to right total time, goal assignment time, individual planning time and parallelization time. Symbol + indicates that the MAP compilation time is very high

Table 3 Total time results in seconds

Scenario	Total Time (s)							
	MAP-LB-EC	MAP-LB	MAP-BC-EC	MAP-BC	SIW	ADP	CMA	LAMA
CorridorH	88.07	1748.38	304.91	1791.85	129.71	484.48	1761.82	95.45
CorridorL	13.37	484.48	33.78	203.00	10.98	85.97	187.57	22.74
ExtremH	639.54	timeout	642.04	timeout	1923.32*	439.58*	timeout	timeout
ExtremL	86.78	5491.00	82.51	5547.00	156.71	402.15	91.24	72.07
MutExH	10.59	1276.18	10.55	1277.47	11.65	6.15	1277.89	6.93
MutExL	2.1	103.40	2.09	97.27	0.81	0.89	96.38	1.06
MazeH	1200.37	timeout	2718.96	timeout	429.87	memlimit	2575.24	2005.42
MazeL	1179.21	882.93	161.03	1213.26	37.92	memlimit	timeout	334.8
Rooms2r	0.19	8.31	2.64	8.42	0.77	2.24	7.32	1.54
Rooms10r	4.7	40.54	4.18	36.06	2.26	2.04	34.21	2.55
Extreme10r	93.79	timeout	98.44	timeout	258.72	169.68	timeout	76.56
Maze10r	36.37	2121.11	29.63	2110.11	61.95	20.04	2099.74	29.7
RoomsH	6180.37*	timeout	6120.28*	timeout	timeout	timeout	timeout	timeout
RoomsL	448.63	timeout	447.56	timeout	286.34	140.1	timeout	132.02

From left to right MAP with estimated-cost information in Load-balance (LB-EC); MAP without estimated cost information in LB; MAP with estimated cost information in Best-cost (BC-EC); MAP without estimated cost information in BC; ADP, SIW and CMA are other multi-agent planners and LAMA is a centralized planner. Symbol * indicates that the planner solved the problem but parallelization over-passed the allotted time (7200s)

Bold values represent which is the configuration that obtained the best score per scenario

indicate that the planner's memory limit was reached before timeout and *parallel* to indicate that the planner solved the problem but the parallelization algorithm could not return a solution in the remaining time to reach 7200s (SIW, ADP in *ExtremH*; MAP-LB-EC, MAP-BC-EC in *RoomsH*).

The best configurations overall regarding plan length are MAP-BC-EC and SIW. Moreover, MAP-LB-EC configuration is generally the best for reducing makespan. Configurations MAP-LB or SIW also obtain good results in specific scenarios. This issue can be explained by the discretization

Table 4 Plan length: from left to right MAP with estimated-cost information in Load-balance (MAP-LB-EC); MAP without estimated cost information in LB; MAP with estimated-cost in Best-Cost (MAP-BC-EC); MAP without estimated cost information in BC; SIW, ADP, CMA and LAMA

	Plan length							
	MAP-LB-EC	MAP-LB	MAP-BC-EC	MAP-BC	SIW	ADP	CMA	LAMA
CorridorH	1511	1512	1653	1556	1543	3545	1541	1471
CorridorL	727	692	784	791	699	627	746	748
ExtremH	3830	timeout	3830	timeout	3580*	8687*	timeout	timeout
ExtremL	1715	1850	1715	1786	1627	2848	1659	1546
MutExH	658	658	658	658	758	773	658	658
MutExL	301	301	301	301	306	302	301	301
MazeH	3358	timeout	3004	timeout	2570	memlimit	1353	2686
MazeL	1599	1387	1434	1441	1236	memlimit	timeout	1345
Rooms2r	142	333	313	315	303	398	302	319
Rooms10r	292	344	287	286	266	318	268	268
Extreme10r	1604	timeout	1627	timeout	1534	2692	timeout	1612
Maze10r	1480	1852	1308	1261	1202	1452	1224	1253
RoomsH	1403*	timeout	1337*	timeout	timeout	timeout	timeout	timeout
RoomsL	370	timeout	336	timeout	366	627	timeout	370

Symbol * indicates that the planner solved the problem but parallelization over-passed the allotted time (7200s)

Bold values represent which is the configuration that obtained the best score per scenario

Table 5 Makespan: from left to right MAP with estimated-cost information in Load-balance (MAP-LB-EC); MAP without estimated cost information in LB; MAP with estimated-cost in Best-Cost (MAP-BC-EC); MAP without estimated cost information in BC; SIW, ADP, CMAP and LAMA

	Makespan							
	MAP-LB-EC	MAP-LB	MAP-BC-EC	MAP-BC	SIW	ADP	CMAP	LAMA
CorridorH	609	583	1353	862	973	1699	717	698
CorridorL	298	256	657	601	265	1073	444	452
ExtremH	905	timeout	905	timeout	parallel	parallel	timeout	timeout
ExtremL	376	702	376	1091	450	1378	564	484
MutExH	117	117	117	117	149	155	117	117
MutExL	58	58	58	58	64	58	58	58
MazeH	1631	timeout	1631	timeout	1941	memlimit	timeout	1369
MazeL	696	993	1035	1384	837	memlimit	1288	1217
Rooms2r	22	262	264	271	214	342	212	237
Rooms10r	47	55	57	54	59	73	67	67
Extreme10r	258	timeout	258	timeout	214	816	timeout	552
Maze10r	246	367	213	246	212	288	221	243
RoomsH	parallel	timeout	parallel	timeout	timeout	timeout	timeout	timeout
RoomsL	240	timeout	242	timeout	180	510	timeout	177

Bold values represent which is the configuration that obtained the best score per scenario

errors from Eq. 9, which are greater when the downsampling rate is bigger. When allocating goals, the estimation costs are the only guide for the MAP algorithm. The consequence of having slightly inaccurate cost estimates results in the allocation of some goals to different agents than the ones that the estimated costs from the relaxation of plans would suggest. However, this issue does not have a big impact on makespan and plan length results.

From the set of planners chosen to compare our approach, SIW obtains the best performance on time, plan length and makespan. SIW is able to solve most of the scenarios due to its serialization of goals. The importance of factorizing a MAP problem is a conclusion that can be extracted after observing Tables 4 and 5, as the planners that do not perform factorization (LAMA, ADP, CMAP, YAHSP) have to solve bigger and more complex tasks.

Regarding our configurations, MAP-BC-EC and MAP-LB-EC perform better in general than equivalent configurations without estimation costs. On the other hand, the lower the number of agents used to plan, the harder the planning task. Total time in BC is usually worse than in LB configurations on scenarios with higher density of waypoints and multiple robots to plan (*CorridorH*, *MazeH*). However, when scenarios are narrowed such as *Extremities* or *Maze*, the opposite effect can be given. When using less agents, the task is easier to solve.

9.3 Experiments Detecting Potential Collisions

In this Section we show the results obtained on the same scenarios as in the previous Section but using instead the PDDL domain that detects collisions described in Section 7.

In Tables 6 and 7, we refer to MAP&R-LB-EC as running the Fig. 8 using the LB strategy. MAP&R-BC-EC runs the BC strategy instead. We have also compared our approach against the same set of planners as in Section 9.2. The maximum time for each planner to solve each scenario is two hours. YAHSP results are not shown in the tables as it was not able to solve any problem.

The aim of this experiment is to analyze the impact of detecting and fixing collisions on makespan and time metrics. Plan length is not relevant on this experiment, as the difficulty relies on the planner's ability to manage several agents and collision avoidance at the same time. That is a feature that directly affects the makespan result.

Regarding time results in Table 6, it can be seen that the number of problems solved decreases considerably. Also, time results have increased in all planners. This is due to the collision avoidance effect. On one hand, centralized approaches can deal with it more easily, as the master agent has the whole control of the agents. However, it is still facing the same issue as in the previous experiments: the tasks are harder to solve and now the search space is bigger.

Our approach is halfway between the centralized and the distributed approach. The first part of our algorithm is distributed while the plan-reuse phase is centralized. Thus, the success of our algorithm depends on the number of collisions and the difficulty of solving them. LPG-ADAPT focus first on reutilizing the greater possible number of the actions from the invalid plan. When a collision is detected, LPG-ADAPT will search for a valid action on the part of the search space that is closer to the invalid action and its current planning state. This causes LPG-ADAPT to

Table 6 Time in seconds from left to right MAP with estimated-cost information in Load-balance (MAP&R-LB-EC); MAP with estimated-cost in Best-Cost (MAP&R-BC-EC); SIW, ADP, CMAP and LAMA

	Total Time (s)					
	MAP&R-LB-EC	MAP&R-BC-EC	SIW	ADP	CMAP	LAMA
CorridorH	timeout	timeout	263.82	timeout	timeout	293.11
CorridorL	73.57	81.51	50.64	193.54	202.16	53.24
ExtremH	timeout	timeout	timeout	timeout	timeout	timeout
ExtremL	timeout	timeout	752.5	timeout	timeout	timeout
MutExH	timeout	timeout	7.61	9.01	1328.77	70.08
MutExL	3.25	3.35	1.35	1.06	104.45	3.25
MazeH	timeout	timeout	986.17	memlimit	timeout	timeout
MazeL	timeout	timeout	67.34	memlimit	945.04	timeout
Rooms2r	4.78	4.71	1.5	5.85	8.47	2.24
Rooms10r	7.11	7.12	4.86	3.51	40.42	9.62
Extreme10r	timeout	timeout	timeout	timeout	timeout	timeout
Maze10r	406.75	timeout	466.03	memlimit	2196.02	560.40
RoomsH	timeout	timeout	timeout	timeout	timeout	timeout
RoomsL	3406.21	3092.80	timeout	timeout	timeout	timeout

iteratively explore the search space starting from a very concrete section. The exploration distance will be increased as long as the valid action is still not found. This approach works well if the collision requires a small change to be fixed i.e. it only affects to a couple of navigation steps - the solution can be found near the search space of the action and current state. But if the way to avoid the collision affects to a bigger part of the plan i.e. robots have to move back several waypoints and change path directions; LPG-ADAPT might get lost on the search space, as it will try to search first on the space closer to the invalid action and the solution might be far away from there. Thus, *timeout* will be reached before a solution is found. Scenarios not solved by our approaches

on Table 6 fail for that reason. On the other hand, Multi-Agent centralized approaches as SIW solve more problems. Next paragraphs contain a discussion on this aspect. We analyze why this particular situation is given with SIW even though MAP centralized approaches are generally worse in performance on big scenarios.

SIW only solves one atomic goal at a time (serialization), which means that goals are not assigned to agents in the first step of the algorithm. The process is interleaved with search. Thus, only one estimation is computed per iteration and current positions of the agents are updated after each goal is reached. Agents work individually but coordination (and thus, collisions) are checked after each iteration. The

Table 7 Makespan: from left to right MAP with estimated-cost information in Load-balance (MAP&R-LB-EC); MAP with estimated-cost in Best-Cost (MAP&R-BC-EC); SIW, ADP, CMAP and LAMA

	Makespan					
	MAP&R-LB-EC	MAP&R-BC-EC	SIW	ADP	CMAP	LAMA
CorridorH	timeout	timeout	1289	timeout	timeout	847
CorridorL	801	806	490	1378	452	424
ExtremH	timeout	timeout	timeout	timeout	timeout	timeout
ExtremL	timeout	timeout	704	timeout	timeout	timeout
MutExH	timeout	timeout	158	211	131	135
MutExL	87	87	80	80	101	95
MazeH	timeout	timeout	1061	memlimit	timeout	timeout
MazeL	timeout	1481	758	memlimit	1139	timeout
Rooms2r	293	293	189	465	228	250
Rooms10r	61	61	61	107	77	54
Extreme10r	timeout	timeout	timeout	timeout	timeout	timeout
Maze10r	293	timeout	432	memlimit	221	282
RoomsH	timeout	timeout	timeout	timeout	timeout	timeout
RoomsL	250	250	timeout	timeout	timeout	timeout

centralized approach followed by SIW is very efficient on the coverage problem. Collisions can be fully avoided because of solving first only one goal at a time and then updating robots' positions. Thus, the algorithm obtains good results in number of problems solved and time. However, when the size of the problem increases, as in the any-shape scenarios, SIW has more difficulties to solve the problem in time. This scenario penalizes SIW because the search space is huge in comparison with circular robot scenarios. Planning one goal at a time following a centralized approach now becomes a worse choice. SIW has to usually deal with the following situation when collisions are given: a set of goals has been reached and the next goal on the list cannot be achieved unless the previous part of the plan is partially modified. In any-shape scenarios this aspect takes more time to fix given the size of the search space.

Regarding our approach, even though we obtain estimation costs through AMs, they are not as effective in guidance as SIW's serialization. The estimation of costs is provided to our approach at the beginning. If agents have to modify their route due to collisions, their estimations and assignment of goals might not be as useful as in the beginning. It can even penalize the agent's performance. Also, if an agent needs to change its route several times, it could mean that the original assignment of goals is completely useless. Also, plan reuse planners are not efficiently prepared to perform an extensive search. They would rather prefer to reutilize actions from previous plans, which in the coverage problem results in generating redundant actions around the planning task. We have tried different features of LPG-ADAPT (low memory, speed mode, increasing different fixed constants...) to check if its performance could be improved but neither of them helped.

Regarding makespan results from Table 7, MAP&R-LB-EC and SIW are the two planners that obtain the best results - they also solve most of the problems. Although the *RoomsL* scenario might seem easier to solve by just looking at Table 1. However, as it works for any-shape robots, the grid of waypoints is bigger and harder to navigate from the planning point of view. The search space is very big and thus centralized approaches are especially penalized. The reason of failing on the *Extremities* and *Maze* scenarios is due to the changes on the robots' paths caused by the collision avoidance or the topology of the scenario. Those scenarios contain narrower areas and large halls where only some robots can reach the end. Thus, robots might spend a lot of time looking for the correct path while at the same time avoiding the rest of the agents.

As a final conclusion, we would like to discuss the overall performance of our contribution. MAP-LB-EC and MAP-BC-EC clearly complement each other on the set of proposed scenarios. This is an advantage, as the algorithm can be easily adapted to different situations and environments. It is

true that we lost some performance on collisions, but we have empirically shown that it is also related to the topology of the scenario and the coverage problem itself. The unexpected advantage of SIW in the experiments detecting collisions has also been addressed and analyzed. The serialization of goals and the nature of the coverage problem where interactions are given occasionally, makes any centralized (Multi-agent or Single-agent) planner to behave well. LAMA is closer to SIW and our approaches in that sense, and it is not even a MAP algorithm. We also want to put in value the scalability of our approach. Through the experiments we have proved that we can successfully deal with different topologies, number of agents, agent's orientation, huge planning tasks, unfeasible goals, independent goal-assignment and pre-processed estimation costs. State-of-the-art planners are not used to satisfy all these features at the same time.

10 Related Work

Multi-Agent Planning (MAP) is an active topic within the planning community as shown by the organization of the 1st CoDMAP³ and the wide range of planners that participated [32]. The planners vary from fully distributed to centralized among other features. Our approach uses first a distributed planner and if plan reuse is needed, a centralized phase.

A MAP approach that uses a preprocessing step is the automated agent decomposition for multi-robot task planning [10]. In that work there is a preprocessing step, prior to actual search, that exploits decompositions of the problem in domains with a lower level of interaction, boosting the final performance. ADP [9] is also related to that work. Our approach factorizes the problem regarding goals and agents involved, creating independent subtasks for each agent before starting to plan.

The methodology that uses morphological operations in order to build the AMs was previously introduced by [27]. Moreover, it has recently been shown that similar transformation can be used to obtain Actuation and Visibility Maps for any-shape robots as well [28].

Earlier in 1997, Ambite and Knoblock [1] worked in a post-processing technique which rewrites some planning rules and local search techniques to make efficient the planning process. Thus, they obtained a low-cost plan for problems that were hard to solve from the point of view of planning.

Another similar problem is the Team Orienteering Problem, where robots maximize the number of covered waypoints visited, with constraints on the total travel time of

³<http://agents.fel.cvut.cz/codmap/>

each robot. In order to solve this problem and optimize the overall planning time, an algorithm was proposed that uses a three-tier graph, interleaving the search for optimal waypoint assignment, ordering of the waypoints and also considering feasible paths between waypoints while avoiding obstacles [33]. This algorithm guarantees optimal solutions, while in our work we focus on sub-optimal planning for multi-robot teams.

Another relevant planning problem in robotics is inspection, which searches for robot paths that can perceive a set of sensing targets. Neural network approaches have been proposed [11], and again a preprocessing method was used, answering visibility queries efficiently both in 2D and 3D scenarios [21]. Here the queries ask for visibility from specific positions. In our work we preprocess the environment to find the overall actuation capabilities of the robot from any reachable position.

There are similar environments to our problem defined in previous planning domains. One it is called *VisitAll*, whose aim is to visit all the waypoints presented in the problem by just navigating through them. It was used for the first time on 7th IPC.⁴ In our domain we added an actuation action (with an associated actuation range), which was not considered in the *VisitAll* domain. Another similar problem is the Rovers domain, but in this case each agent (rover) can execute a bigger set of actions when it is placed on a waypoint. Some other examples were mentioned in the Introduction.

When robots with many degrees of freedom execute successive motions in the same environment, it usually requires many complex planning instances. By applying a preprocessing technique, it is possible to improve efficiency of path planning for those robots [22]

Finally, a very common robotic application, coverage path planning, has been widely studied [16]. However, the many cell decomposition-based strategies still do not objectively consider heterogeneity and thus we do not exactly know how it creates different feasible tasks for each robot when assigning goals.

11 Conclusions and Future Work

In this paper we showed how to combine information from Actuation Maps (AMs) with Multi-Agent Planning (MAP) to solve a multi-robot path planning problem more efficiently skipping the computation of estimated cost during planning. We used AMs in a preprocessing step to determine the feasibility of pairs robot-goal and to extract an estimated cost. That cost is used later to avoid the computation of relaxed plans during Goal-Assignment. The environment map was discretized into a grid of waypoints.

The goals were distributed thanks to a goal-allocation algorithm and unfeasible goals identified and discarded from the planning task. Then, the planning task was factorized for each robot. They generate their individual paths that result in a maximal space coverage in terms of actuation. We also proposed a new version of the MAP algorithm that is able to fix agents' interactions after the individual planning phase. On the experiments we have designed a total of eight scenarios, seven for circular robots and one for any-shape robots, which is another contribution to the paper.

Our approach greatly reduces the GA time, and because GA is one of the bottlenecks of MAP, we were able to also reduce the overall planning time when preprocessed information was provided to the MAP algorithm. The gains in performance depend greatly on the topology of the environment and the characteristics of each robot.

Experiments also show that when solving big size multi-agent problems using planning, it is very helpful to first factorize the problem into subtasks. Thus, the total planning time will be smaller than when trying to solve the whole problem at once. Also, factorization is essential when working on problems that explicitly involve agents' interactions. Experiments on collision avoidance show the importance of task factorization and the topology of the scenario in order to successfully fix collisions.

As future work, we would like to extend the preprocessing technique to other domains and consider different -robot or agent- models. Our approach can be easily extended to path planning tasks or real-time strategy videogames. We gave some examples of the former such as surveillance tasks or search and rescue tasks. The latter domain could be interesting when designing bots that play automatically. Our approach could improve the player/bot performance when extracting information from the map to decide which goals are more relevant to achieve first. We also want to improve the performance of fixing interactions. Plan reuse works well when collisions only affect to a couple of actions. For biggest plan modifications plan-reuse is not enough.

We would also like to study in the future the possibility of using plan-reuse in order to deal with dynamic environment during plan execution. For the case of changes during plan execution, we would also need to create a technique to efficiency update the PDDL problem file according to the changes detected in the environment regarding obstacles.

Acknowledgements This work has been partially funded by FEDER/Ministerio de Ciencia, Innovación y Universidades - Agencia Estatal de Investigación/TIN2017-88476-C2-2-R and MINECO/TIN2014-55637-C2-1-R. I has been also financed by the ERDF – European Regional Development Fund through the Operational Programme for Competitiveness and Internationalisation - COMPETE 2020 Programme within project «POCI-01-0145-FEDER-006961», and by National Funds through the FCT – Fundação para a Ciência e a Tecnologia (Portuguese Foundation for Science and Technology) as part of project UID/EEA/50014/2013, and FCT grant SFRH/BD/52158/2013 through Carnegie Mellon Portugal Program.

⁴<http://www.plg.inf.uc3m.es/ipc2011-deterministic>

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

References

- Ambite, J.L., Knoblock, C.A.: Planning by rewriting: Efficiently generating high-quality plans (1997)
- Borrajó, D.: Plan sharing for multi-agent planning. In: Nissim, R., Kovacs, D.L., Brafman, R. (eds.) *Proceedings of the ICAPS'13 DMAP Workshop*, pp. 57–65 (2013)
- Borrajó, D., Fernández, S.: MAPR and CMAP. In: *Proceedings of the Competition of Distributed and Multi-Agent Planners (CoDMAP-15)*. Jerusalem. <http://agents.fel.cvut.cz/codmap/results/CoDMAP15-proceedings.pdf> (2015)
- Borrajó, D., Fernández, S.: Efficient approaches for multi-agent planning. *KAIS* **14**, 253–302 (2018)
- Borrajó, D., Veloso, M.M.: Probabilistically reusing plans in deterministic planning. In: *Inproceedings of ICAPS'12 Workshop on Heuristics and Search for Domain Independent Planning* (2012)
- Borrajó, D., Roubířková, A., Serina, I.: Progress in case-based planning. *ACM Comput Surv* **47**(2), 35:1–35:39 (2015). <https://doi.org/10.1145/2674024>
- Brafman, R.I., Domshlak, C.: On the complexity of planning for agent teams and its implications for single agent planning. *Artif. Intell.* **198**, 52–71 (2013). <https://doi.org/10.1016/j.artint.2012.08.005>
- Conitzer, V.: Comparing multiagent systems research in combinatorial auctions and voting. *AMAI* **58**(3–4), 239–259 (2010). <https://doi.org/10.1007/s10472-010-9205-y>
- Crosby, M.: Adp an agent decomposition planner. In: *Proceedings of the CoDMAP-15*, p. 4 (2015)
- Crosby, M., Rovatsos, M., Petrick, R.P.: Automated agent decomposition for classical planning. In: *ICAPS* (2013)
- Faigl, J.: Approximate solution of the multiple watchman routes problem with restricted visibility range. *IEEE Trans. Neural Netw.* **21**(10), 1668–79 (2010)
- Fikes, R.E., Nilsson, N.J.: Strips: A new approach to the application of theorem proving to problem solving. In: *Proceedings of IJCAI*, pp. 608–620. <http://dl.acm.org/citation.cfm?id=1622876.1622939> (1971)
- Foulser, D., Li, M., Yang, Q.: Theory and algorithms for plan merging. *Artif. Intell.* **57**(2–3), 143–181 (1992)
- Fox, M., Long, D.: PDDL2.1: An extension to PDDL for expressing temporal planning domains. *JAIR* **20** (2003)
- Fox, M., Gerevini, A., Long, D., Serina, I.: Plan stability: replanning versus plan repair. In: *Proceedings of ICAPS'06*, pp. 212–221 (2006)
- Galceran, E., Carreras, M.: A survey on coverage path planning for robotics. *Robot. Auton. Syst.* **61**(12), 1258–1276 (2013)
- García-Olaya, A., de la Rosa, T., Borrajó, D.: Using the relaxed plan heuristic to select goals in oversubscription planning problems. In: Lozano, J.A., Gámez, J.A., Moreno, J.A. (eds.) *Advances in Artificial Intelligence*, pp. 183–192 (2011)
- Gerkey, B.P., Mataric, M.: A formal analysis and taxonomy of task allocation in multi-robot systems. *Int. J. Robot. Res.* **23**(9), 939–954 (2004)
- Hoffmann, J., Nebel, B.: The FF planning system: Fast plan generation through heuristic search. *JAIR* **14**, 253–302 (2001)
- Howey, R., Long, D., Fox, M.: VAL: Automatic plan validation, continuous effects and mixed initiative planning using PDDL. In: *ICTAI 2004*, pp. 294–301 (2004)
- Janousek, P., Faigl, J.: Speeding up coverage queries in 3D multi-goal path planning. *ICRA* **1**, 5082–5087 (2013). <https://doi.org/10.1109/ICRA.2013.6631303>
- Kavraki, L., Latombe, J.C.: Randomized preprocessing of configuration for fast path planning. In: *Proceedings of the International Conference on Robotics and Automation*, pp. 2138–2145. IEEE (1994)
- Krulwich, B.: Planning for soft goals. In: Hendler, J. (ed.) *Artificial Intelligence Planning Systems*, pp. 289–290 (1992). <https://doi.org/10.1016/B978-0-08-049944-4.50047-1>. <https://www.sciencedirect.com/science/article/pii/B9780080499444500471>
- Luis, N., Borrajó, D.: Plan merging by reuse for multi-agent planning. In: *Proceedings of the 2nd ICAPS Distributed and Multi-Agent Planning Workshop (DMAP)* (2014)
- Mali, A.D.: Plan merging & plan reuse as satisfiability. In: *Proceedings of the 5th European Conference on Planning: Recent Advances in AI Planning. ECP'99*, pp. 84–96 (2000)
- Muise, C., Lipovetzky, N., Ramirez, M.: MAP-LAPKT: Omnipotent multi-agent planning via compilation to classical planning. In: *Competition of Distributed and Multiagent Planners*. http://www.haz.ca/papers/muise_CoDMAP15.pdf (2015)
- Pereira, T., Veloso, M., Moreira, A.: Multi-robot planning using robot-dependent reachability maps. Springer (2015)
- Pereira, T., Veloso, M., Moreira, A.P.: Visibility maps for any-shape robots. In: *IROS'16, the IEEE/RSJ International Conference on Intelligent Robots and Systems*, vol. 29, pp. 428–459 (2016)
- Pereira, T., Luis, N., Moreira, A., Borrajó, D., Veloso, M., Fernandez, S.: Heterogeneous multi-agent planning using actuation maps. In: *2018 IEEE International Conference on Autonomous Robot Systems and Competitions (ICARSC)*, pp. 219–224. <https://doi.org/10.1109/ICARSC.2018.8374186> (2018)
- Richter, S., Westphal, M.: The LAMA planner: Guiding cost-based anytime planning with landmarks. *JAIR* **39**, 127–177 (2010)
- Smith, D.E.: Choosing objectives in over-subscription planning. In: *ICAPS* (2004)
- Štolba, M., Komenda, A., Kovacs, D.L. (eds.): *Proceedings of the Competition of Distributed and Multi-Agent Planners (CoDMAP-15)* (2015)
- Thakur, D., Likhachev, M., Keller, J., Kumar, V., Dobrokhodov, V., Jones, K., Wurz, J.: Kammer I (2013) Planning for opportunistic surveillance with multiple robots. In: *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 5750–5757. IEEE
- Torreño, A., Onaindia, E., Komenda, A., Štolba, M.: Cooperative multi-agent planning: A survey. *ACM Comput. Surv.* **50**(6), 84:1–84:32 (2017). <https://doi.org/10.1145/3128584>
- Veloso, M.M., Pérez, M.A., Carbonell, J.G.: Nonlinear planning with parallel resource allocation. In: *Proceedings of the DARPA Workshop on Innovative Approaches to Planning, Scheduling, and Control*, pp. 207–212. Morgan Kaufmann (1990)
- Vidal, V.: YAHSP3 and YAHSP3-MT in the 8th international planning competition. In: *8th International Planning Competition (IPC-2014)*, pp. 64–65 (2014)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Nerea Luis is a PhD student at the Universidad Carlos III de Madrid. Her thesis is about combining multi-agent planning and plan-reuse to easily adapt to different scenarios. In 2016, she was a visitor student at Manuela Veloso's CORAL group for 4 months in Carnegie Mellon University (USA). Her current research interests are multi-agent planning, machine learning and robotics. She was awarded with the Women Techmakers scholarship by Google.

Tiago Pereira is a Ph.D. candidate in the CMU-Portugal Program, being a student in the Electrical and Computer Engineering Department both at Carnegie Mellon University and Faculty of Engineering of University of Porto. He is also a research assistant at INESC TEC. He received a B.S and M.S. in Electrical and Computer Engineering from University of Porto in 2012. His research interests include motion planning and coordination of heterogeneous multi-robot systems.

Susana Fernández is currently an Associate Professor at the Universidad Carlos III de Madrid, since 2007. She received the B.A. degree in Physics (Universidad Complutense de Madrid, 1989), and a M.Phil. degree in "Logic, Text and Information Technology" (University of Dundee, Scotland, 1991). She worked for 10 years as a software engineer in two private companies, INDRA and Eliop, S.A. She received the Ph.D. degree in Computer Science (Universidad Carlos III de Madrid, 2006). Her current research interest is in Artificial Intelligence, particularly automated planning, machine learning and multi-agent systems. She has participated in several competitive projects related to automated planning, problem solving or agents. She has published more than 20 papers.

António Moreira is currently an Associated Professor in Electrical Engineering at University of Porto, developing his research within the Robotic and Intelligent Systems Centre of INESC TEC (Centre Coordinator). He graduated with a degree in Electrical Engineering from the University of Porto in 1986. He then pursued graduate studies at the University of Porto, completing a M.Sc. degree in Electrical Engineering - Systems in 1991 and a Ph.D. degree in Electrical Engineering in 1998. From 1986 to 1998 he also worked as an assistant lecturer in the Electrical Engineering Department of the University of Porto. His main research areas are Process Control and Robotics.

Daniel Borrajo received his Ph.D. from Universidad Politécnica de Madrid (1990). Currently, he is a Professor of Computer Science at Universidad Carlos III de Madrid. He has published over 200 journal and conference papers, mainly in the fields of problem solving methods (heuristic search and automated planning) and machine learning. He has been the Program Co-chair of the International Conference of Automated Planning and Scheduling (ICAPS'13), Conference co-chair of the Symposium of Combinatorial Search (SoCS'12, SoCS'11) and ICAPS'06, and PC member (including Area chair and Senior PC) of the main conferences on Artificial Intelligence (e.g., IJCAI, AAAI, ICAPS, ...). His current research interests lie in goal reasoning, Multi-agent Planning, and Machine Learning applied to Planning.

Manuela Veloso is the Herbert A. Simon University Professor in the School of Computer Science at Carnegie Mellon University. She is the former Head of the Machine Learning Department. Currently on leave, she is now the Head of Artificial Intelligence at JP Morgan. She researches in Artificial Intelligence, Robotics, and Machine Learning. She founded and directs the CORAL research laboratory, for the study of autonomous agents that Collaborate, Observe, Reason, Act, and Learn, www.cs.cmu.edu/~coral. Professor Veloso is AAAI, Fellow, ACM Fellow, AAAS Fellow, and IEEE Fellow, and the past President of AAAI and of RoboCup. Professor Veloso and her students research a variety of autonomous robots, including mobile service robots and soccer robots. See www.cs.cmu.edu/~mmv for further information, including publications.