Scheduling Multi-robot Missions with Joint Tasks and Heterogeneous Robot Teams

Gricel Vázquez, Radu Calinescu, and Javier Cámara

Department of Computer Science, University of York, UK

Abstract. We present a work-in-progress approach to scheduling multi-robot missions comprising tasks that need to be performed by multiple robots. Our approach (1) supports the scheduling of such missions for heterogeneous robots, (2) can take into account dependability, performance and other nonfunctional requirements, and (3) guarantees compliance with mission requirements by using a combination of formal techniques to allocate the mission tasks to individual robots, and to plan the order in which each robot will execute its allocated tasks. We show the effectiveness of our approach by applying it to the scheduling of a multi-robot mission in a hospital-support application.

Keywords: Multi-robot systems · Task allocation and planning · Constraint solving · Probabilistic model checking

1 Introduction

Multi-robots systems (MRS) have the potential to perform missions that humans find too dangerous, tedious or costly. Examples of such missions include search and rescue [10], hospital and care-home support [1,3], and inspection of critical infrastructure [13]. However, scheduling MRS missions is very challenging due to the complexity of their constraints and requirements. These missions must achieve strict dependability, performance and other nonfunctional requirements, and may need to be carried out by teams of heterogeneous robots. No existing MRS-mission scheduling solution [11] can support them together. The use of probabilistic model checkers as planners is specially useful to provide behavioural, performance and safety guarantees [4,8,14]; as well as capturing, for example: the probability of succeeding with a task [9], spatial distribution of the tasks [2], multiple decompositions of tasks [12], and partial knowledge of the environment [5]. Most studies simultaneously solve the allocation of tasks and planning problems. However, they do not consider complex task dependencies that we capture (via separating these two problems), such as tasks that require more than one robot to be completed, ordered and consecutive tasks.

Our paper introduces a work-in-progress approach for the scheduling of heterogeneous-robot MRS missions comprising ordered and joint tasks, where these missions need to satisfy nonfunctional requirements such as cost minimisation. Our approach supports the *high-level scheduling* of MRS missions, i.e., we assume that the robots can navigate through their environment, avoid obstacles, etc., and we use (a) constraint solving to allocate tasks (e.g., 'R1 cleans hospital room A' and 'R3 disinfects room C') to individual robots; and (b) probabilistic model checking to decide the execution order for these tasks (e.g., 'robot R1 cleans room A, then rearranges the furniture in room D together with robot R5').

2 MRS Mission Scheduling Approach

As shown in Fig. 1, our MRS mission scheduling approach takes four inputs. First, domain experts provide a *task specification* that defines the *types of tasks* for the application domain/organisation using the MRS. This includes *atomic tasks* with their properties (mean execution time, number of robots needed, etc.), and *compound tasks*, i.e., lists of atomic and/or other compound (sub)tasks that may need to be executed in order and/or consecutively. Next, an "MRS team" of engineers provides: (i) a *world model* defining the physical layout of the environment where the MRS missions will be performed, and (ii) a *robot specification* describing the *capabilities*, initial location and other characteristics of every available robot. Each capability of a robot indicates a type of task which that robot can execute, and provides details about the performance, reliability, energy use, etc. with which the robot would execute the task. Finally, the MRS users provide a *mission specification* defining the combination of tasks that need to be performed by the available robots, at specific locations and with given timing/cost/etc. constraints and optimisation objectives.

Given these inputs, we use a two-stage approach to generate individual robot plans whose execution ensures the correct completion of the specified mission. Stage 1 of the approach uses a *constraint solver* such as the Alloy analyzer [6] to distribute the tasks of the mission among the available robots, such that all the constraints from the task specification *and* the mission specification are satisfied. This involves using a *constraint problem generator* to encode these constraints in a format that the constraint solver can use to generate feasible task allocations.

Stage 2 of the approach optimises the order in which each robot will execute its tasks. Optimal robot plans are produced for each feasible task allocation

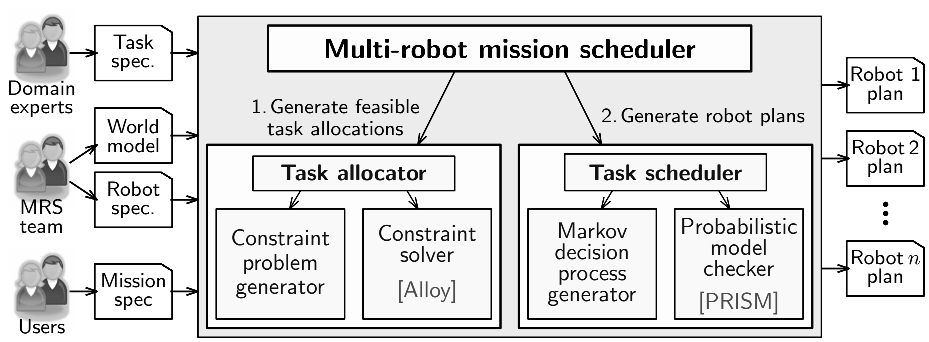

Fig. 1. Two-stage MRS mission scheduling approach

from Stage 1, and the best combination of plans across all task allocations is adopted. To generate the optimal robot plans, we use a *Markov decision process (MDP) generator* to encode the task-order optimisation as an MDP policy syn-

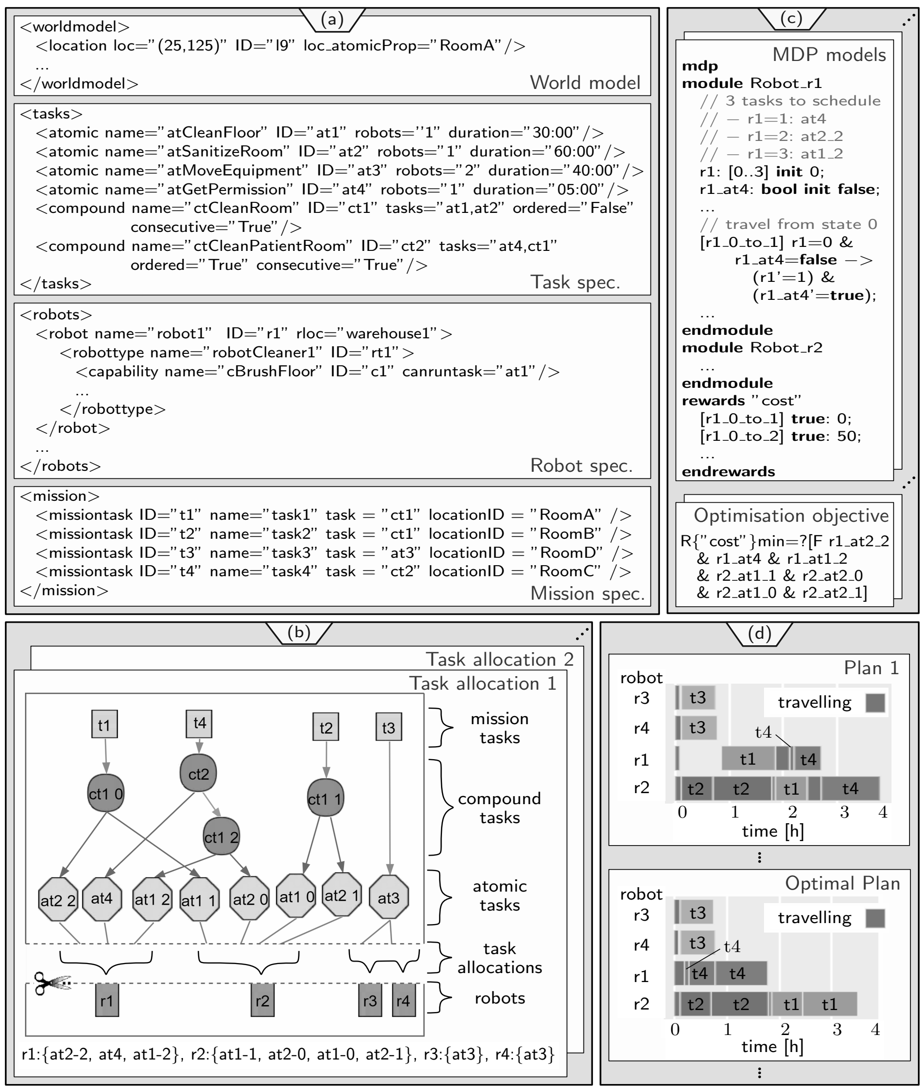

Fig. 2. Application of our MRS mission scheduling approach to a hospital case study, showing: (a) the problem specification (world model, tasks, robots and mission); (b) the Alloy-generated task allocations (the robots are shown in red at the bottom and the mission tasks in grey at the top); (c) the MDP models for each subset of robots allocated interdependent tasks; and (d) the robot plans obtained through MDP policy synthesis

thesis problem that we then solve using a probabilistic model checker such as PRISM [7]. For increased efficiency, a separate, small MDP is generated for each subset of robots that were allocated interdependent tasks, e.g., joint tasks, or tasks of a compound task with an order or sequence (`https://git.io/JGLRZ`).

3 Implementation and Case Study Summary

We developed a first version of our MRS mission scheduler using the Alloy analyser [6] for the task allocation, and the model checker PRISM [7] for the robot plan generation. Fig. 2 shows the use of our solution to schedule an MRS mission in a hospital scenario. The mission is carried out in an area comprising four rooms (A to D), and consists of four tasks (t1 to t4): cleaning empty rooms A and B (t1 and t2); moving medical equipment within room D (t3); and cleaning patient room C (t4). Room cleaning is a compound task (ct2) requiring patient permission (at4) (unless the room is empty – ct1), floor cleaning (at1), and sanitizing (at2). To move medical equipment (at3), two robots are needed. Four robots are available, two cleaner robots (r1, r2) and two pick-and-place robots (r3, r4). This information is encoded in XML (Fig. 2a) and supplied to our task allocator, which uses the Alloy Analyser as a constraint solver to create 672 feasible task allocation models. These allocation models (Fig. 2b) fulfil a set of constrains, called *facts* in Alloy language. For example: a) every atomic task is linked to a specific capability, and b) every atomic task states how many (different) robots needs to be completed

Each of the allocations are passed to the Task Scheduler which applies *transitive closure* to divide each allocation into independent robot groups (groups of robots that do not have tasks in common and do not share constrained tasks), and generates their corresponding MDP encodings (Fig. 2c). Finally, correct-by-construction robot plans are obtained through optimal MDP policy synthesis (Fig. 2d). The optimisation objective used in our hospital case study (specified as a PRISM reward property at the bottom of Fig. 2c) is the minimisation of the overall robot travelling cost. We provide a detailed description of the case study, and the specifications, models, intermediate results and robot plans from Fig. 2 in our project's GitHub repository `https://git.io/Js1Yj`.

4 Conclusions and Discussion

We introduced a new approach for the scheduling of multi-robot missions comprising joint, ordered and consecutive tasks that need to be executed by teams of heterogeneous robots. By using a combination of constraint solving and MDP policy synthesis, our approach generates correct-by-construction robot plans. In future work, we will leverage the capabilities of probabilistic model checkers to expand the range of optimisation objectives supported by our MRS mission scheduling so that they include minimising mission cost and robot-team size, maximising mission reliability, etc. Additionally, we will improve the scalability of the Alloy task allocation by (a) adding constraints that preclude the generation of permutations of the same task allocation, (b) combining it with AI techniques for a faster identification of optimal or nearly optimal task allocations, and (c) optimising the allocation of tasks considering, for instance, the spatial distance between tasks (e.g., group the tasks by capabilities in a certain area and assign them to a single robot, similar to [2]).

Understanding the computational complexity of our approach is another area of future work for the project. Analysing the complexity of the approach is non-trivial, as it depends on the configuration of the Alloy Analyser's SAT solver (MiniSat, SAT4J, ZChaff, etc.), and of the PRISM engine (MTBDD, sparse, hybrid, explicit). The time to find an optimal solution depends on these configurations, and on the size of the MDP model; which in turn depends on the number of robots, number of tasks and task dependencies. As mentioned on our GitHub page, most task dependencies are modelled within the MDP in a way that reduces the state space. In addition, the evaluation comparing the system to related solutions (e.g., [2,12]) is planned for the full-paper version of this work. Finally, we will extend our approach to support adaptation of the robot plans as they are executed, so that robot failures, mission changes, etc. are supported.

References

1. Benavidez, P., Kumar, M., Agaian, S., Jamshidi, M.: Design of a home multi-robot system for the elderly and disabled. In: SoSE'15. pp. 392–397 (2015)
2. Carreno, Y., Pairet, È., Petillot, Y., Petrick, R.P.: Task allocation strategy for heterogeneous robot teams in offshore missions. In: AAMAS'19. pp. 222–230 (2020)
3. Das, G.P., McGinnity, e.a.: A distributed task allocation algorithm for a multi-robot system in healthcare facilities. JINT **80**(1), 33–58 (2015)
4. Gavran, I., Majumdar, R., Saha, I.: Antlab: a multi-robot task server. ACM Transactions on Embedded Computing Systems (TECS) **16**(5s), 1–19 (2017)
5. Guo, M., Dimarogonas, D.V.: Multi-agent plan reconfiguration under local ltl specifications. The International Journal of Robotics Research **34**(2), 218–235 (2015)
6. Jackson, D.: Alloy: a lightweight object modelling notation. ACM Transactions on Software Engineering and Methodology (TOSEM) **11**(2), 256–290 (2002)
7. Kwiatkowska, M., Norman, G., Parker, D.: PRISM 4.0: Verification of probabilistic real-time systems. In: CAV'11. pp. 585–591 (2011)
8. Lacerda, B., Faruq, F., Parker, D., Hawes, N.: Probabilistic planning with formal performance guarantees for mobile service robots. IJRR **38**(9), 1098–1123 (2019)
9. Lacerda, B., Parker, D., Hawes, N.: Optimal policy generation for partially satisfiable co-safe ltl specifications. In: IJCA'24 (2015)
10. Pujol-Gonzalez, M., Cerquides, J., Meseguer, P., Rodríguez-Aguilar, J.A.: Efficient inter-team task allocation in RoboCup rescue. In: AAMAS. pp. 413–421 (2015)
11. Rizk, Y., Awad, M., Tunstel, E.W.: Cooperative heterogeneous multi-robot systems: A survey. ACM Computing Surveys (CSUR) **52**(2), 1–31 (2019)
12. Schillinger, P., et al.: Simultaneous task allocation and planning for temporal logic goals in heterogeneous multi-robot systems. IJRR (2018)
13. Sukkar, F., Best, G., Yoo, C., Fitch, R.: Multi-robot region-of-interest reconstruction with dec-mcts. In: ICRA. pp. 9101–9107 (2019)
14. Yu, P., Dimarogonas, D.V.: Distributed motion coordination for multi-robot systems under ltl specifications. arXiv preprint arXiv:2103.09111 (2021)