

Optimal Multi-robot Task Planning: from Synthesis to Execution (and Back)

Francesco Leofante*

RWTH Aachen University, Germany

University of Genoa, Italy

leofante@cs.rwth-aachen.de

Abstract

Integrated task planning and execution is a challenging problem with several applications in AI and robotics. In this work we consider the problem of generating and executing optimal plans for multi-robot systems under temporal and ordering constraints. More specifically, we propose an approach that unites the power of Optimization Modulo Theories with the flexibility of an on-line executive, providing optimal solutions for task planning, and runtime feedback on their execution.

1 Motivation

Lights-out manufacturing is no longer a distant dream. Modern factories are increasingly shifting towards production paradigms where activities and material flows are handled entirely by autonomous robotic systems. As the abilities and the complexity of such systems grow, providing *guarantees* and *explanations* for their behaviors becomes crucial.

The RoboCup Logistics League (RCLL) [Niemueller *et al.*, 2015] has been proposed as a simplified, yet realistic, testbed to study the above mentioned problem at a manageable scale. There, two teams of autonomous robots compete to handle the logistics of materials through several dynamic stages to produce goods in a smart factory scenario – see Fig. 1.

Successful *heuristic* approaches have been proposed to solve the planning problem underlying the RCLL – see, *e.g.*, [Hofmann *et al.*, 2016; Niemueller *et al.*, 2013]. However, an important limitation of these methods is that they cannot provide guarantees on the quality of the solutions they produce. To account for this problem, we propose to rely on the recently emerging field of Optimization Modulo Theories (OMT), where Satisfiability Modulo Theories (SMT) solving is extended with optimization functionalities – see, *e.g.*, [Sebastiani and Tomasi, 2015]. By combining symbolic reachability techniques and optimization, OMT solvers such as *vZ* [Bjørner *et al.*, 2015] can be leveraged to generate plans with *formal* performance guarantees by reasoning on expressive models that combine temporal and ordering constraints on tasks and robots.

*Joint work with E. Ábrahám, G. Lakemeyer, T. Niemueller and A. Tacchella.




Figure 1: Simulated RCLL arena [Niemueller *et al.*, 2015].

If *synthesizing* plans for the class of problems considered can be challenging, so it can be *acting* according to such plans. This is especially true in multi-robot systems, where execution can fail due to several reasons (*e.g.*, unexpected events, breaks in communication). To cater for this, the dynamics that occur during execution must be considered, making *integrated* approaches essential to maintain efficiency.

The goal of this research is to address the problems introduced above by combining techniques from AI and Formal Methods. We propose to integrate OMT-based plan synthesis into an on-line execution and monitoring system based on CLIPS [Wygant, 1989], a rule-based production system using forward chaining inference. In the following, we will present what we achieved so far together with ongoing work and future directions inspired by current results.

2 Contributions

To generate optimal plans, we started from a traditional *Planning as Satisfiability* (PaS) [Kautz and Selman, 1992] framework and extended it to enable optimization over reward structures expressed in first-order arithmetic theories. This idea was first presented in [Leofante *et al.*, 2017], where we show how a formal model for the *exploration phase* of the RCLL can be defined and represented symbolically using Boolean combinations of linear constraints over real variables. During exploration, teams of robots roam the factory shop floor and determine where the team’s own machines are positioned. The task is then to compute a plan that instructs multi-robot teams to explore different areas efficiently. As shown in the paper, although the formulation of this problem looks simple, computing a solution efficiently proved to be challenging due to the heavy combinatorial nature of the

problem. To tackle this, we proposed a new structure for representing PaS problems which produces more compact encodings by combining abstract state representations with relaxations of ordering properties in the formula encoding. The absence of explicit encoding of ordering properties requires an increased effort during solving. While for smaller problems this effort overweighs the advantages of our relaxations, results rapidly change with increasing problem size.

Meanwhile, we focused on how this approach could be integrated into the CLIPS-based reasoning agent presented in [Niemueller *et al.*, 2013] and currently used by the RCLL world champions. Preliminary work in this direction was presented in [Niemueller *et al.*, 2017], where together with other colleagues we reasoned on how plans should be represented internally to ease the exchange between the OMT synthesis module and the execution agent. As a result, a prototypical implementation of a system integrating our synthesis approach into the on-line execution agent was presented.

As anticipated in [Leofante, 2018], we then shifted our attention to the main challenge presented by the RCLL, that is, the *production phase*. During this phase of the game robots in the team receive orders dynamically and cooperate to deliver finished products within specified temporal windows. This phase of the game presents several challenges as it requires careful modeling of interactions between multiple robots. Furthermore, time windows make task planning hard because they require to encode logically both spatial and temporal relationships among tasks and robots within the team.

To solve this problem, we defined a new logical model of production processes accounting for robots, machines, orders and rewards. Then, the integrated system presented in [Niemueller *et al.*, 2017] was extended to support the reasoning needed to execute and monitor the new plans. A detailed description of this work is presented in [Leofante *et al.*, 2018], where an extensive experimental analysis shows how the approach we propose compares to other approaches – namely, PDDL-based planning and a greedy rule-based reactive execution used by the RCLL world champions.

3 Current and Anticipated Progress

Motivated by the results obtained, we identified three concrete directions along which this work could be extended.

Goal Reasoning. Providing robots with the ability to reason about their goals can increase performances, especially in dynamic domains like the RCLL. So far, we only exploited simple goal rewards – *e.g.*, minimize time to delivery. In a more encompassing view, goal structures can be used to, *e.g.*, select among actions having different priorities, assess risk-reward trade offs. Given that OMT allows ample flexibility in the definition of goal structures, we would like to study if and how OMT-based reasoning could be used to make informed choices on what goals robots should pursue in the future.

Explaining Plans. The problem of generating explanations for decisions taken by autonomous robotic systems is very pressing, one of the main limitation of these systems being their current inability to explain their decisions and actions in a human-readable way. In [Leofante *et al.*, 2018] we started looking into this problem, as we believe OMT-based synthe-

sis builds on techniques that have the potential to ease explaining and facilitate understanding of the underlying decision process – *i.e.*, generation of *unsat cores*, possibility to combine several metrics to assess the quality of a plan.

Novel Relaxations. Our experiments with different encodings of planning problems into OMT indicate that considerable progress can be made by considering novel kinds of relaxations. Beyond computational concerns, new relaxations can be of great interest from a representational standpoint. A prototypical implementation of a domain-independent planner based on our findings is under way and will be used to assess the impact of this research on the broader field of AI planning.

References

- [Björner *et al.*, 2015] Nikolaj Björner, Anh-Dung Phan, and Lars Fleckenstein. *vz* - an optimizing SMT solver. In *TACAS*, pages 194–199, 2015.
- [Hofmann *et al.*, 2016] Till Hofmann, Tim Niemueller, Jens Claßen, and Gerhard Lakemeyer. Continual planning in Golog. In *AAAI*, pages 3346–3353, 2016.
- [Kautz and Selman, 1992] Henry A. Kautz and Bart Selman. Planning as satisfiability. In *ECAI*, pages 359–363, 1992.
- [Leofante *et al.*, 2017] Francesco Leofante, Erika Ábrahám, Tim Niemueller, Gerhard Lakemeyer, and Armando Tacchella. On the synthesis of guaranteed-quality plans for robot fleets in logistics scenarios via optimization modulo theories. In *IRI*, pages 403–410, 2017.
- [Leofante *et al.*, 2018] Francesco Leofante, Erika Ábrahám, Tim Niemueller, Gerhard Lakemeyer, and Armando Tacchella. Integrated synthesis and execution of optimal plans for multi-robot systems in logistics. *Information Systems Frontiers*, 2018.
- [Leofante, 2018] Francesco Leofante. Guaranteed plans for multi-robot systems via Optimization Modulo Theories. In *AAAI*, 2018.
- [Niemueller *et al.*, 2013] Tim Niemueller, Gerhard Lakemeyer, and Alexander Ferrein. Incremental Task-level Reasoning in a Competitive Factory Automation Scenario. In *AAAI Spring Symposium - Designing Intelligent Robots: Reintegrating AI*, 2013.
- [Niemueller *et al.*, 2015] Tim Niemueller, Gerhard Lakemeyer, and Alexander Ferrein. The RoboCup Logistics League as a benchmark for planning in robotics. In *PlanRob@ICAPS*, 2015.
- [Niemueller *et al.*, 2017] Tim Niemueller, Gerhard Lakemeyer, Francesco Leofante, and Erika Ábrahám. Towards CLIPS-based task execution and monitoring with SMT-based decision optimization. In *PlanRob@ICAPS*, 2017.
- [Sebastiani and Tomasi, 2015] Roberto Sebastiani and Silvia Tomasi. Optimization modulo theories with linear rational costs. *ACM Trans. Comput. Log.*, 16(2), 2015.
- [Wygant, 1989] Robert M. Wygant. CLIPS: A powerful development and delivery expert system tool. *Computers & Industrial Engineering*, 17(1–4), 1989.