FOND Planning for LTL_f and PLTL Goals

Francesco Fuggitti
York University
Toronto, Canada
fuggitti@eecs.yorku.ca

Planning for temporally extended goals with *deterministic* actions has been well studied during the years starting from (Bacchus and Kabanza 1998) and (Doherty and Kvarnstram 2001). Two main reasons why temporally extended goals have been considered over the classical goals, viewed as a desirable set of final states to be reached, are because they are not limited in what they can specify and they allow us to restrict the manner used by the plan to reach the goals. Indeed, temporally extended goals are fundamentals for the specification of a collection of real-world planning problems. Yet, many of these real-world planning problems have a *non-deterministic* behavior owing to unpredictable environmental conditions. However, planning for temporally extended goals with *non-deterministic* actions is a more challenging problem and has been of increasing interest only in recent years with (Camacho et al. 2017; De Giacomo and Rubin 2018).

In this scenario, we aim at proposing a new solution to the problem of *fully observable* non deterministic domains (FOND) (Ghallab, Nau, and Traverso 2004; Geffner and Bonet 2013), where temporally extended goals are expressed not only using Linear Temporal Logic on finite traces (LTL_f) as in previous researches, but also employing its counterpart Past Linear Temporal Logic (PLTL). Indeed, an extended goal specification, expressed in PLTL, can be thought as reaching a final state such that the history leading to such a state satisfies a PLTL formula. For example, in the Triangle Tireworld domain we can think of a PLTL goal like $\varphi = vehicleAt(l22) \wedge \Diamond\!\!\!\!-(vehicleAt(l31))$, namely reach location *l22* passing through location *l31* once. Furthermore, the investigation on planning for PLTL goals may have a computational advantage since PLTL formulas can be reduced to the corresponding Deterministic Finite-state Automaton (DFA) in single-exponential time (vs. double-exponential time of LTL_f formulas) (Chandra, Kozen, and Stockmeyer 1981).

The objective of this work is not only to allow past goal specifications (i.e. using PLTL goals), but also to provide a new formalization of those temporally extended goals in the Planning Domain Definition Language (PDDL). This new approach, called FOND4LTL_f/PLTL, stems from the research in (Camacho et al. 2017) and (De Giacomo and Rubin 2018), that, basically, propose automata-theoretic foundations of FOND planning for LTL/LTL_f goals.

Drawing from these previous works, we exploit the translation of temporal formulas to automata, but in a different way. In particular, our approach performs the following steps:

1. given an LTL_f/PLTL formula representing the goal, we translate it into its corresponding minimized DFA;
2. we encode such a DFA into the non-deterministic planning PDDL domain;
3. we change the planning PDDL problem accordingly;
4. we use an off-the-shelf (FOND) planner to solve the problem;
5. we refine the potential policy found.

The translation of temporal goals into the corresponding DFA is done through a tool, called LTL_f2DFA[1] (Fuggitti 2018), that employs MONA (Elgaard, Klarlund, and Møller 1998)[2]. Following (De Giacomo and Vardi 2013; Zhu, Pu, and Vardi 2019) works, LTL_f2DFA first translates the LTL_f/PLTL input formula into First Order Logic (FOL) over finite sequences, then it encodes the result in a suitable input for MONA, and finally it outputs the expected minimized DFA. It has been shown by (Zhu et al. 2017) that this approach outperforms explicit tools such as SPOT (Duret-Lutz et al. 2016).

Subsequently, the domain is modified introducing the encoding of a general representation (i.e. replacing atomic propositions with variables) of the resulting DFA. This modification allows us to introduce a new PDDL action in the domain, called `trans`, which simulates the automata transitions. Moreover, we introduce an additional predicate, `turnDomain`, that enables us to switch between domain actions and the `trans` action. After that, we augment the initial state and goal state of the problem instance taking into account the changes made in the domain.

As a result, we obtain a new PDDL domain and instance that represent a new world in which the agent is asked to achieve the goal satisfying the temporal specification. In

[1] LTL_f2DFA is also available online at http://ltlf2dfa.diag.uniroma1.it/

[2] MONA is a tool that translates from the Weak Second-order Theory of One or Two successors (WS1S/WS2S) to symbolic DFA

other words, we compile extended temporal goals together with the original planning domain and instance, specified in PDDL, reducing the original problem to a classic FOND planning problem. Therefore, we can feed any off-the-shelf FOND planner with our new classical FOND planning problem and get a policy, if one exists.

Finally, we refine the potential policy found deleting each occurrence of the `trans` action.

In future work, we plan to extend this approach to *partially observable* non deterministic domains and further investigate the possible computational advantage in using temporally extended goals with past modalities.

Acknowledgments

This work has been done as a part of my Master Thesis, under the supervision of Prof. Giuseppe De Giacomo.

References

Bacchus, F., and Kabanza, F. 1998. Planning for temporally extended goals. *Annals of Mathematics and Artificial Intelligence* 22(1-2):5–27.

Camacho, A.; Triantafillou, E.; Muise, C. J.; Baier, J. A.; and McIlraith, S. A. 2017. Non-deterministic planning with temporally extended goals: LTL over finite and infinite traces. In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence (AAAI)*, 3716–3724.

Chandra, A.; Kozen, D.; and Stockmeyer, L. 1981. Alternation. *J. ACM* 28(114-133):531–543.

De Giacomo, G., and Rubin, S. 2018. Automata-theoretic foundations of fond planning for ltlf and ldlf goals. In *IJCAI*, 4729–4735.

De Giacomo, G., and Vardi, M. Y. 2013. Linear temporal logic and linear dynamic logic on finite traces. In *IJCAI*, volume 13, 854–860.

Doherty, P., and Kvarnstram, J. 2001. Talplanner: A temporal logic-based planner. *AI Magazine* 22(3):95.

Duret-Lutz, A.; Lewkowicz, A.; Fauchille, A.; Michaud, T.; Renault, E.; and Xu, L. 2016. Spot 2.0—a framework for ltl and $omega$-automata manipulation. In *International Symposium on Automated Technology for Verification and Analysis*, 122–129. Springer.

Elgaard, J.; Klarlund, N.; and Møller, A. 1998. MONA 1.x: new techniques for WS1S and WS2S. In *Proc. 10th International Conference on Computer-Aided Verification (CAV)*, volume 1427 of *LNCS*, 516–520. Springer-Verlag.

Fuggitti, F. 2018. LTLf2DFA. http://ltlf2dfa.diag.uniroma1.it/.

Geffner, H., and Bonet, B. 2013. A concise introduction to models and methods for automated planning. *Synthesis Lectures on Artificial Intelligence and Machine Learning* 8(1):1–141.

Ghallab, M.; Nau, D.; and Traverso, P. 2004. *Automated Planning: theory and practice*. Elsevier.

Zhu, S.; Tabajara, L. M.; Li, J.; Pu, G.; and Vardi, M. Y. 2017. Symbolic LTLf synthesis. *arXiv preprint arXiv:1705.08426*.

Zhu, S.; Pu, G.; and Vardi, M. Y. 2019. First-order vs. second-order encodings for ltlf-to-automata translation. *arXiv preprint arXiv:1901.06108*.