Bridging the Gap Between LTL Synthesis and Automated Planning

Alberto Camacho[†], Jorge A. Baier[‡], Christian Muise[⋆], Sheila A. McIlraith[†]
[†]Department of Computer Science, University of Toronto
[‡]Pontificia Universidad Católica de Chile, and Chilean Center for Semantic Web Research
[⋆]CSAIL, Massachusetts Institute of Technology
[†]{acamacho,sheila}@cs.toronto.edu, [‡]jabaier@ing.puc.cl, [⋆]cjmuise@mit.edu

Abstract

Linear Temporal Logic (LTL) synthesis can be understood as the problem of building a controller that defines a winning strategy, for a two-player game against the environment, where the objective is to satisfy a given LTL formula. It is an important problem with applications in software synthesis, including controller synthesis. Recent work has explored the close connection between automated planning and LTL synthesis but has not provided a full mapping between the two problems nor have its practical implications been explored. In this paper we establish the correspondence between LTL synthesis and fully observable non-deterministic (FOND) planning. We study LTL interpreted over both finite and infinite traces. We also provide the first explicit compilation that translates an LTL synthesis problem to a FOND problem. Experiments with state-of-the-art LTL FOND and synthesis solvers show automated planning to be a viable and effective tool for highly structured LTL synthesis problems.

1 Introduction

The problem of synthesizing software, including controllers, from logical specification is a fundamental problem in AI and computer science more generally. Church's synthesis problem was first posed by Church in 1957 in the context of synthesizing digital circuits from a logical specification (Church 1957) and is considered one of the most challenging problems in reactive systems (Piterman *et al.* 2006). Two common approaches to solving the problem have emerged: reducing the problem to the emptiness problem of tree automata, and characterizing the problem as a two-player game.

In 1989, Pnueli and Rosner examined the problem of reactive synthesis using Linear Temporal Logic (LTL) (Pnueli 1977) as the specification language (what we refer to here as "LTL synthesis") viewing the problem as a two-player game, and showing that this problem was 2EXPTIME-complete (Pnueli and Rosner 1989). Over the years, this discouraging result has been mitigated by the identification of several restricted classes of LTL for which the complexity of the synthesis problem need not be so high (e.g., (Asarin *et al.* 1998; Alur and La Torre 2004)). More recently Piterman, Pnueli, and Sa'ar examined the synthesis of reactive designs when the LTL specification was restricted to the class of so-called *Generalized Reactivity(1)* (GR1) formulae, presenting an N^3-time algorithm which checks whether the formula is realizable, and in the case where it is, constructs an automaton representing one of the possible implementing circuits (Piterman *et al.* 2006). Today, a number of synthesis tools exist with varying effectiveness (e.g., Acacia+ (Bohy *et al.* 2012), Lily (Jobstmann and Bloem 2006)).

Recent work has explored various connections between automated planning and synthesis (e.g., (De Giacomo *et al.* 2010; Patrizi *et al.* 2013; Sardiña and D'Ippolito 2015; De Giacomo and Vardi 2015)) but has not provided a full mapping between the two problems, nor have the practical implications of such a mapping been explored from an automated planning perspective. In this paper we investigate the relationship between (LTL) synthesis and automated planning, and in particular (LTL) Fully Observable Non-Deterministic (FOND) planning. We do so by leveraging a correspondence between FOND and 2-player games. This work is inspired by significant recent advances in the computational efficiency of FOND planning that have produced FOND planners that scale well in many domains (e.g., NDP (Alford *et al.* 2014), FIP (Fu *et al.* 2011), myND (Mattmüller *et al.* 2010) and PRP (Muise *et al.* 2012)). Our insights are that just as SAT can be (and has been) used as a black-box solver for a myriad of problems that can be reduced to SAT, so too can FOND be used as a black-box solver for suitable problems. Establishing the connection between FOND and 2-player games not only provides a connection to LTL synthesis – the primary subject of this exploration – it also provides the key to leveraging FOND for other problems.

In Section 3 we establish the correspondence between LTL synthesis and strong solutions to FOND planning. This is followed in Section 4 by the first approach to automatically translate a realizability problem, given by an LTL specification, into a planning problem, described in the Planning Domain Definition Language (PDDL), the de facto standard input language for automated planners. Experiments with state-of-the-art LTL synthesis and FOND solvers illustrate that the choice of formalism and solver technology for a problem can have a dramatic impact. We elucidate some of the properties that would indicate why one technique should be used over the other. As a general rule-of-thumb, if the problem is highly structured and the uncertainty largely restricted, planning-based approaches will excel. Such highly

structured problems are evident in synthesis problems for physical devices.

2 Preliminaries

2.1 FOND

A FOND planning problem is a tuple $\langle \mathcal{F}, \mathcal{I}, \mathcal{G}, \mathcal{A} \rangle$, where $\mathcal{F}$ is a set of *fluents*; $\mathcal{I} \subseteq \mathcal{F}$ characterizes what holds initially; $\mathcal{G} \subseteq \mathcal{F}$ characterizes what must hold for the goal to be achieved; and $\mathcal{A}$ is the set of actions. The set of literals of $\mathcal{F}$ is $Lits(\mathcal{F}) = \mathcal{F} \cup \{\neg f \mid f \in \mathcal{F}\}$.

Each action $a \in \mathcal{A}$ is associated with $\langle Pre_a, Eff_a \rangle$, where $Pre_a \subseteq Lits(\mathcal{F})$ is the precondition and Eff_a is a set of outcomes of a. We sometimes write $oneof(Eff_a)$ to emphasize that Eff_a is non-deterministic. Each outcome $e \in Eff_a$ is a set of conditional effects of the form $(C \to \ell)$, where $C \subseteq Lits(\mathcal{F})$ and $\ell \in Lits(\mathcal{F})$. Given a planning state $s \subseteq \mathcal{F}$ and a fluent $f \in \mathcal{F}$, we say that s satisfies f, denoted $s \models f$, iff $f \in s$. In addition $s \models \neg f$ if $f \notin s$, and $s \models L$ for a set of literals L, if $s \models \ell$ for every $\ell \in L$.

Action a is *applicable* in state s if $s \models Pre_a$. We say *s' is a result of applying a in s* iff, for one outcome e in Eff_a, s' is equal to $s \setminus \{f \mid (C \to \neg f) \in e, s \models C\} \cup \{f \mid (C \to f) \in e, s \models C\}$. A *policy* p, is a partial function from states to actions such that if $p(s) = a$, then a is applicable in s. An *execution* π of a policy p in state s is a sequence $s_0, a_0, s_1, a_1, \ldots$ (finite or infinite), where $s_0 = s$, and such that every state-action-state substring s, a, s' are such that $p(s) = a$ and s' is a result of applying a in s. Finite executions ending in a state s are such that $p(s)$ is undefined.

A finite execution π *achieves* a set of literals L if its ending state s is such that $s \models L$. An infinite execution π *achieves* a set of literals L if there exists a state s that appears infinitely often in π and that is such that $s \models \mathcal{G}$. An infinite execution σ is *fair* iff whenever s, a occurs infinitely often within σ, then so does s, a, s', for every s' that is a result of applying a in s (Geffner and Bonet 2013). Note this implies that finite executions are fair. A policy p is a *strong-cyclic plan* for a FOND problem $P = \langle \mathcal{F}, \mathcal{I}, \mathcal{G}, \mathcal{A} \rangle$, iff every fair execution of p over $\mathcal{I}$ satisfies the goal. A policy p is a *strong plan* for P iff every execution of p over $\mathcal{I}$ satisfies $\mathcal{G}$.

2.2 Linear Temporal Logic and Automata

Linear Temporal Logic (LTL) is a propositional logic extended with temporal modal operators *next* ($\circ$) and *until* (U). The set of LTL formulae over a set of propositions $\mathcal{P}$ is defined inductively as follows. p is a formula if $p \in \mathcal{P}$ or the constant $\top$. If φ_1 and φ_2 are LTL formulas, then so are $\neg\varphi_1$, $\varphi_1 \wedge \varphi_2$, $\circ\varphi_1$ and $\varphi_1 \,\mathsf{U}\, \varphi_2$. Let $\sigma = s_0, s_1, \ldots$ be an infinite sequence of subsets of $\mathcal{P}$, and φ be an LTL formula. Then σ *satisfies* φ, denoted as $\sigma \models \varphi$ iff $\sigma, 0 \models \varphi$, where:

- $\sigma, i \models p$, for each $p \in \mathcal{P} \cup \{\top\}$ iff $s_i \models p$.
- $\sigma, i \models \neg\varphi$ iff $\sigma, i \models \varphi$ does not hold.
- $\sigma, i \models \varphi_1 \wedge \varphi_2$ iff $\sigma, i \models \varphi_1$ and $\sigma, i \models \varphi_2$.
- $\sigma, i \models \circ\varphi$ iff $\sigma, (i+1) \models \varphi$.
- $\sigma, i \models \varphi_1 \,\mathsf{U}\, \varphi_2$ iff there exists a $j \geq i$ such that $\sigma, j \models \varphi_2$, and $\sigma, k \models \varphi_1$, for each $k \in \{i, i+1, \ldots, j-1\}$.

Intuitively, the *next* operator tells what needs to hold in the next time step, and the *until* operator tells what needs to hold until something else holds. The modal operators *eventually* ($\Diamond$) and *always* ($\Box$) are defined by $\Diamond\varphi \equiv \top \,\mathsf{U}\, \varphi$, $\Box\varphi \equiv \neg\Diamond\neg\varphi$. Additional constants and operators are defined by following conventional rules as follows $\bot \equiv \neg\top$, $\varphi_1 \vee \varphi_2 \equiv \neg(\neg\varphi_1 \wedge \neg\varphi_2)$, $\varphi_1 \to \varphi_2 \equiv \neg\varphi_1 \vee \varphi_2$, $\varphi_1 \leftrightarrow \varphi_2 \equiv (\varphi_1 \to \varphi_2) \wedge (\varphi_2 \to \varphi_1)$.

LTL over finite traces: Variants of LTL interpreted over finite traces have been studied in the verification, model checking and planning communities. Within the planning community, LTL interpreted over finite traces has been used to specify temporally extended goals (e.g. (Bacchus and Kabanza 1998; 2000; Baier and McIlraith 2006b; 2006a; Edelkamp 2006; Camacho *et al.* 2017)), and temporally extended preferences (e.g., (Baier *et al.* 2009; Bienvenu *et al.* 2011; Coles and Coles 2011)), and a number of planners exist. $\mathsf{LTL_f}$ is one of the most recent and popular examples (De Giacomo and Vardi 2013). Whereas $\mathsf{LTL_f}$ and LTL share the same syntax, its interpretations can be rather different. For example, the $\mathsf{LTL_f}$ formula $\Diamond\neg\circ\top$ is true in a finite trace, whereas the same formula in LTL evaluates false on an infinite traces. Similarly, *weak next* must often replace *next* to avoid unintended interpretations of LTL over finite traces.

Automata: There is a well-established correspondence between LTL and automata. A *Non-deterministic Büchi Automaton* (NBA) is a tuple $M = (Q, \Sigma, \delta, q_0, Q_{Fin})$, where Q is the set of automaton states, Σ is the alphabet, $\delta \subseteq Q \times \Sigma \times Q$ is the transition relation, q_0 is the initial state of the automaton, and $Q_{Fin} \subseteq Q$ is the set of accepting states. The automaton is *deterministic* (DBA) when for each $q \in Q$, and $s \in \Sigma$, there exists a unique $q' \in Q$ such that $(q, s, q') \in \delta$. A *run* of M on an infinite word $\sigma = s_0, s_1, \ldots$ of elements in Σ is a sequence of automaton states, starting in q_0, such that $(q_i, s_i, q_{i+1}) \in \delta$ for all $i \geq 0$. A run is accepting if it visits an infinite number of accepting states. Finally, we say that M accepts σ if there is an accepting run of M on σ. *Non-deterministic Finite-state Automata* (NFAs) differ from NBAs in that the acceptance condition is defined on *finite* words: a word $\sigma = s_0, s_1, \ldots, s_m$ is accepting if $q_{m+1} \in Q_{Fin}$. Finally, *Deterministic Finite-state Automata* are NFAs where the transition relation is deterministic.

Given an LTL formula φ, it is possible to construct an NBA A_φ that accepts σ iff $\sigma \models \varphi$. The construction is worst-case exponential in the size of φ (Vardi and Wolper 1994). It is not always possible to construct a DBA, and the construction is double exponential. Similar results hold for $\mathsf{LTL_f}$: it is always possible to construct an NFA (resp. DFA) that accepts σ iff $\sigma \models \varphi$, and the construction is worst-case exponential (resp. double exponential) (Baier and McIlraith 2006b). Figure 1 depicts an NBA corresponding to LTL formula $\Diamond\Box(x \leftrightarrow y)$, which does not have a DFA representation. Automaton states are represented by circles (double-ringed in accepting states), and transitions are represented with arrows.

2.3 LTL FOND

Recently Camacho *et al.* (2017) extended FOND with LTL goals. An LTL FOND problem is a tuple $\langle \mathcal{F}, \mathcal{I}, \mathcal{G}, \mathcal{A} \rangle$,

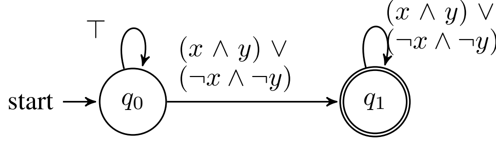

Figure 1: An NBA corresponding to $\Diamond\Box(x \leftrightarrow y)$. q_1 is an accepting state, and q_0 has non-deterministic transitions.

where $\mathcal{G}$ is an LTL or $\mathsf{LTL_f}$ formula over $\mathcal{F}$, and $\mathcal{F}, \mathcal{I}, \mathcal{A}$ are defined as in FOND planning. In short, LTL FOND executions are defined just like in FOND, and a policy is a strong-cyclic (resp. strong) plan for problem P if each fair (resp. unrestricted) execution π results in a sequence of states σ such that $\sigma \models \mathcal{G}$.

2.4 LTL Synthesis

The LTL synthesis problem (Pnueli and Rosner 1989) intuitively describes a two-player game between a controller and the environment. The game consists of an infinite sequence of turns. In each turn the environment chooses an action, and the controller then chooses another. Each action actually corresponds to setting the values of some variables. The controller has a winning strategy if, no matter how the environment plays, the sequences of states generated satisfy a given LTL formula φ. Formally, a synthesis problem is a tuple $\langle \mathcal{X}, \mathcal{Y}, \varphi \rangle$, where $\mathcal{X} = \{x_1, \ldots, x_n\}$, the environment variables, and $\mathcal{Y} = \{y_1, \ldots, y_m\}$, the controller variables, are disjoint sets. An LTL formula over $\mathcal{X} \cup \mathcal{Y}$ is *realizable* if there exists a function $f : (2^{\mathcal{X}})^* \to 2^{\mathcal{Y}}$ such that for every infinite sequence of subsets of $\mathcal{X}$, $X_1X_2\ldots$, it holds that $\pi = (X_1 \cup f(X_1)), (X_2 \cup f(X_1X_2)) \ldots$ is such that $\pi \models \varphi$. Intuitively, no matter what the choice of the environment is, which is given by the sequence $X_1X_2\ldots$, the controller has a strategy, given by f, that ensures formula φ is satisfied in the resulting game. The *synthesis* problem corresponds to actually computing function f. Some authors have studied the synthesis problem as a game over *finite* sequences of turns, using $\mathsf{LTL_f}$ to describe specifications (e.g. (De Giacomo and Vardi 2015)). In the rest of the paper, we write LTL synthesis to also refer to $\mathsf{LTL_f}$ synthesis, and make the distinction explicit only when necessary.

3 Relationship Between FOND and Synthesis

Both LTL synthesis and FOND are related to two-player games: in both problems an agent (or controller) seeks a solution that achieves a condition no matter what choices are taken by the environment. There are however two important differences. First, in LTL synthesis the controller reacts to the environment; in other words, the environment "plays first", while the controller "plays second". Rather, in FOND, the play sequence is inverted since the environment decides the outcome of an action, which is in turn defined by the agent (controller). Second, state-of-the-art FOND solvers find strong-cyclic solutions, and indeed those types of solutions are considered standard. This assumes fairness in the environment, which is not an assumption inherent to LTL synthesis. Thus a correct mapping between FOND and Synthesis should handle fairness correctly.

Previous work has explored the relation between FOND and synthesis. Sardiña and D'Ippolito (2015) show how to translate FOND as a reactive synthesis problem by expressing fairness constraints as temporal logic formulae. De Giacomo and Vardi (2013) sketches a mapping from FOND to LTL synthesis, in which the effects of actions are specified using LTL. This approach, however, does not dive into the details of the inverted turns. Neither do the works by De Giacomo *et al.*; Kissmann and Edelkamp (2010; 2009), which show a correspondence between two-player game structures and FOND planning.

In the rest of the section we provide an explicit mapping between LTL FOND and LTL synthesis. We aim at a correct mapping between both problems. Efficiency is the focus of the next section.

To establish a correspondence between LTL synthesis and LTL FOND, we address the inverted turns by considering the negation of realizability. Observe that an instance $\langle \mathcal{X}, \mathcal{Y}, \varphi \rangle$ is not realizable iff there exists a sequence $X_1X_2X_3\ldots$ of subsets of $\mathcal{X}$ such that. for every function $f : (2^{\mathcal{X}})^* \to 2^{\mathcal{Y}}$:

$$X_1 \cup f(X_1), X_2 \cup f(X_1X_2), X_3 \cup f(X_1X_2X_3) \ldots \models \neg\varphi$$

Note that what comes after the "iff" maps directly into an instance of LTL FOND: we define the problem $P_\varphi = \langle \mathcal{F}, \mathcal{I}, \mathcal{G}, \mathcal{A} \rangle$ such that fluents are the union of all variables (i.e., $\mathcal{F} = \mathcal{X} \cup \mathcal{Y}$), and the set of actions is the set of subsets of $\mathcal{X}$ (i.e., $\mathcal{A} = \{a_x \mid x \subseteq \mathcal{X}\}$). Intuitively action a_x is always executable (has empty preconditions) and deterministically sets to true the variables in x and to false the variables in $\mathcal{X} \setminus x$. In addition, it non-deterministically sets the values of variables in $\mathcal{Y}$ to every possible combination. Formally, $\mathit{Eff}_{a_x} = \{e_{x,y} \mid y \subseteq \mathcal{Y}\}$, where each $e_{x,y} = \{f \mid f \in x \cup y\} \cup \{\neg f \mid f \in (\mathcal{X} \cup \mathcal{Y}) \setminus (x \cup y)\}$. Finally, we set $\mathcal{I} = \{\}$ and $\mathcal{G} = \bigcirc\neg\varphi$.

A more involved argument follows for $\mathsf{LTL_f}$ synthesis. In this case, an instance $\langle \mathcal{X}, \mathcal{Y}, \varphi \rangle$ is not realizable iff for every finite m there exists a sequence $X_1X_2\ldots X_m$ of subsets of $\mathcal{X}$ such that, for every function $f : (2^{\mathcal{X}})^* \to 2^{\mathcal{Y}}$:

$$X_1 \cup f(X_1), \ldots, (X_1 \ldots X_m) \cup f(X_1 \ldots X_m) \models \neg\varphi$$

What follows after the "iff" cannot be directly mapped into an instance of LTL FOND, because the formula above has to hold for all m. We can mitigate for this by adding a new variable to $\mathcal{P}_\varphi$, y_{ok}, that acts like any other variable in $\mathcal{Y}$. The goal of $\mathcal{P}_\varphi$ is the $\mathsf{LTL_f}$ formula $\mathcal{G} = \bigcirc(\neg\varphi \wedge \Diamond(y_{ok} \wedge \neg\bigcirc\top)) \vee \Diamond(\neg y_{ok} \wedge \neg\bigcirc\top)$.

Theorem 1. *An* LTL *synthesis problem* $\langle \mathcal{X}, \mathcal{Y}, \varphi \rangle$ *is realizable iff* P_φ *has no strong plan.*

In the other direction, let $P = \langle \mathcal{F}, \mathcal{I}, \mathcal{G}, \mathcal{A} \rangle$ be an LTL FOND problem. We now construct a synthesis problem $\langle \mathcal{X}_P, \mathcal{Y}_P, \varphi_P \rangle$ following well-known encodings of planning into SAT (Rintanen *et al.* 2006); we no longer play inverted turns, and we use LTL to establish a connection between a state and its successors, instead of different variables, and that we consider explicitly that actions have a number of outcomes. The specification φ_P is defined by:

$$\varphi_P := \varphi_{init} \to (\varphi_{env} \to (\varphi_{agt} \wedge \varphi_g))$$

Intuitively, φ_{init} models the initial state $\mathcal{I}$, φ_{env} and φ_{agt} model the dynamics in $\mathcal{P}$, and φ_g is the LTL goal formula $\mathcal{G}$.

For each action in $a \in \mathcal{A}$, we create a variable $a \in \mathcal{X}_P$. Each fluent $f \in \mathcal{F}$ is also a variable in $\mathcal{X}_P$. Variables in $\mathcal{Y}_P$ are used to choose one of the non-deterministic outcomes of each action; this way if the action with the largest number of outcomes has n outcomes, we create $\lceil \log n \rceil$ variables, whose objective is to "choose" the outcome for an action. To model the preconditions of the action, we conjoin in φ_{env}, for each action a the formula $\Box(a \rightarrow \bigwedge_{\ell \in Pre_a} \ell)$. We express the fact that only one action can execute at a time by conjoining to φ_{env} the formulae $\Box \bigvee_{a \in \mathcal{A}} a$, and $\Box(a \rightarrow \neg a')$, for each $a' \in \mathcal{A}$ different from a. To model the fact that the environment selects the outcome being performed, for each action outcome e we create a variable a_e in $\mathcal{X}_P$. For each action $a \in \mathcal{A}$ and outcome $e \in \mathit{Eff}_a$, φ_{agt} has formulae of the form $\Box(a \wedge \chi_{a,e} \rightarrow a_e)$, where $\chi_{a,e}$ is a formula over $\mathcal{Y}_P$, which intuitively "selects" outcome e for action a. For space, we do not go into the details of how to encode $\chi_{a,e}$. However, these formulae have the following property: for any action a, given an assignment for $\mathcal{Y}_P$ variables there is exactly one $e \in \mathit{Eff}_a$ for which $\chi_{a,e}$ becomes true. This models the fact that the $\mathcal{Y}_P$ variables are used to select the outcomes.

Finally, we now conjoin to φ_{env} formulae to express the dynamics of the domain. Specifically we add successor-state-axiom-like expressions (Reiter 2001) of the form:

$$\Box(\circ f \equiv (\phi_f^+ \vee (f \wedge \neg \phi_f^-))), \quad \text{for each } f \in \mathcal{F}$$

where ϕ_f^+ is a formula that encodes the conditions under which f becomes true after an outcome has occurred, and where ϕ_f^- encodes the conditions under which f becomes false in the next state. Both of these formulae can be computed from Eff_a (Reiter 2001), and have fluents a_e for $e \in \mathit{Eff}_a$. Finally, φ_{init} is the conjunction of the fluents in the initial state $\mathcal{I}$, and φ_g is the goal formula, $\mathcal{G}$. When the goal of $\mathcal{P}$ is an $\mathsf{LTL_f}$ formula, the construction conjoins $\circ\top$ to the successor state axioms in φ_{env}.

Now, it is not hard to see that there exists a strong solution to the LTL problem P iff there exists a (finite for $\mathsf{LTL_f}$ goals, infinite for LTL) sequence of settings of the $\mathcal{X}_P$ variables, such that for every sequence of settings of the $\mathcal{Y}$ variables (i.e., for every function $f : (2^{\mathcal{X}})^* \rightarrow 2^{\mathcal{Y}}$), it holds that:

$$(X_1 \cup Y_1), (X_2 \cup Y_2), (X_3 \cup Y_3), \ldots \models \varphi_P$$

Theorem 2. *An LTL FOND problem $P = \langle \mathcal{F}, \mathcal{I}, \mathcal{G}, \mathcal{A} \rangle$ has a strong plan iff $\langle \mathcal{X}_P, \mathcal{Y}_P, \neg\varphi_P \rangle$ is not realizable.*

4 Approach

In Section 3 we established the correspondence between existence of solutions to LTL synthesis, and existence of strong solutions to LTL FOND planning. In this section we introduce the first translation from LTL synthesis into FOND planning (and by inclusion, into LTL FOND), and a translation for $\mathsf{LTL_f}$ specifications.

4.1 Compiling LTL Synthesis to FOND

Our approach to solve an LTL synthesis problem $\mathcal{P} = \langle \mathcal{X}, \mathcal{Y}, \varphi \rangle$ as FOND consists of three stages. First, we pre-process $\mathcal{P}$. Second, we compile it into a standard FOND problem $\mathcal{P}'$. Finally, we solve $\mathcal{P}'$ with a strong-cyclic planner. Extracting a strategy for $\mathcal{P}$ from a solution to $\mathcal{P}'$ is straightforward, and we omit the details for lack of space.

Automaton Transformation: In a pre-processing stage, we first simplify the specification, if possible, by removing from $\mathcal{X}$ and $\mathcal{Y}$ those variables that do not appear in φ (cf. Section 4.2). Then, we transform φ into an automaton, $A_\varphi = (Q, \Sigma, \delta, q_0, Q_{Fin})$, that can be an DBA when the LTL formula is interpreted over infinite traces, or an NFA (or DFA, by inclusion) when the specification is an $\mathsf{LTL_f}$ formula. In addition to DBAs, our algorithm can seamlessly handle NBAs at the cost of losing its completeness guarantee. NBAs are a good alternative to DBAs as they are usually more compact, and only a subset of LTL formulae can be transformed into DBAs. The transition relation δ in A_φ implicitly defines the conditions under which the automaton in state q is allowed to transition to state q'. These conditions are known as *guards*. Formally, $guard(q, q') = \bigvee_{(q,s,q') \in \delta} s$. In our case, elements of the alphabet Σ are conjunctions of boolean variables, that allow for guard formulae to be described in a compact symbolic form. In what follows, we assume guard formulae $guard(q, q') = \bigvee_m c_m$ are given in DNF, where each clause c_m is a conjunction of boolean state variables. We denote as $\delta^\star$ the set of tuples $T_m = (q, c_m, q')$ for each pair (q, q') with $guard(q, q') \neq \bot$, and for each clause c_m in $guard(q, q')$. For convenience, we write $guard(T_m) = c_m$, and refer to elements of $\delta^\star$ as *transitions*. Wherever convenient, we drop the subindex of transitions and simply write T. In the NBA of Figure 1, guards are the labels of edges connecting two automaton states, and $\delta^\star$ contains, among others, the transitions $(q_0, x \wedge y, q_1)$ and $(q_0, \neg x \wedge \neg y, q_1)$.

In the second stage, we compile $\mathcal{P} = \langle \mathcal{X}, \mathcal{Y}, \varphi \rangle$ with automaton A_φ into a parametrized FOND problem $\mathcal{P}'(\mathcal{X}, \mathcal{Y}, A_\varphi, H) = \langle \mathcal{F}, \mathcal{I}, \mathcal{G}, \mathcal{A} \rangle$ that integrates the dynamics of A_φ with a two-player game between the environment and the agent. Before introducing the technical details of the compilation, we first describe its dynamics in a high level. The compilation simulates automaton states by means of fluents q, one for each automaton state with the same name. Planning states s have the following property: *an automaton fluent q is true in s iff for some σ, a run σ of A_φ finishes in q.* Notably, the input word σ can be obtained directly from the state-action plan that leads the initial state to s in the search tree. When the input of the algorithm is a non-deterministic automaton (NBA or NFA), planning states can simultaneously capture multiple runs of the automaton in parallel by simultaneously having multiple q fluents set to true.

The acceptance condition behaves differently for Büchi and non-Büchi automata, and this is also reflected in our compilation. For Büchi automata, the planning process expands a graph that simulates the environment and agent moves, and the search for solutions becomes a search for strong cyclic components that hit an accepting state infinitely often. The latter property is checked by means of *tokenized* fluents q^t, one for each q. Intuitively, the truth of $q \wedge q^t$ in state s indicates a commitment to progress runs

finishing in q into runs that reach an accepting state. Conversely, $s \models q \wedge \neg q^t$ represents that such a commitment has been accomplished. The role of the parameter H is twofold: it indirectly bounds the horizon in the search for cycles, and it allows the use of strong-cyclic solvers to find solutions to a problem whose non-determinism has unrestricted fairness.

The dynamics of the compilation run in two modes that alternate sequentially and simulate each two-player turn. The *environment mode* simulates the environment moves, which are non-deterministic and uncontrollable to the agent. In *automaton mode*, the agent moves are simulated and the automaton state fluents are synchronized according to valid transitions in $\delta^\star$. Auxiliary fluents q^s and $q^{s,t}$ are used to store the value of automaton state fluents q and q^t prior to synchronization, so that more than one transition can be simulated in the case of non-deterministic automata compilations.[1] When an accepting state q is recognized, the agent can set the token fluents q^t to commit to progress the runs that finish in q into a run that hits another accepting state.

The dynamics of the compilation are similar for non-Büchi automata. The exception is that accepting runs are recognized whenever an accepting automaton state fluent is reached, and there is no need to commit to reaching another accepting state. Consequently, tokenized fluents q^t and $q^{s,t}$ are not needed. For generality, we include those fluents in the algorithm presented below, but we apprise the reader that any occurrence of these fluents can be safely removed from the compilation without affecting the soundness and completeness of the approach for $\mathsf{LTL_f}$ synthesis.

The sets of fluents, $\mathcal{F}$, and actions, $\mathcal{A}$, of the problem are listed below. In what follows, we describe the technical details of the compilation.

$$
\begin{aligned}
\mathcal{F} = & \left\{q, q^s, q^t, q^{s,t} \mid q \in Q\right\} \cup \{goal\} \cup \\
& \{env_mode, aut_mode, can_switch, can_accept\} \cup \\
& \{at_horizon(h)\}_{0 \le h \le H} \cup \{next(h{+}1, h)\}_{0 \le h < H} \\
& \{turn_k\}_{1 \le x \le |\mathcal{X}|} \cup \{v_x, v_{\neg x}\}_{x \in \mathcal{X}} \cup \{v_y, v_{\neg y}\}_{y \in \mathcal{Y}} \\
\mathcal{A} = & \{move_k\}_{1 \le k \le |\mathcal{X}|} \cup \{trans_T\}_{T \in \delta^\star} \cup \\
& \{switch2aut, switch2env, accept\}
\end{aligned}
$$

Environment Mode In the environment mode, the dynamics of the problem simulates the move of the environment. As this move is uncontrollable by the agent, it can be simulated with a non-deterministic action that has $2^{|\mathcal{X}|}$ effects, each one simulating an assignment to variables in $\mathcal{X}$. Fluents v_l simulate the truth value of variables in $\mathcal{X} \cup \mathcal{Y}$. More precisely, v_x (resp. $v_{\neg x}$) indicates that $x \in \mathcal{X}$ is made true (resp. false), and similarly for $y \in \mathcal{Y}$. In order to reduce the explosion in non-deterministic action effects, we simulate the environment's move with a cascade of non-deterministic actions $move_k$, each one setting (v_{xk}) or unsetting ($v_{\neg xk}$) the value of a variable x_k in $\mathcal{X}$.

$$
\begin{aligned}
Pre_{move_k} &= \{env_mode, turn_k\} \\
Eff_{move_k} &= oneof(\{v_{xk}, \neg v_{\neg xk}\}, \{\neg v_{xk}, v_{\neg xk}\}) \cup \Psi_k \\
\Psi_k &= \begin{cases} \{turn_{k+1}, \neg turn_k\}, & \text{if } k < |\mathcal{X}| \\ \{can_switch, \neg turn_k\}, & \text{if } k = |\mathcal{X}| \end{cases}
\end{aligned}
$$

After the environment's move has been simulated, the *switch2aut* action switches the dynamics to automaton mode, and the automaton configuration (represented by fluents of the form q and q^t) is *frozen* into copies q^s and $q^{s,t}$. Special predicates $at_horizon(h)$ capture the number of turns from the last recognized accepting state in the plan. If $h < H$, the horizon value is incremented by one.

$$
\begin{aligned}
Pre_{switch2aut}(h, h') = & \{env_mode, can_switch\} \cup \\
& \{at_horizon(h), next(h', h)\} \\
Eff_{switch2aut}(h, h') = & \{at_horizon(h'), \neg at_horizon(h)\} \\
& \cup \{aut_mode, \neg env_mode, \neg can_switch\} \cup \\
& \left\{q \rightarrow \{q^s, \neg q\}, q^t \rightarrow \left\{q^{s,t}, \neg q^t\right\} \mid q \in Q\right\}
\end{aligned}
$$

Automaton Mode The automaton mode simulates the assignment to variables in $\mathcal{Y}$ and the automaton state transitions. Whereas the update in the automaton configuration is usually understood as a *response* to the observation to variables in $\mathcal{X} \cup \mathcal{Y}$, the dynamics of the encoding take a different perspective: the agent can decide which automaton transitions to perform, and then set the variables in $\mathcal{Y}$ so that the transition guards are satisfied. Such transitions are simulated by means of $trans_T$ actions, one for each $T = (q_i, guard(T), q_j) \in \delta^\star$.

$$
\begin{aligned}
Pre_{trans_T} &= \{aut_mode, q_i^s, \neg q_j\} \cup \{\neg v_{\neg l}\}_{l \in guard(T)} \\
Eff_{trans_T} &= \{q_j\} \cup \{v_l\}_{l \in guard(T)} \cup \Psi_{trans_T} \\
\Psi_{trans_T} &= \begin{cases} \left\{q_i^{s,t} \rightarrow q_j^t\right\}, & \text{if } q_j \notin Q_{Fin} \\ \{can_accept\}, & \text{if } q_j \in Q_{Fin} \end{cases}
\end{aligned}
$$

A transition $T = (q_i, guard(T), q_j)$ can be simulated when there exists a run of the automaton finishing in q_i (as such, q_i had to be *frozen* into q_i^s by means of *switch2aut*). Preconditions include the set $\{\neg v_{\neg l} \mid l \in guard(T)\}$, that checks that the transition guard is not violated by the current assignment to variables. Here, we abuse notation and write $l \in guard(T)$ if the literal l appears in $guard(T)$. As usual, we use the equivalence $\neg(\neg l) = l$. The effects $\{v_l \mid l \in guard(T)\}$ set the variables in $\mathcal{Y}$ so that the guard is satisfied and T can be fired. In parallel, the automaton state fluent q_j is set, as to reflect the transition T. According to the semantics of the tokenized fluents, when $q_i^{s,t}$ holds in the current state the token is progressed into q_j^t to denote a commitment to reach an accepting state. If q_j is indeed an accepting state, then the tokenized fluent is not propagated and instead the fluent can_accept is set. Notably, the conditional effects $q_i^{s,t} \rightarrow q_j^t$ do *not* delete the copies q_i^s and $q_i^{s,t}$. This allows the agent to simulate more than one transition when the automaton is an non-deterministic, thereby capturing multiple runs of the automaton in the planning state (although it is not obliged to simulate all transitions). When the

[1] When the automaton is deterministic, the compilation can be slightly modified so that fluents q^s and $q^{s,t}$ are no longer needed.

automaton is deterministic, the effects of $trans_T$ allow for at most one transition can be simulated. Finally, the fluent q_j appears negated in the preconditions of $trans_T$ merely for efficiency purposes, as executing $trans_T$ when q_j is true has no value to the plan (and $trans_T$ can be safely pruned).

The agent has two action mechanisms to switch back to environment mode: $switch2env$ and $accept$. The agent can, at any time in the automaton mode, execute *switch2env* causing all *frozen* copies q^s and $q^{s,t}$ to be deleted.[2] The purpose of *Regularize*, which is optional, is to improve the search performance as described in Section 4.2.

$$\begin{aligned}
Pre_{switch2env} &= \{aut_mode\} \\
Eff_{switch2env} &= \{env_mode, \neg aut_mode\} \cup \{turn_1\} \cup \\
&\quad \{\neg q^s, \neg q^{s,t} \mid q \in Q\} \cup \\
&\quad \cup\, Regularize \\
Regularize &= \{\neg v_z, \neg v_{\neg z} \mid z \in \mathcal{X} \cup \mathcal{Y}\}
\end{aligned}$$

The *accept* action is useful to compilations based on Büchi automata, and recognizes runs that have satisfied a commitment to hit an accepting state. At least one of these runs exist if fluent *can_accept* (which is part of the preconditions) holds true. By executing *accept*, the agent *forgets* those runs that did not satisfy the commitment to hit an accepting state, and commits to progress the rest of the runs into runs that hit another accepting state. The agent can postpone action *accept* as much as necessary in order to progress all relevant runs into runs that hit an accepting state. Action *accept* has a non-deterministic effect *goal*, introduced artificially as a method to find infinite plans that visit accepting states infinitely often. The interested reader can find full details in Camacho *et al.* (2017).

$$\begin{aligned}
Pre_{accept}(h) &= \{aut_mode, can_accept, at_horizon(h)\} \\
Eff_{accept}(h) &= oneof(\{goal\}, \\
&\quad \{turn_1, at_horizon(0), \neg at_horizon(h)\} \cup \\
&\quad \{env_mode, \neg aut_mode, \neg can_accept\} \cup \\
&\quad \{q^s \rightarrow \neg q^s, q^{s,t} \rightarrow \neg q^{s,t} \mid q \in Q\} \cup \\
&\quad \{q \wedge q^t \rightarrow \{\neg q, \neg q^t\} \mid q \in (Q \setminus Q_{Fin})\} \cup \\
&\quad \{q \wedge \neg q^t \rightarrow q^t \mid q \in Q\} \cup Regularize)
\end{aligned}$$

Initial and Goal States The initial state of the problem is $\mathcal{I} = \{q_0, env_mode, turn_1, at_horizon(0)\} \cup \{next(h+1, h) \mid h \in 0 \ldots H\}$. When the input of the algorithm is a Büchi automaton, the goal is $\mathcal{G} = \{goal\}$. For NFAs and DFAs, the goal is $\mathcal{G} = \{can_accept\}$.

To summarize, we define our compilation method from LTL synthesis into FOND, that we call Syn2FOND.

[2] Switching to environment mode without having applied any $trans_T$ action inevitably leads to a deadend state. The compilation presented here can be slightly modified to require the application of at least one $trans_T$ action in agent mode. However, doing this did not affect the planner's performance.

Definition 1 (Syn2FOND compilation)**.** *For an* LTL *synthesis problem $\mathcal{P} = \langle \mathcal{X}, \mathcal{Y}, \varphi \rangle$, (NBA, DBA, NFA, or DFA) automaton A_φ, and parameter H, the* Syn2FOND compilation *constructs the FOND problem $\mathcal{P}'(\mathcal{X}, \mathcal{Y}, A_\varphi, H) = \langle \mathcal{F}, \mathcal{I}, \mathcal{G}, \mathcal{A} \rangle$ as described above.*

Solutions to the compiled problem $\mathcal{P}'$ yield solutions to $\mathcal{P}$ (cf. Theorem 3). The iterated search of solutions to Syn2FOND compilations (with $H = 1, 2 \ldots, 2^{|Q|}$) is guaranteed to succeed, if $\mathcal{P}$ is realizable, when the input automaton is a DBA, NFA, or DFA (cf. Theorem 4). This follows, intuitively, from the fact that if a solution exists, then a strong cyclic policy can be unfolded and simulated in a Syn2FOND compilation search graph. If the agent's strategy cannot always guarantee hitting an accepting state within $H \leq 2^{|Q|}$ turns, then the environment can force a non-accepting cycle – i.e., the environment has a winning strategy that prevents the agent from satisfying the specification. With deterministic automata, the bound can be lowered to $H \leq |Q|$. We illustrate below with a counter-example that completeness is not guaranteed for NBAs.

Theorem 3 (soundness)**.** *Strong-cyclic plans to the Syn2FOND compilation $\mathcal{P}'$ correspond to solutions for $\mathcal{P}$.*

Theorem 4 (completeness)**.** *For a* realizable LTL *synthesis problem $\mathcal{P} = \langle \mathcal{X}, \mathcal{Y}, \varphi \rangle$, and NFA automaton A_φ, the* Syn2FOND compilation *$\mathcal{P}'$ is solvable for some $H \leq 2^{|Q|}$. When A_φ is a DBA or DFA, the bound can be lowered to $H \leq |Q|$.*

Illustrative Example Let $\mathcal{P} = \langle \mathcal{X}, \mathcal{Y}, \varphi \rangle$ be an LTL synthesis problem, with $\mathcal{X} = \{x\}$, $\mathcal{Y} = \{y\}$, and $\varphi = \Diamond\Box(x \leftrightarrow y)$, and let A_φ be the NBA in Figure 1. The dynamics of the compilation are as follows. The initial state starts in environment mode and automaton state q_0. The non-deterministic environment moves generate two different states, in which v_x and $v_{\neg x}$ hold, after which the agent switches to automaton mode and freezes the automaton state q_0 into a copy q_0^c. When one of these states is processed (say, $s_1 \models \{q_0^c, v_x, aut_mode\}$) the agent can choose to apply action $trans_{T_1}$, with $T_1 = (q_0, x \wedge y, q_1)$ to transition to state q_1. Additionally, the agent can also choose to apply $trans_{T_2}$, with $T_2 = (q_0, \top, q_0)$ to self-transition to q_0. If the former action is performed, the fluent can_accept is held true because q_1 is accepting. Once at this point, the agent transitions to environment mode, and can choose to recognizably *accept* the runs that do not have token associated with the automaton fluents. Subsequently, the automaton states are *tokenized* – i.e. q_0 and q_1 become q_0^t and q_1^t, respectively. In this particular example, maintaining multiple runs in a single planning state (in this case, q_0 and q_1) reflects the non-determinism of the NBA, and allows the agent to defer the satisfaction of the sub-formula $x \leftrightarrow y$ a finite, but unbounded number of time steps. The latter property cannot be captured with deterministic automata.

Incompleteness of the NBA-based compilation The NBA-based compilation is not guaranteed to preserve solutions. Let $\varphi = \bigcirc(\Box x \vee \Diamond \neg x)$, and consider

the NBA A_φ with states $Q = \{q_0, q_1, q_2, q_3\}$, $\delta^\star = \{(q_0, \top, q_1), (q_0, \top, q_2), (q_1, x, q_1), (q_2, x, q_2), (q_3, \top, q_3)\}$, and $Q_{Fin} = \{q_1, q_3\}$. The environment can consistently play x a finite, but unbounded number of times before playing $\neg x$ – at which point the runs of the automaton that finish in q_2 must not have been forgotten. There is no bounded parameter H that can satisfy such requirement.

4.2 Exploiting Relevance

The exploitation of relevance has been extensively studied within automated planning as a means of improving the efficiency of plan generation and execution monitoring (e.g., (Haslum *et al.* 2013)). Here, we present a suite of ideas in a similar vein that leverages state relevance in service of efficient synthesis via FOND planning.

Global Irrelevance We say that a set of variables $\mathcal{Z} \subseteq \mathcal{X} \cup \mathcal{Y}$ is *globally irrelevant* wrt φ if, for every sequence $\pi = \{\mathcal{Z}_n\}_n$ of assignments to variables in $\mathcal{X} \cup \mathcal{Y}$, for every k, and for every subset $\mathcal{Z}' \subseteq \mathcal{Z}$, it holds that $\pi \models \varphi$ iff $\pi' \models \varphi$, where $\pi' = \{\mathcal{Z}'_n\}_n$ is such that $\mathcal{Z}'_n = \mathcal{Z}_n$ if $n \neq k$, and $\mathcal{Z}'_k = (Z_k \setminus \mathcal{Z}) \cup \mathcal{Z}'$. In other words, the truth of variables in $\mathcal{Z}$ in a sequence π does not influence whether π satisfies φ.

Theorem 5. *Let $\mathcal{P} = \langle \mathcal{X}, \mathcal{Y}, \varphi \rangle$ be an* LTL *synthesis problem, and let $vars(\varphi) \subseteq \mathcal{X} \cup \mathcal{Y}$ be the set of variables that appear in φ. Then, $\mathcal{Z} = (\mathcal{X} \cup \mathcal{Y}) \setminus vars(\varphi)$ is a globally irrelevant set of variables with respect to φ.*

Corollary 1. *The pre-process stage in Syn2FOND prunes a set of globally irrelevant variables in $\mathcal{X} \cup \mathcal{Y}$ wrt φ.*

An immediate consequence of Theorem 5 – which follows trivially from the observation that variables in $\mathcal{Z}$ do not appear in the specification φ – is that the pruning of (the globally irrelevant set of) variables $\mathcal{Z} = (\mathcal{X} \cup \mathcal{Y}) \setminus vars(\varphi)$ performed by the Syn2FOND compilation does not affect the correctness of the approach.

As we illustrate in the experimental section, reducing the number of irrelevant sets of variables in the problem makes it, in general, more tractable, and not doing so may induce problems of scalability. Commonly used techniques in automated planning reason about goal reachability by exploiting the compact description of world change based on the preconditions and effects of the actions. They prune variables (and also actions) that are not relevant to the achievement of the goal (cf. (Helmert 2006; Palacios and Geffner 2007)). By solving a synthesis problem as planning, we can benefit from variable pruning techniques existing in automated planning tools.

Local Irrelevance The dynamics of the Syn2FOND compilation, presented in Section 4, make it possible to prune what we can intuitively consider to be locally irrelevant variables at planning time. In particular, when the agent decides to apply a $trans(T)$ action, only variables in $guard(T)$ are implicitly considered to be relevant. This has two benefits. First, by computing plans that only consider a subset of the uncontrollable variables in $\mathcal{X}$, the planner reduces the exponential blowup (in $|\mathcal{X}|$) in the search caused by all possible environment's moves. Second, by assigning values to a reduced subset of the controllable variables in $\mathcal{Y}$ (only those that are necessary to transition in the NBA accordingly), the planner is able to compute more compact policies that are conditioned on only what is relevant.

State Regularization The state regularization performed by the *switch2env* and *accept* actions leverage the fact that, at the next planning step, the value of the variables $v_x, v_{\neg x}, v_y, v_{\neg y}$ (that simulate the variables in $\mathcal{X}$ and $\mathcal{Y}$) is not relevant. By abstracting all planning states, the search space is notably reduced. Besides savings in search space, this also reduces the planning run time, because it is more likely for a strong-cyclic planner to find a loop in the abstract space than in the concrete space.

5 Evaluation

Our main objective for the evaluation is to give a sense of when to choose one formalism over another. Although the same problems can be represented as either LTL synthesis or FOND planning, the choice of formalism can have a dramatic impact on the time required to find a solution. We would expect the FOND setting to be better suited for problems with more "structure", and our results serve to illustrate this hypothesis.

We consider four natural sources of problem encodings: (1) problems encoded directly as LTL synthesis using the TLSF format (Jacobs *et al.* 2016); (2) problems encoded directly as FOND planning using the PDDL format; (3) problems from (2) that have been converted to the LTL synthesis setting following the proposal in (De Giacomo and Vardi 2013) modified to better capture the FOND setting; and (4) problems from (1) that have been converted automatically using the method described in Section 4. In our experiments, we used state-of-the-art synthesis and FOND tools Acacia+ (Bohy *et al.* 2012) and PRP (Muise *et al.* 2012). Acacia+ can be configured to return either all solutions or a single one, and for our experiments we use the more efficient latter option. In our Syn2FOND tool, we use `spot` (Duret-Lutz *et al.* 2016) to transform the specification φ into an NBA and PRP as FOND planner.

First, we consider some representative problems from both synthesis and FOND perspectives. The first group of problems – **lily** and **loadcomp** (*load*, for short) – come from the synthesis community. The **lily** problems are a variety of demo benchmarks from the testing set of the synthesis tool *Lily* (Jobstmann and Bloem 2006). The **load** problems, from the testing set of the LTL synthesis tool *Unbeast* (Ehlers 2011), implement a load balancer that receives jobs and distributes them to a (parametrized) number of servers.

The second group of problems – **ttw** and **ctw** – come from the FOND benchmark tireworld. The tireworld domain requires a car to drive along a predefined map forming a triangle (**ttw**) or a chain (**ctw**). The car can only move if the tire is working. Driving non-deterministically breaks a tire (i.e., the environment's play), and certain key locations have a spare tire that can be used to repair the car. We engineered a compact LTL specification that models this group of problems. For instance, the dynamics of *ttw-p3* are captured in

	Acacia$^+$		Syn2FOND(PRP)			PRP
Problem	Aut ($N \times \lvert Q\rvert$)	Syn	H	Aut ($\lvert Q\rvert$)	Search	Search
lilly-p4	0.11 (3×30)	0.02	3	0.02 (20)	0.6	N/A
lilly-p5	0.14 (3×26)	0.01	3	0.02 (17)	0.38	N/A
lilly-p6	0.16 (3×34)	0.01	3	0.02 (22)	0.62	N/A
lilly-p7	0.15 (3×24)	0.00	3	0.02 (13)	0.14	N/A
load-p2	0.39 (8×35)	0.01	3	0.02 (9)	0.14	N/A
load-p3	0.15 (12×37)	0.30	3	0.05 (10)	0.18	N/A
load-p4	0.39 (18×37)	7.83	3	0.02 (11)	0.26	N/A
ctw-p3	81.0 (1×82)	1.79	3	1.08 (15)	0.12	0.01
ctw-p4	TLE	–	4	25.8 (19)	0.22	0.01
ctw-p5	TLE	–	–	MLE	–	0.01
ttw-p3	TLE	–	–	MLE	–	0.01
build-2	0.05 (2×70)	0.04	1	0.02 (6)	1.82	
build-3	0.42 (2×415)	3.45	1	0.13 (7)	13.3	0.08
build-4	56.2 (2×2682)	1.65	1	2.01 (8)	160	1.06
build-irr-4-2	60.2 (2×3202)	5.18	–	9.05 (8)	MLE	0.32
build-irr-4-3	113 (2×3202)	14.4	–	20.2 (8)	MLE	1.24
build-irr-4-4	212 (2×3202)	50.9	–	44.9 (8)	MLE	3.52

Table 1: Performance of LTL synthesis and FOND planning tools. (Aut) Time (in seconds) to compile the specification into automata. (N) size of automata and ($|Q|$) average size of each automaton. (Syn) Time (in seconds) for synthesis. (Search) Total time taken during search for a solution in PRP (including bookkeeping and policy maintenance time). MLE - Memory Limit Exceeded; TLE - Time Limit Exceeded.

about 30 lines – where each line contains a cube describing φ_{init}, φ_{env}, φ_{agt}, and φ_g. This description scales linearly with the size of the problem, whereas the domain description in PDDL is more compact and remains constant.

The final set of problems is a newly introduced domain, **build**, encoded directly (and as compactly as possible) in both LTL and FOND. The **build** domain addresses the problem of building maintenance, and requires the agent to maintain which rooms have their lights on or off depending on the time of day and whether or not people are in the room. The environment controls the transition between day and night, as well as when people enter or leave a room. The agent must control the lights in response. Problems **build-p#** have # rooms, and problems **build-irr-p#-n** introduce n rooms that may non-deterministically be vacuumed at night (controlled by the environment). Note that this is irrelevant to the task of turning lights on or off (the vacuuming can be done in any lighting condition).

Table 1 summarizes the results of our experiments with Acacia$^+$ and our synthesis tool Syn2FOND equipped with PRP as strong cyclic planner. Additionally, we tested PRP in some problems directly encoded as FOND. The first thing to note is the drastic performance hit that can occur converting from one formalism to another. Going from LTL synthesis to FOND is workable in some instances, but the opposite direction proved impossible for even the simplest **ttw** problems that can be solved very efficiently in its native FOND formulation. Automata transformations become a bottleneck in the synthesis tools, causing time (TLE, 30 min) and memory (MLE, 512MB) limit exceptions. This is because the specification requires complex constraints to properly maintain the reachable state-space. It is this "structure" that we conjecture the synthesis tools struggle with, and test separately below.

	Acacia$^+$		Syn2FOND(PRP)			PRP
Problem	Aut ($N \times \lvert Q\rvert$)	Syn	H	Aut ($\lvert Q\rvert$)	Search	Search
p4-0	212 (6 × 131)	0.01	3	0.16 (13)	3.1	0
p4-1	11.1 (6 × 195)	0.03	3	0.17 (13)	2.56	0.32
p4-2	1.55 (6 × 181)	0.02	3	0.16 (13)	1.26	0.66
p4-3	0.58 (6 × 46)	0.02	3	0.1 (12)	1.04	0.22
p5-0	TLE	—	3	1.33 (15)	1.02	0
p5-1	TLE	—	3	1.59 (15)	0.88	3.3
p5-2	16991 (7 × 548)	0.02	3	1.25 (15)	0.94	2.14
p5-3	533 (7 × 452)	0.08	3	1.34 (15)	15.8	3.82
p5-4	111 (7 × 236)	0.06	3	0.59 (14)	10.22	4.66
p6-0	TLE	—	3	9.11 (17)	1.9	0
p6-1	TLE	—	3	11.8 (17)	1.8	4.9
p6-2	TLE	—	3	11.4 (17)	1.78	5.54
p6-3	TLE	—	3	10.7 (17)	1.52	15.66
p6-4	TLE	—	3	10.4 (17)	3.66	25.44
p6-5	TLE	—	3	6.15 (16)	3.32	26.18

Table 2: Performance of LTL synthesis and FOND planning tools for the **switches** problems.

For the **build-irr** domain, we can see that the synthesis tools scale far worse as the number of rooms increases (both in generating automata and performing the synthesis). Further, as the number of rooms that need to be vacuumed (which is irrelevant to computing a controller), the relevance reasoning present in the FOND planner is able to cope by largely ignoring the irrelevant aspects of the environment. Conversely, the synthesis component of Acacia+ struggles a great deal. This highlights the strength of the FOND tools for leveraging state relevance to solve problems efficiently.

Finally, we created a synthetic domain that lets us tune the level of "structure" in a problem: more structure leads to fewer possibilities for the environment to act without violating a constraint or assumption. In the **switches** domain, a total of n switches, $s_1 \ldots s_n$, initially switched on, need to be all switched off eventually. The environment affects the state of the switches non-deterministically. However, the dynamics of the environment is such that immediately after the agent switches off s_k, the environmental non-determinism can only affect the state of a certain number of switches $s_{k'}$, with $k' > k$. A trivial strategy for the agent is to switch off s_1 to s_n in that order.

We encoded a series of **switches** problems, natively as LTL specifications in TLSF format and also as FOND. Table 2 shows how Acacia+, Syn2FOND, and PRP coped with the range of problems. They are in three distinct sets (each of increasing number of switches), and within each group the problems range from most structured to least (by varying k). The problems in the first two groups are solved quite

readily by PRP, and so the trend is less clear, but for the larger problems we find that PRP struggles when there is less structure and the environment can flip many switches without violating an assumption. This trend also manifests in the Syn2FOND compilations. On the other side, we find that the most structured problems are the most difficult for Acacia+, and the compilation into automata becomes the bottleneck again. On the other hand, the synthesis becomes easier when there is less structure (i.e., more switches can be flipped).

The structure we tune in the **switches** domain is one property of a problem that may favour FOND technology over the synthesis tools. Another is the presence of irrelevance we discuss in Section 4.2 and surfaced in the variation of the building maintenance benchmark. Other notions, such as causal structure in the problem, may also play an important role in features that separate the effectiveness of the two formalisms. We plan to investigate these possibilities in future work.

6 Concluding Remarks

LTL synthesis is both an important and challenging problem for which broadly effective tools have remained largely elusive. Motivated by recent advances in the efficiency of FOND planning, this work sought to examine the viability of FOND planning as a computational tool for the realization of LTL synthesis. To this end, we established the theoretical correspondence between LTL synthesis and strong solutions to FOND planning. We also provided the first approach to automatically translate a realizability problem, given by an specification in LTL or LTL_f, into a planning problem described in PDDL. Experiments with state-of-the-art LTL synthesis and FOND solvers highlighted properties that challenged or supported each of the solvers. Our experiments show automated planning to be a viable and effective tool for highly structured LTL synthesis problems.

Acknowledgements The authors gratefully acknowledge funding from the Natural Sciences and Engineering Research Council of Canada (NSERC) and from Fondecyt grant number 1150328.

References

Ron Alford, Ugur Kuter, Dana Nau, and Robert P Goldman. Plan aggregation for strong cyclic planning in nondeterministic domains. *Artificial Intelligence*, 216:206–232, 2014.

Rajeev Alur and Salvatore La Torre. Deterministic generators and games for LTL fragments. *ACM Trans. Comput. Logic*, 5(1):1–25, January 2004.

Eugene Asarin, Oded Maler, Amir Pnueli, and Joseph Sifakis. Controller synthesis for timed automata. In *Proceedings of the IFAC Symposium on System Structure and Control*, pages 469–474. Elsevier, 1998.

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

Fahiem Bacchus and Froduald Kabanza. Using temporal logics to express search control knowledge for planning. *Artificial Intelligence*, 116(1-2):123–191, 2000.

Jorge A. Baier and Sheila A. McIlraith. Planning with first-order temporally extended goals using heuristic search. In *Proceedings of the 21st National Conference on Artificial Intelligence (AAAI)*, pages 788–795, Boston, MA, July 2006.

Jorge A. Baier and Sheila A. McIlraith. Planning with temporally extended goals using heuristic search. In *Proceedings of the 16th International Conference on Automated Planning and Scheduling (ICAPS)*, pages 342–345, 2006.

Jorge A. Baier, Fahiem Bacchus, and Sheila A. McIlraith. A heuristic search approach to planning with temporally extended preferences. *Artificial Intelligence*, 173(5-6):593–618, 2009.

Meghyn Bienvenu, Christian Fritz, and Sheila A. McIlraith. Specifying and computing preferred plans. *Artificial Intelligence*, 175(7–8):1308–1345, 2011.

Aaron Bohy, Véronique Bruyère, Emmanuel Filiot, Naiyong Jin, and Jean-François Raskin. Acacia+, a tool for LTL synthesis. In *Proceedings of the 24th International Conference on Computer Aided Verification (CAV)*, pages 652–657, 2012.

Alberto Camacho, Eleni Triantafillou, Christian Muise, Jorge A. Baier, and Sheila A. McIlraith. Non-deterministic planning with temporally extended goals: LTL over finite and infinite traces. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI)*, 2017.

Alonzo Church. Applications of recursive arithmetic to the problem of circuit synthesis. *Summaries of the Summer Institute of Symbolic Logic, Cornell University 1957*, 1:3–50, 1957.

Amanda Coles and Andrew Coles. LPRPG-P: relaxed plan heuristics for planning with preferences. In *Proceedings of the 21st International Conference on Automated Planning and Sched. (ICAPS)*, 2011.

Giuseppe De Giacomo and Moshe Y. Vardi. Linear temporal logic and linear dynamic logic on finite traces. In *Proceedings of the 23rd International Joint Conference on Artificial Intelligence (IJCAI)*, 2013.

Giuseppe De Giacomo and Moshe Y. Vardi. Synthesis for LTL and LDL on finite traces. In *Proceedings of the 24th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 1558–1564, 2015.

Giuseppe De Giacomo, Paolo Felli, Fabio Patrizi, and Sebastian Sardiña. Two-player game structures for generalized planning and agent composition. In *Proceedings of the 24th AAAI Conference on Artificial Intelligence (AAAI)*, 2010.

Alexandre Duret-Lutz, Alexandre Lewkowicz, Amaury Fauchille, Thibaud Michaud, Etienne Renault, and Laurent Xu. Spot 2.0 — a framework for LTL and ω-automata manipulation. In *Proceedings of the 14th International Symposium on on Automated Technology for Verification and Analysis (ATVA)*, volume 9938 of *Lecture Notes in Computer Science*, pages 122–129. Springer, October 2016.

Stefan Edelkamp. On the compilation of plan constraints and preferences. In *Proceedings of the 16th International Conference on Automated Planning and Scheduling (ICAPS)*, pages 374–377, 2006.

Rüdiger Ehlers. Unbeast: Symbolic bounded synthesis. In *Proceedings of the 17th International Conference on Tools and Algorithms for the Construction and Analysis of Systems (TACAS)*, pages 272–275, 2011.

Jicheng Fu, Vincent Ng, Farokh B. Bastani, and I-Ling Yen. Simple and fast strong cyclic planning for fully-observable nondeterministic planning problems. In *Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI)*, pages 1949–1954, 2011.

Hector Geffner and Blai Bonet. A Concise Introduction to Models and Methods for Automated Planning. *Synthesis Lectures on Artificial Intelligence and Machine Learning*, 7(2):1–141, 2013.

Patrik Haslum, Malte Helmert, and A Jonsson. Safe, strong and tractable relevance analysis for planning. In *Proceedings of the 23nd International Conference on Automated Planning and Sched. (ICAPS)*, pages 317–321, 2013.

Malte Helmert. The Fast Downward planning system. *Journal of Artificial Intelligence Research*, 26:191–246, 2006.

Swen Jacobs, Felix Klein, and Sebastian Schirmer. A high-level LTL synthesis format: TLSF v1.1. In *Proceedings Fifth Workshop on Synthesis, SYNT@CAV 2016, Toronto, Canada, July 17-18, 2016.*, pages 112–132, 2016.

Barbara Jobstmann and Roderick Bloem. Optimizations for LTL synthesis. In *Formal Methods in Computer-Aided Design, 6th International Conference, FMCAD 2006, San Jose, California, USA, November 12-16, 2006, Proceedings*, pages 117–124, 2006.

Peter Kissmann and Stefan Edelkamp. Solving fully-observable non-deterministic planning problems via translation into a general game. In *Proceedings of the 32nd Annual German Conference on AI (KI09)*, pages 1–8, 2009.

Robert Mattmüller, Manuela Ortlieb, Malte Helmert, and Pascal Bercher. Pattern database heuristics for fully observable nondeterministic planning. In *Proceedings of the 20th International Conference on Automated Planning and Sched. (ICAPS)*, pages 105–112, 2010.

Christian Muise, Sheila A. McIlraith, and J. Christopher Beck. Improved Non-deterministic Planning by Exploiting State Relevance. In *Proceedings of the 22nd International Conference on Automated Planning and Sched. (ICAPS)*, pages 172–180, 2012.

Hector Palacios and Hector Geffner. From Conformant into Classical Planning: Efficient Translations that May Be Complete Too. In *Proceedings of the 17th International Conference on Automated Planning and Scheduling (ICAPS)*, pages 264–271, 2007.

Fabio Patrizi, Nir Lipovetzky, and Hector Geffner. Fair LTL synthesis for non-deterministic systems using strong cyclic planners. In *Proceedings of the 23rd International Joint Conference on Artificial Intelligence (IJCAI)*, 2013.

Nir Piterman, Amir Pnueli, and Yaniv Sa'ar. Synthesis of reactive(1) designs. In *Verification, Model Checking, and Abstract Interpretation, 7th International Conference, VMCAI 2006, Charleston, SC, USA, January 8-10, 2006, Proceedings*, pages 364–380, 2006.

Amir Pnueli and Roni Rosner. On the synthesis of a reactive module. In *Conference Record of the Sixteenth Annual ACM Symposium on Principles of Programming Languages, Austin, Texas, USA, January 11-13, 1989*, pages 179–190, 1989.

Amir Pnueli. The temporal logic of programs. In *Proceedings of the 18th IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 46–57, 1977.

Raymond Reiter. *Knowledge in Action: Logical Foundations for Specifying and Implementing Dynamical Systems*. MIT Press, Cambridge, MA, 2001.

Jussi Rintanen, Keijo Heljanko, and Ilkka Niemelä. Planning as satisfiability: parallel plans and algorithms for plan search. *Artificial Intelligence*, 170(12-13):1031–1080, 2006.

Sebastian Sardiña and Nicolás D'Ippolito. Towards fully observable non-deterministic planning as assumption-based automatic synthesis. In *Proceedings of the 24th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 3200–3206, 2015.

Moshe Y. Vardi and Pierre Wolper. Reasoning about infinite computations. *Information and Computation*, 115(1):1–37, 1994.