A New Algorithm for Strategy Synthesis in LTL Games

Aidan Harding[1], Mark Ryan[1], and Pierre-Yves Schobbens[2]

[1] School of Computer Science, The University of Birmingham, Edgbaston,
Birmingham B15 2TT, UK
[2] Institut d'Informatique, Facultés Universitaires de Namur, Rue Grandgagnage 21,
5000 Namur, Belgium

Abstract. The automatic synthesis of programs from their specifications has been a dream of many researchers for decades. If we restrict to open finite-state reactive systems, the specification is often presented as an ATL or LTL formula interpreted over a finite-state game. The required program is then a strategy for winning this game. A theoretically optimal solution to this problem was proposed by Pnueli and Rosner, but has never given good results in practice. This is due to the 2EXPTIME-complete complexity of the problem, and the intricate nature of Pnueli and Rosner's solution. A key difficulty in their procedure is the determinisation of Büchi automata. In this paper we look at an alternative approach which avoids determinisation, using instead a procedure that is amenable to symbolic methods. Using an implementation based on the BDD package CuDD, we demonstrate its scalability in a number of examples. Furthermore, we show a class of problems for which our algorithm is singly exponential. Our solution, however, is not complete; we prove a condition which guarantees completeness and argue by empirical evidence that examples for which it is not complete are rare enough to make our solution a useful tool.

1 Introduction

Finite-state reactive systems occur in many critical areas of computing. They can be found in places such as network communication protocols, digital circuits, and industrial control systems. Their use in systems which involve concurrency and their interaction with unpredictable or hostile environments makes reactive systems notoriously hard to write correctly. By considering such systems as games we can distinguish between events that we can control (inside the program) and events that we cannot (the environment). This gives a more realistic framework for reasoning about them than the conventional approach of "closing" the system by adding a restricted environment and treating all choices uniformly.

We take the stance that closing an open system for verification or synthesis is imprecise and that reasoning with game semantics provides a much better solution. This stance has been advocated by many other researchers [1, 9, 10, 8], but there are some verification and synthesis problems that become much

harder in the game-theoretic world. In particular the problems of synthesis and verification for games with LTL winning conditions are 2EXPTIME-complete [12, 13]. This high complexity and the intricacy of the solution offered by Pnueli and Rosner have meant that despite the wealth of potential applications, there has been no implementation of synthesis for LTL games. We address this problem by providing a novel algorithm which avoids a major difficulty of the classical approach: the determinisation of Büchi automata. The best known method for this determinisation is due to Safra [14] and this method has been proven to be optimal [11] but has resisted efforts at efficient symbolic implementation [17]. Instead of trying to determinise a Büchi automaton, our algorithm uses a "shift automaton" to track the possible states that the Büchi automaton could be in and retake non-deterministic choices if they turn out to be wrong. The shift automaton is of roughly equal size to the deterministic automaton produced by Safra's algorithm, but it can be constructed symbolically. This has allowed for the construction of an efficient implementation using BDDs. In this paper we describe in detail a new algorithm for the synthesis of strategies in LTL games; we describe some small problems that can be solved by the implementation of this algorithm; and finally we give some performance data obtained by parameterising the given examples.

2 ω-Automata and Infinite Games

We quickly review the definitions and establish notations for writing about ω-automata and infinite games. Detailed information on ω-automata can be found in [18]; and information on infinite games can be found from [8] and [19].

Given an alphabet Σ, we denote the set of all finite words made from letters in Σ as Σ^*, and the set of all ω-words (infinte words) as Σ^ω. For a word $\lambda \in \Sigma^\omega$, we write $\lambda[i]$ for the i-th letter, $\lambda[i, j]$ for the finite section of the word from point i to j, and $\lambda[i, \infty]$ for the infinte suffix from point i. ω-automata provide a way of recognising sets of ω-words. An ω-automaton $\mathcal{A} = \langle Q_\mathcal{A}, i_\mathcal{A}, \delta_\mathcal{A}, Acc \rangle$ is a tuple where the component parts are as follows: $Q_\mathcal{A}$ is a finite set of states; $i_\mathcal{A}$ is an initial state; $\delta_\mathcal{A} : Q_\mathcal{A} \times \Sigma \mapsto 2^{Q_\mathcal{A}}$ is a transition function (we may define $\delta_\mathcal{A} : Q_\mathcal{A} \times \Sigma \mapsto Q_\mathcal{A}$ for deterministic automata); and Acc is an acceptance condition. A run ρ of an ω-automaton on a word λ is an infinite sequence of states such that $\rho[0] = i_\mathcal{A}$ and for all $i \geq 0$ $\rho[i + 1] \in \delta_\mathcal{A}(\rho[i], \lambda[i])$. We denote the set of states that occur infintely often on a run ρ by $In(\rho)$. In this paper we are concerned with two types of ω-automata: Büchi automata and Rabin automata. We write DB for a deterministic Büchi automaton and NB for a non-deterministic Büchi automaton. The acceptance condition in a Büchi automaton is a set $F \subseteq Q_\mathcal{A}$ and a word λ is accepted if and only if there is a run ρ on λ such that $In(\rho) \cap F \neq \emptyset$. The acceptance condition on a Rabin automaton is a set of pairs $\{(E_0, F_0), \ldots, (E_n, F_n)\}$ and a word λ is accepted if and only if there is a run ρ on λ such that there exists $i \in [0, n]$ such that $In(\rho) \cap F_i \neq \emptyset$ and $In(\rho) \cap E_i = \emptyset$.

An infinite game is a tuple $G = \langle Q_G, i_G, Q_P, Q_A, \delta_G \rangle$ representing a two-player game between the protagonist P and the antagonist A. Q_G is a set of states from which i_G is the initial one; Q_P and Q_A partition Q_G into turns for P and A, respectively; $\delta_G : Q_G \mapsto 2^{Q_G}$ is a transition function such that $\forall q \in Q_P \ \delta_G(q) \subseteq Q_A$ and $\forall q \in Q_A \ \delta_G(q) \subseteq Q_P$ i.e. the players alternate turns. A play of the game is an infinite sequence λ of states from Q_G such that $\lambda[0] = i_G$ and for all $i \geq 0 \ \lambda[i+1] \in \delta_G(\lambda[i])$. We formalise the capabilites of players in the game with the notion of strategies. A strategy $f : Q_G^+ \mapsto 2^{Q_G}$ for P restricts the choices of P by prescribing how he should play his moves. We require that for any play λ and all $i \geq 0 \ f(\lambda[0,i]) \subseteq \delta_G(\lambda[i])$. The set of outcomes $out(f, q_0)$ of playing a strategy f from a state q_0 is defined as

$$out(f, q_0) = \{q_0, q_1, \ldots \mid \forall i \geq 0 \text{ if } (q_i \in Q_P) \ q_{i+1} \in f(q_0, \ldots, q_i)$$
$$\text{else } q_{i+1} \in \delta_G(q_i)\}$$

We also use partial strategies which are partial functions with the same type as a normal strategy. The set of outcomes of a partial strategy is defined as:

$$out(f, q_0) = \{q_0, q_1, \ldots \mid \forall i \geq 0 \text{ if } (q_i \in Q_P \wedge f(q_0, \ldots, q_i) \text{ defined})$$
$$q_{i+1} \in f(q_0, \ldots, q_i) \text{ else } q_{i+1} \in \delta_G(q_i)\}$$

When a game is provided with a winning condition W, we say that P can win the game if and only if there is a strategy (partial or complete) for P such that all outcomes of the strategy satisfy W.

3 Synthesis for NB Games

The main algorithm in our synthesis procedure takes as input a game $G = \langle Q_G, i_G, Q_P, Q_A, \delta_G \rangle$ and a winning condition in the form of a Büchi automaton, $B = \langle Q_B, i_B, \delta_B, F_B \rangle$. The algorithm identifies a set of states that are winning for P and produces a partial strategy (partial because it may be undefined from states which are not winning for P) such that the set of outcomes of the strategy from any winning state are accepted by B.

To see why the conventional approach uses determinisation, let us consider an approach which uses an obvious extension of the algorithm used for games with Büchi winning conditions [19]. A Büchi winning condition (as opposed to a winning condition specified by a NB) specifies a set $F_G \subseteq Q_G$ such that plays are winning for P if and only if they visit F_G infinitely often. The solution for these games offered in [19] is a game-theoretic extension of the algorithm proposed by Emerson and Lei for finding fair strongly connected components [7]. This algorithm is attractive as it has shown to be quite efficient when compared against other symbolic fair cycle detection algorithms [16]. It works by finding the set of states from where P has a strategy to reach F_G, then the states where P has a strategy to reach F_G and from there has a strategy to reach F_G again etc. An increasing number of visits to F_G are required until a fixed-point is achieved whereupon we know that P has a strategy to visit F_G infinitely often.

The game-theoretic aspect of the algorithm in [19] works by adapting the notion of predecessor to enforce that a state in Q_P is winning if there exists a winning successor, and a state in Q_A is winning if all successors are winning.

We could try to perform the same computation over $G \times B$, always evaluating the B component existentially. However, there is an assumption in this algorithm that winning is transitive i.e. the states along a winning path are winning themselves. Whilst this may seem like a natural thing to expect, it is not actually true. The game and specification in Figure 1 give an example of how winning can fail to be transitive in this sense. P has only one choice which comes at state 3; he can win this game by always choosing $3 \to 4$ if he gets that choice. Although there is a winning strategy from $(1, t_0)$, if we follow that strategy and, at the same time, try to construct a winning run from the B part, we cannot be sure to reach a state from where there is another winning strategy. The opponent can stay in $\{(1, t_?), (2, t_?)\}$ as long as he likes and we must choose what to do with the Büchi component. We cannot allow the Büchi automaton to visit t_1 in case the opponent later chooses 3. So we either have a losing run in the Büchi component, $((1, t_0)(2, t_0))^\omega$ or reach $\{(1, t_1), (2, t_1)\}$ from where there is no winning strategy (A chooses 3 and B is stuck). On this basis, $(1, t_0)$ would not be identified as winning because P cannot be sure to reach a state in $Q_G \times F_B$ from where he has a winning strategy.

(a) Game   (b) Winning Condition
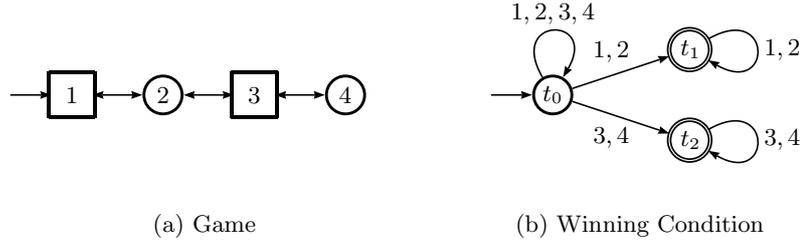
Fig. 1. A game where the winning condition "needs" determinisation. Square states are P's moves, circles are A's.

Our solution to this is to allow some "shifting" between Büchi states. If the Büchi automaton is in a dead-end state, we now allow the transition relation to be overridden by making a shift i.e make a transition as if the Büchi automaton were in a different state. We maintain the set of reachable Büchi states at all times, and this provides the justification for shifts – whenever a shift is made, it is made to some reachable state and, thus, is equivalent to retaking some earlier non-deterministic choices. The set of reachable states is provided by the shift automaton, $S = \langle Q_S, i_S, \delta_S \rangle$, a deterministic automaton derived from B with the subset construction where each state in Q_S represents a set of states from Q_B:

$$ Q_S = \mathcal{P}(Q_B) \qquad i_S = \{i_B\} \qquad \delta_S(\phi, q) = \{t' \mid \exists t \in \phi . t' \in \delta_B(t, q)\} \qquad (1) $$

Accordingly, we define the synthesis algorithm over $G \times B \times S$. We call this product the *composite*, C, and define some shorthands to ease the burden of

notation:

$$Q_C = Q_G \times Q_B \times Q_S \tag{2}$$

$$F_C = \{(q, t, \phi) \in Q_C \mid t \in F_B\} \tag{3}$$

$$Q_{CP} = \{(q, t, \phi) \in Q_C \mid q \in Q_{GP}\} \tag{4}$$

$$Q_{CA} = \{(q, t, \phi) \in Q_C \mid q \in Q_{GA}\} \tag{5}$$

Now, if we are in a situation such as the one in Figure 1, we can make winning transitive by allowing some shifts. The informal argument is as follows: From $(1, t_0, \{t_0\})$ we always go to $(2, t_1, \{t_0, t_1\})$ because we are optimistic about getting an accepting run in the B component. If the opponent always chooses $2 \to 1$, then we have an accepting run made up of $(1, t_0, \{t_0\})$ $((2, t_1, \{t_0, t_1\})$ $(1, t_0, \{t_0, t_1\}))^\omega$. If the opponent eventually chooses $2 \to 3$, then we take a shift: In state $(2, t_1, \{t_0, t_1\})$ the B component could have been in t_0, so when the opponent chooses 3 we make a shift and take the $t_0 \to^3 t_2$ transition i.e. we go to $(3, t_2, \{t_0, t_2\})$. From here, P can win and we end up with an overall winning path made up by $(1, t_0, \{t_0\})((2, t_1, \{t_0, t_1\})(1, t_0, \{t_0, t_1\}))^*(2, t_1, \{t_0, t_1\})$ $((3, t_2, \{t_0, t_2\}) (4, t_2, \{t_0, t_2\}))^\omega$.

Shifting helps the issue of completeness, but allowing an infinite number of shifts would be unsound. However, if shifting is only allowed finitely often, the language is not changed. Informally we justify this on the basis that acceptance is evaluated over infinite paths, and although shifting may allow finitely many extra visits to F_C, paths must eventually have no more shifts and thus would be accepting without any shifting. The soundness of finite shifting is implied by Theorem 2.

To write down the main synthesis algorithm with finite shifting, we first define two predecessor functions pre_P and pre_A. These are evaluated over the triple state-space of $G \times B \times S$, respecting the alternation of the game and allowing for shifting. Unlike a conventional predecessor function, two arguments are supplied. The second argument is a set that we allow shifts into. $pre_P(X, W)$ is the set of transitions which obey the game and shift automaton, and either have a transition in B to reach X ($t' \in \delta_B(t, q') \land (q', t', \phi') \in X$ in Equation 6) or have a shift justified by the shift automaton to reach W ($t' \in \phi' \land (q', t', \phi') \in W$ in Equation 6). $pre_A(X, W)$ simply uses $pre_P(X, W)$ as an approximation, and then makes sure that there is a good transition for every possible game-successor.

$$
\begin{aligned}
pre_P(X, W) = \{\langle(q, t, \phi), (q', t', \phi')\rangle \mid q' \in \delta_G(q), \phi' = \delta_S(\phi, q'), \\
(t' \in \delta_B(t, q') \land (q', t', \phi') \in X) \lor (t' \in \phi' \land (q', t', \phi') \in W)\}
\end{aligned} \tag{6}
$$

$$
\begin{aligned}
pre_A(X, W) = \{\langle(q, t, \phi), (q', t', \phi')\rangle \in pre_P(X, W) \mid \forall q'_2 \in \delta_G(q) \\
\exists \langle(q, t, \phi), (q'_2, t'_2, \phi'_2)\rangle \in pre_P(X, W)\}
\end{aligned} \tag{7}
$$

Using these definitions, we write the main algorithm in Figure 2[1]. To understand how the synthesis algorithm works, consider each of the variables in turn:

---

[1] We denote the k-th projection of a tuple T by $\pi_k(T)$

```
1    w_0 := ∅;
2    δ_{F,0,∞,∞} := ∅;
3    repeat counted by j = 1 . . .
4        z_{j,0} := Q_C;
5        repeat counted by k = 1 . . .
6            τ_{j,k} := z_{j,k−1} ∩ F_C;
7            s_{j,k,0} := ∅;
8            δ_{Fj,k,0} := δ_{Fj−1,∞,∞};
9            repeat counted by l = 1 . . .
10               u_A := pre_A(τ_{j,k} ∪ s_{j,k,l−1}, w_{j−1}) ∩ (Q_{CA} × Q_{CV});
11               u_V := pre_V(τ_{j,k} ∪ s_{j,k,l−1}, w_{j−1}) ∩ (Q_{CV} × Q_{CA});
12               δ_{Fj,k,l} := δ_{Fj,k,l−1} ∪ {⟨(q, t, φ), (q′, t′, φ′)⟩ ∈ u_A ∪ u_V
                                   | (q, t, φ) ∉ π_1(δ_{Fj,k,l−1})};
13               s_{j,k,l} := s_{j,k,l−1} ∪ π_1(u_A) ∪ π_1(u_V);
14           until s_{j,k,l} = s_{j,k,l−1}
15           z_{j,k} := z_{j,k−1} ∩ s_{j,k,∞};
16       until z_{j,k} = z_{j,k−1}
17       w_j := w_{j−1} ∪ z_{j,∞};
18   until w_j = w_{j−1}
```

Fig. 2. Synthesis algorithm with finite shifting

- w_j: At the end of the algorithm, this will contain the set of winning states. The j subscript is the maximum number of shifts required to win from a state in w_j.
- $z_{j,k}$: At the end of the middle loop, this is the set of states from where every outcome reaches $z_{j,k} \cap F_C$ infinitely often with no shifting or just reaches $w_{j−1}$ (possibly by shifting). During the middle loop, every outcome reaches $z_{j,k−1} \cap F_C$ with no shifting or $w_{j−1}$ (possibly by shifting).
- $\tau_{j,k}$: The "target" for the innermost loop. This variable could be substituted for its definition at each use, it is clearer (and more efficient in implementation) to write separately.
- $s_{j,k,l}$: The set of states from where P can be sure to reach $\tau_{j,k}$ in l steps with no shifting or $w_{j−1}$ in l steps with a shift.
- $\delta_{F,j,k,l}$: The partial strategy as it is synthesised. On the first j-loop it will be a strategy to win with no shifting. On line 12 we must be careful not to overwrite old moves. On iteration l of the inner loop, when a transition is first added to the strategy it must go into $s_{j,k,l−1} \cup \tau_{j,k}$ or $w_{j−1}$. However, this state will be rediscovered on later iteration of the inner loop and u_P/u_A may contain transitions which do not make progress towards an accepting state and we must therefore keep the transition from the first discovery. Having built a strategy with no shifts, we carry this forward to the next iteration of the j-loop. Here another strategy is built up, but this time it allows the possibility of a shift to $w_{j−1}$ since we already have a winning strategy from there. New moves are written for states in $w_j − w_{j−1}$, but as soon as the strategy reaches $w_{j−1}$ the old strategy takes over.

The algorithm in Figure 2 is computing nested fixed-points which can be characterised as $\mu w.\nu z.\mu s.\pi_1((pre_P(s \vee (z \wedge F_C), w) \wedge (Q_{CP} \times Q_{CA})) \vee (pre_A(s \vee (z \wedge F_C), w) \wedge (Q_{CA} \times Q_{CP}))$. We note that on the first iteration of the outer loop in this algorithm, the computation performed is the naïve extension of the solution for Büchi games i.e. $\nu z.\mu s.\pi_1((pre_P(s \vee (z \wedge F_C), \perp) \wedge (Q_{CP} \times Q_{CA})) \vee (pre_A(s \vee (z \wedge F_C), \perp) \wedge (Q_{CA} \times Q_{CP}))$. This calculation does not depend on the shift automaton and, in this way, we are sometimes able to perform the strategy synthesis without having to generate the shift automaton. We can give a precise condition which assures this by first defining trivially determinisable Büchi automata.

Definition 1. *A Büchi automaton is trivially determinisable if and only if it can be made deterministic by removing 0 or more transitions without changing its language.*

Using this definition it is possible to prove the following theorem:

Theorem 1. *For any game G, with a winning condition specified by a Büchi automaton B, if B is trivially determinisable, then all winning states for P in Q_C satisfy $\nu z.\mu s.\pi_1((pre_P(s \vee (z \wedge F_C), \perp) \wedge (Q_{CP} \times Q_{CA})) \vee (pre_A(s \vee (z \wedge F_C), \perp) \wedge (Q_{CA} \times Q_{CP}))$.*

Whilst this definition includes the shift automaton, it is clear that the predecessor functions do not depend on the shift automaton when W is empty, so this proves that if B is trivially determinisable, the algorithm can succeed without generating the shift automaton. Intuitively, this theorem holds because the naïve algorithm is complete for deterministic Büchi automata and since the transitions of B are evaluated existentially, a trivially determinisable Büchi automaton is as good as a deterministic one. In the long version of this paper, this theorem and all the other theorems that follow are proven in the appendix.

Since pre_P and pre_A are monotonic functions and the state-spaces involved in the algorithm are finite, it follows that the algorithm terminates. The algorithm's soundness is asserted by the following theorem:

Theorem 2. *Once the algorithm has terminated, for all $(q, t, \phi) \in w_\infty$, $\delta_{\mathcal{F}\infty,\infty,\infty}$ is a partial strategy such that $\forall \lambda \in out(\delta_{\mathcal{F}\infty,\infty,\infty}, (q, t, \phi))$ $\pi_1(\lambda[1, \infty]) \in \mathcal{L}(B, \phi)$.*

In much the same way as the completeness condition in Theorem 1, we can give a condition for the algorithm in Figure 2. To do this, we introduce the concept of the *generalised Rabin expansion* of a Büchi automaton. Intuitively, this automaton encodes the idea of finite shifting by its structure and winning condition.

Definition 2. *Let B be a Büchi automaton, and S be the corresponding shift automaton for B. The* generalised Rabin expansion, $R = \langle Q_R, i_R, \delta_R, F_R, E_R \rangle$,

of B and S is defined as

$$Q_R = Q_B \times Q_S$$
$$i_R = i_B \times i_S$$
$$\delta_R = \{((t,\phi), a, (t', \phi')) \in Q_R \times \Sigma \times Q_R \mid \phi' \in \delta_S(\phi, a), t' \in \phi'\}$$
$$F_R = \{(t,\phi) \in Q_R \mid t \in F_B\}$$
$$E_R = \{((t,\phi), a, (t', \phi')) \in \delta_R \mid t' \notin \delta_B(t,a)\}$$

where, in the usual way, Q_R, i_R, and δ_R are the state-space, initial state, and transition function, respectively. F_R and E_R are used to define the winning condition of R: A run ρ on a word λ is winning if and only if $\exists^\infty i \geq 0.\rho[i] \in F_R$ and $\exists j \geq 0.\forall k \geq j \ (\rho[k], \lambda[k], \rho[k+1]) \notin E_R$.

It is convenient to specify a winning condition on transitions rather than states, but it is easy to translate such an automaton into a conventional Rabin automaton. The translation could be done as follows: create a second copy of every state; make every transition in E_R go instead to the copy; make every transition in the copy go back into the original; finally, set the Rabin condition to have infinitely many visits to F_R in the original and only finitely many visits to the copied states. We also note that for any reachable state (t, ϕ) in R, the invariant is maintained that $t \in \phi$.

With this definition in place, it is possible to prove the following theorem about completeness for the synthesis algorithm.

Theorem 3. *For any game G, with a winning condition specified by a Büchi automaton B, if B's generalised Rabin expansion is trivially determinisable then all winning states for P in Q_C satisfy $\mu w.\nu z.\mu s.\pi_1((pre_P(s \vee (z \wedge F_C), w) \wedge (Q_{CP} \times Q_{CA})) \vee (pre_A(s \vee (z \wedge F_C), w) \wedge (Q_{CA} \times Q_{CP}))$.*

We note that this is a safe approximation of the class of problems for which the algorithm will be complete. In fact, the structure of the game is also crucial to completeness. Providing a characterisation which uses the structure of both the game and the specification would be an interesting avenue for future research.

4 Synthesis for LTL Games

In the previous section we provided an algorithm for solving games with NB winning conditions that was complete under a condition on the form of the NB. We can perform synthesis for LTL games by using the tableau method to translate an LTL specification into a NB and then using the algorithm in Figure 2. With the restriction of Theorem 3 and our goal of symbolic implementation in mind, our choice of translation from LTL to NB must be made wisely. The method that we use is based on the symbolic construction of [4], with three changes: First, we deal with formulae in negation normal form rather than using a minimal set of temporal operators – as noted by [15], this provides us with a slight efficiency improvement as safety formulae do not have to be treated as negated liveness

formulae. Secondly, we define a weaker transition formula than [4] – this allows
some types of formulae to be translated into trivially determinisable Büchi au-
tomata and allows for the last change. Finally, we split variables into "required"
and "optional" forms meaning that the fairness constraints of optional formulae
do not necessarily have to be met as long as some other fairness constraints are
whilst the constraints of required formulae must always be met.

First we recall the syntax and semantics of LTL, a more thorough review can
be found in [6]. Syntactically, we consider an LTL formula ψ in negation normal
form to obey the following grammar

$$\Pi ::= p \mid \neg \Pi \mid \Pi \vee \Pi \qquad \psi ::= \Pi \mid \psi \wedge \psi \mid \psi \vee \psi \mid X\psi \mid \psi \, \mathcal{U} \, \psi \mid \psi \, \mathcal{R} \, \psi$$

where p is a member of the set of atomic propositions. The semantics of LTL
are defined inductively over infinite paths λ.

- $\lambda \vDash p$ iff $p \in \lambda[0]$.
- $\lambda \vDash \neg\psi$ iff $\lambda \nvDash \psi$.
- $\lambda \vDash \psi_1 \vee \psi_2$ iff $\lambda \vDash \psi_1$ or $\lambda \vDash \psi_2$.
- $\lambda \vDash \psi_1 \, \mathcal{U} \, \psi_2$ iff $\exists i \geq 0 . \lambda[i, \infty] \vDash \psi_2$ and $\forall j \in [0, i-1] \; \lambda[j, \infty] \vDash \psi_1$.
- $\lambda \vDash \psi_1 \, \mathcal{R} \, \psi_2$ iff $\forall i \geq 0$ either $\lambda[i, \infty] \vDash \psi_2$ or there exists $j \in [0, i]$ such that
  $\lambda[j, \infty] \vDash \psi_1 \wedge \psi_2$.

Like [4], we define a function el() to return the set of elementary sub-formulae
of an LTL formula. The set el(ψ) forms the set of propositions in the tableau for
ψ.

- el(p) = $\{p\}$
- el($\neg\psi$) = el(ψ)
- el($\psi_1 \wedge \psi_2$) = el(ψ_1) $\cup$ el(ψ_2)
- el($X\psi_1$) = $\{(X\psi_1)^r\} \cup$ el(ψ_1)
- el($\psi_1 \vee \psi_2$) = el(ψ_1) $\cup$ el(ψ_2) $\cup \{x^o \mid x^r \in$ el(ψ_1) $\cup$ el(ψ_2)$\}$
- el($\psi_1 \, \mathcal{U} \, \psi_2$) = $\{(X(\psi_1 \, \mathcal{U} \, \psi_2))^r\} \cup$ el(ψ_1) $\cup$ el(ψ_2)
- el($\psi_1 \, \mathcal{R} \, \psi_2$) = $\{(X(\psi_1 \, \mathcal{R} \, \psi_2))^r\} \cup$ el(ψ_1) $\cup$ el(ψ_2)

We see from this that the propositions arising from formulae under an $\vee$ appear
in optional and required forms. This is how the optional formation is used –
in a formula $\psi_1 \vee \psi_2$ the conventional tableau construction would generate an
automaton which chooses between three covering formulae: $\psi_1 \wedge \neg\psi_2$, $\neg\psi_1 \wedge \psi_2$,
and $\psi_1 \wedge \psi_2$. In a game where the opponent can infinitely often choose between
satisfying the fairness constraints of ψ_1 or ψ_2, this splitting can necessitate in-
finite shifting for the algorithm in Figure 2. By making the fairness constraints
optional, we allow the tableau to follow the $\psi_1 \wedge \psi_2$ path as long as the play
is consistent with the safety requirements of $\psi_1 \wedge \psi_2$ and consider the play to
be accepted if it satisfies either the fairness constraints of ψ_1 or the fairness
constraints of ψ_2.

Again like [4], we define a function sat() which takes an LTL formula and
returns a formula representing the set of states in the tableau for which out-
going fair paths are labelled by plays which satisfy the LTL formula. It works

uniformly for all $t \in \{r, o\}$. The only change from the standard definition is in $\mathrm{sat}((\psi_1 \vee \psi_2)^r)$; here we always allow the possibility of the optional versions being taken instead of the required ones. Since the clause for the optional variables is $(\mathrm{sat}(\psi_1^o) \wedge \mathrm{sat}(\psi_2^o))$, the structure of the tableau ensures that paths are consistent with both formulae (i.e. they satisfy the non-fairness part of the formulae). Our new definition of the fairness constraints will allow one or the other to be satisfied. For all $t \in \{r, o\}$, $\mathrm{sat}(\psi^t)$ is defined as:

- $\mathrm{sat}(\Pi^t) = \Pi$
- $\mathrm{sat}((\psi_1 \wedge \psi_2)^t) = \mathrm{sat}(\psi_1^t) \wedge \mathrm{sat}(\psi_2^t)$
- $\mathrm{sat}((\psi_1 \vee \psi_2)^t) = \mathrm{sat}(\psi_1^t) \vee \mathrm{sat}(\psi_2^t) \vee (\mathrm{sat}(\psi_1^o) \wedge \mathrm{sat}(\psi_2^o))$
- $\mathrm{sat}((X\psi_1)^t) = (X\psi_1)^t$
- $\mathrm{sat}((\psi_1 \; \mathcal{U} \; \psi_2)^t) = \mathrm{sat}(\psi_2^t) \vee (\mathrm{sat}(\psi_1^t) \wedge X(\psi_1 \; \mathcal{U} \; \psi_2)^t)$
- $\mathrm{sat}((\psi_1 \; \mathcal{R} \; \psi_2)^t) = (\mathrm{sat}(\psi_1^t) \wedge \mathrm{sat}(\psi_2^t)) \vee (\mathrm{sat}(\psi_2^t) \wedge X(\psi_1 \; \mathcal{R} \; \psi_2)^t)$

The transition relation ensures that if $X\psi$ occurs in a state, all fair paths from all successors of that state satisfy ψ.

$$\bigwedge_{t \in \{r, o\}} \bigwedge_{(X\psi_1)^t \in \mathrm{el}(\psi^r)} (X\psi_1)^t \Rightarrow \mathrm{sat}(\psi_1^t)' \tag{8}$$

This differs from [4] by the inclusions of optional/required tags and by using $\Rightarrow$ instead of $\Leftrightarrow$. This relaxation is possible because the input formulae are in negation normal form and its soundness is implied by Theorem 4.

The fairness constraints on the tableau are defined by another new function, fsat(). Conventionally, the fairness constraints for a formula ψ would require that for each sub-formula of the form $\psi_1 \; \mathcal{U} \; \psi_2$ a fair path infinitely often has either $\neg X(\psi_1 \; \mathcal{U} \; \psi_2)$ or $\mathrm{sat}(\psi_2)$ i.e. at any point, either $\psi_1 \; \mathcal{U} \; \psi_2$ is not required or it is eventually satisfied. Our definition of fsat is based on this notion, but allows for the special case of optional variables. Fairness on a path π is defined by the following function:

- $\mathrm{fsat}(\Pi^t) = \top$
- $\mathrm{fsat}((\psi_1 \vee \psi_2)^t) = \mathrm{fsat}(\psi_1^t) \wedge \mathrm{fsat}(\psi_2^t) \wedge (\mathrm{fsat}(\psi_1^o) \vee \mathrm{fsat}(\psi_2^o))$
- $\mathrm{fsat}((\psi_1 \wedge \psi_2)^t) = \mathrm{fsat}(\psi_1^t) \wedge \mathrm{fsat}(\psi_2^t)$
- $\mathrm{fsat}((X\psi_1)^t) = \mathrm{fsat}(\psi_1^t)$
- $\mathrm{fsat}((\psi_1 \; \mathcal{U} \; \psi_2)^t) = \mathrm{fsat}(\psi_1^t) \wedge \mathrm{fsat}(\psi_2^t)$
  $$\wedge \; \exists^{\infty} i \geq 0 . \pi[i] \in \mathrm{sat}(\psi_2^t) \vee \neg(X(\psi_1 \; \mathcal{U} \; \psi_2))^t$$
- $\mathrm{fsat}((\psi_1 \; \mathcal{R} \; \psi_2)^t) = \mathrm{fsat}(\psi_1^t) \wedge \mathrm{fsat}(\psi_2^t)$

We see that the only departure from the conventional usage (a convention observed by [4]) is in allowing one or the other of a pair of optional variables to be satisfied.

Since our construction is a relaxation of the one given by [4], we do not prove its completeness. However, its soundness is asserted by the following theorem:

Theorem 4. *Let ψ be an LTL formula over a set, $\mathcal{P}$, of atomic propositions. Let T be the symbolic tableau automaton for ψ constructed as above. For any $\psi_1 \in \mathrm{sub}(\psi)$, any $t \in \{r, o\}$ and any path, π, in T, for all $i \geq 0$ if $\pi[i] \in \mathrm{sat}(\psi_1^t)$ and $\pi[i, \infty]$ satisfies $\mathrm{fsat}(\psi_1^t)$ then $\pi[i, \infty] \vDash \psi_1$.*

Using this tableau construction we can perform the entire synthesis procedure symbolically. First the LTL formula is translated in the manner described above. This tableau is then used as the Büchi automaton in the algorithm from Figure 2 which can also be computed symbolically. Despite the doubly exponential worst-case complexity of this procedure we shall see in the following section that useful results can be computed. First, we note that for request-response specifications [20], strategies can be synthesised without generating a shift automaton i.e. the tableau for such specifications is trivially determinisable. Request-response specifications can be written in LTL as $G(p \wedge (r_0 \Rightarrow Fs_0) \wedge \ldots \wedge (r_n \Rightarrow Fs_n))$ and we prove the following proposition in the appendix.

Proposition 1. *The tableau for LTL formulae of the form $G(p \wedge (r_0 \Rightarrow Fs_0) \wedge \ldots \wedge (r_n \Rightarrow Fs_n))$ (where p, r_i, and s_i are propositional formulae) is always trivially determinisable.*

5 Implementation

The synthesis algorithm has been implemented in Java using native calls to the CuDD [5] library to handle BDDs. As input, the program takes an XML file containing a symbolic description of the game and an LTL specification. After successful synthesis, a number of output options are possible: the program can print the set of winning states; produce an explicit graph of the winning strategy with dot; show an interactive, expandable tree of the strategy so that the user can play it out; and convert the strategy into a program in the language of the Cadence SMV model checker [3]. The output to SMV can be used to check the correctness of the implementation by checking against the original LTL specification (once a winning strategy has been synthesised, it is possible to view the strategy as a closed system and check for correctness on all paths).

5.1 Examples

Mutual Exclusion In this example we synthesise a controller to enforce mutual exclusion. We solve this problem for various numbers of processes in order to get a measure of the scalibility of the implementation, using n as the parameter for size.

The game is modelled with the following boolean variables: u indicates the current turn and alternates between moves. When u is true, it is the user processes turn, when it is false it is the system's turn. $r1, r2, \ldots, rn$ indicate that a process requesting access to its critical section. When a request is made, it cannot be withdrawn until the critical section is reached. Furthermore, when it is the system's turn, the request variables must keep their old values. $c1, c2, \ldots, cn$ indicate that a process is in its critical section. ci can only become true if it is the system's turn and ri is true. When it is true, ci will become false in the next turn.

The specification of mutual exclusion and liveness can easily be written in in LTL as

$$G(\bigwedge_{i \in [1,n]} ((ri \Rightarrow Fci) \wedge \bigwedge_{j \in [i+1,n]} \neg(ci \wedge cj)))$$

For two processes, the synthesis implementation ran in about 30ms and produced a strategy which plays out Peterson's algorithm for mutual exclusion. Whenever both processes request at once, one process (say process 1) is chosen non-deterministically to proceed. From then until process 2 gets to enter its critical section, if both processes request at once then process 2 will be favoured. This behaviour is symmetric for both processes, guaranteeing that neither process starves. Strategies were found for problems with between 2 and 80 processes, recording the "peak live node usage" and time for each problem. The peak live node usage is a statistic gathered by the CuDD library denoting the maximum number of BDD nodes that have been used during a computation. The tests were run on a 1.5GHz Pentium 4 system with 256MB of RAM running Linux. The results are shown in Figure 3. We note that even with 80 processes, which gives a state-space of $2^{(80 \times 2)+1} \approx 10^{48}$ and a formula with 80 liveness sub-formulae, the time taken was about 50 minutes. This is quite reasonable for a model checking tool and the overall growth in the two plots shows the feasibility of the approach.
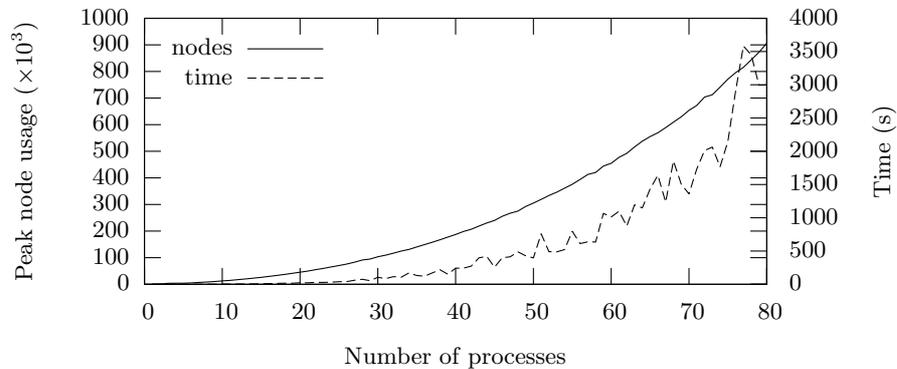
Fig. 3. Mutual exclusion performance data

Lift system Here we synthesise a controller for a lift (elevator) system. The game describes some general behaviour regarding the physical situation of a lift (there are user-controlled buttons, movement between floors is consecutive etc.), and the LTL specification puts requirements on the actual controller. We model the lift system with various numbers of floors, using n as a parameter for the size; the variables used to describe the system are: u which indicates the current turn and works as in mutual exclusion. f which indicates the current floor. This is modelled with $\lceil log(n) \rceil$ variables which we treat as a single integer variable. We write $f[0], f[1], \ldots f[n]$ to denote floors. Initially, the floor is 0; we require that transitions between floors are consecutive and that the floor does not change on users turns. $b[0], b[1], \ldots, b[n]$ are the button variables. These boolean variables are controlled by the users to simulate requests for the lift. Initially, all buttons are off; once lit, a button stays on until the lift arrives and the buttons do not change on the system's turns. up is a boolean variable which observes the

transition between floors and is true if the lift is going up. Initially *up* is false. The specification that we use for the lift system is as follows:

$$G\Big(\big(\bigwedge_{i\in[0,n]} b[i] \Rightarrow Ff[i]\big) \wedge \bigwedge_{i\in[0,n]} f[i] \wedge \neg sb \Rightarrow f[i] \; \mathcal{U} \; (sb \; \mathcal{R} \; (F(f[0] \vee sb) \wedge \neg up))\big)\Big)$$

Where sb is an abbreviation for $\bigvee_{i\in[1,n]} b[i]$ to mean "some button is lit". The first conjunct says that every request is eventually answered, the second demands that the lift should park when it is idle. If we synthesise a strategy without the parking specification, we find that the lift does answer all calls, but it does so by moving up and down continuously regardless of what calls are made (this was apparent from playing out the strategy and verified formally using SMV). In the second conjunct, the release formula is the actual parking action: the lift should go to $f[0]$ by going continuously down, unless some button is pressed. The rest of the formula can be read as: if the lift is on floor i and no button (other than 0) is pressed, then remain at floor i until parking can commence. Synthesising a strategy for the entire specification, we find that the lift now behaves as expected and, once again, we can verify that the strategy implements the specification by using SMV. The results for a range of sizes are plotted in Figure 4 – we see the time taken and the number of nodes used rising dramatically. This is due to the size and complexity of the specification resulting in an extremely large tableau. The fluctuations in time taken and nodes used are hard to explain. One possible cause could be the heuristic nature of the re-ordering algorithms in CuDD. Similar fluctuations are seen in the shifting example below, but they seem to be magnified here by the size of the state-space increase as each extra floor is added to the lift problem.
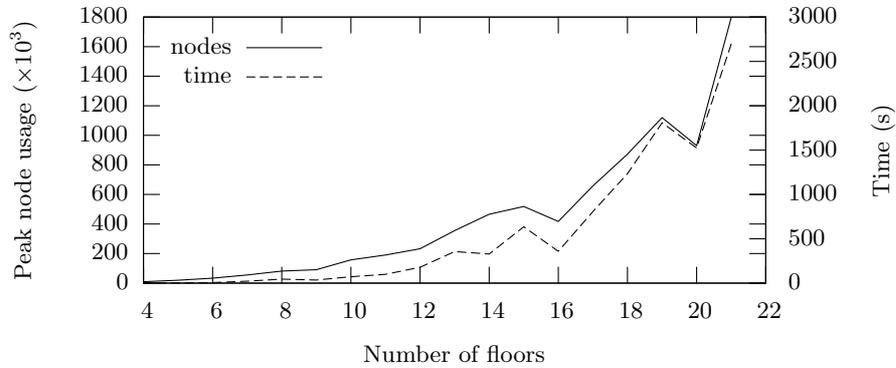
Fig. 4. Lift performance data

A Shifting Problem The previous two examples were computed without having to use any shifting. Although a number of specifications were tried in the context of these systems, shifting was never required. It is clear that shifting is necessary for some specifications, though, so we use an abstract example based

on the one from Figure 1 to measure performance in such cases. The game is shown in Figure 5.
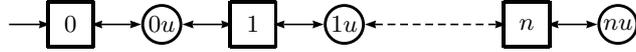
Fig. 5. Game for shifting example

The specification used is simply $FG0 \vee FG1 \vee \ldots \vee FGn$. For $n = 2$, this produces a tableau with the same property as the Büchi automaton in Figure 1 – it has two disjoint regions, one for $FG0$ and one for $FG1$. In order to solve this game, the algorithm needs to consider shifting. As expected, the synthesis algorithm generates a shift automaton and terminates using one shift. The winning strategy simply keeps play in 1 if the opponent ever chooses it, otherwise it has no choice but to stay in 0. Figure 6 shows the performance over a range of sizes. Even with a length of 95, and having to generate the doubly exponential shift automaton, the implementation ran in 1 hour 14 minutes.
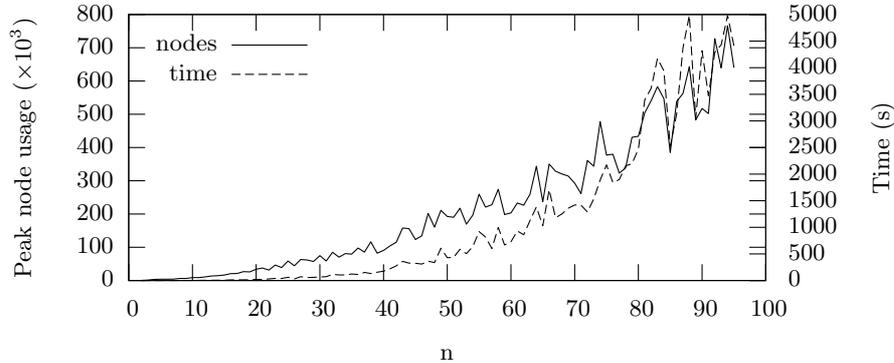
Fig. 6. Shift performance data

6 Conclusion

We have provided a new algorithm and corresponding symbolic implementation to solve the problem of strategy synthesis in LTL games. Whilst this algorithm is not complete, it has performed well in the test-cases that were given to it and holds enough promise to warrant use on larger, real-world problems in the future. The separation between the partial solution to NB games and the symbolic tableau method itself offers the prospect of future improvements to the completeness result by changes to the tableau method. It would be interesting to see whether the completeness result given in terms of automata could be related back to a fragment of LTL so that it might be compared with such work as [2]. At present, there are formulae for which we know that the algorithm here does not perform optimally and there are formulae which cannot be dealt with in fragments solved by [2], but can be dealt with by our work.

References

1. R. Alur, T. A. Henzinger, and O. Kupferman. Alternating-time temporal logic. *Journal of the ACM*, 49(5):672–713, 2002.
2. R. Alur and S. La Torre. Deterministic generators and games for LTL fragments. *ACM Transactions on Computational Logic*, 5(1):1–25, January 2004.
3. SMV 10-11-02p1. `http://www-cad.eecs.berkeley.edu/~kenmcmil/smv/`, November 2002.
4. E. Clarke, O. Grumberg, and K. Hamaguchi. Another look at LTL model checking. In David L. Dill, editor, *CAV94*, volume 818, pages 415–427. Springer-Verlag, 1994.
5. CuDD: Colorado university decision diagram package, release 2.30. `http://vlsi.colorado.edu/~fabio/CUDD/`, February 2001.
6. E. A. Emerson. *Handbook of Theoretical Computer Science Volume B*, chapter Temporal and Modal Logic, pages 995–1072. Elsevier, 1990.
7. E. A. Emerson and C. Lei. Efficient model checking in fragments of the propositional model mu-calculus. In *IEEE Symposium on Logic in Computer Science*, pages 267–278, June 1986.
8. E. Grädel, W. Thomas, and T. Wilke, editors. *Automata, Logic, and Infinite Games*. Number 2500 in Lecture Notes In Computer Science. Springer, 2002.
9. S. Kremer and J.-F. Raskin. A game-based verification of non-repudiation and fair exchange protocols. *Journal Of Computer Security*, 11(3):399–429, 2003.
10. O. Kupferman and M. Y. Vardi. Module checking. In *8th Conference on Computer-Aided Verification*, volume 1102 of *LNCS*, pages 75–86, 1996.
11. C. Löding. Optimal bounds for the transformation of ω-automata. In *Proceedings of the 19th Conference on Foundations of Software Technology and Theoretical Computer Science*, number 1738 in LNCS, pages 97–109. Springer, 1999.
12. A. Pnueli and R. Rosner. On the synthesis of a reactive module. In *Proceedings Of 16th ACM Symposium On Principles Of Programming Languages*, pages 179–190, 1989.
13. R. Rosner. *Modular Synthesis of Reactive Systems*. PhD thesis, Weizmann Institute of Science, Rehovot, Israel, 1992.
14. S. Safra. *Complexity of Automata on Infinite Objects*. PhD thesis, The Weizmann Institute of Science, Rehovot, Israel, March 1989.
15. K. Schneider. Improving automata generation for linear temporal logic by considering the automaton hierarchy. *Lecture Notes in Computer Science*, 2250, 2001.
16. F. Somenzi, K. Ravi, and R. Bloem. Analysis of symbolic scc hull algorithms. In *Proceedings of the 4th International Conference on Formal Methods in Computer-Aided Design*, pages 88–105. Springer-Verlag, 2002.
17. S. Tasiran, R. Hojati, and R. K. Brayton. Language containment using non-deterministic omega-automata. In *Proceedings of Advanced Research Working Conference on Correct Hardware Design and Verification Methods, CHARME'95*, number 987 in Lecture Notes In Computer Science, pages 261–277, 1995.
18. W. Thomas. *Handbook of Theoretical Computer Science Volume B*, chapter Automata on Infinite Objects, pages 133–192. Elsevier, 1990.
19. W. Thomas. On the synthesis of strategies in infinite games. In *Symposium on Theoretical Aspects of Computer Science*, pages 1–13, 1995.
20. N. Wallmeier, P. Hütten, and W. Thomas. Symbolic synthesis of finite-state controllers for request-response specifications. In *Proceedings of the 8th International Conference on the Implementation and Application of Automata, CIAA*, volume 2759 of *Lecture Notes in Computer Science*, pages 11–22, 2003.

A Appendix

A.1 Proof of Theorem 1

Theorem 1. *For any game G, with a winning condition specified by a Büchi automaton B, if B is trivially determinisable, then all winning states for P in Q_C satisfy $\nu z.\mu s.\pi_1((pre_P(s \vee (z \wedge F_C), \bot) \wedge (Q_{CP} \times Q_{CA})) \vee (pre_A(s \vee (z \wedge F_C), \bot) \wedge (Q_{CA} \times Q_{CP}))$.*

Proof. Let $B = \langle Q_B, i_B, \delta_B, F_B \rangle$ be a trivially determinisable Büchi automaton and $G = \langle Q_G, i_G, Q_P, Q_A, \delta_G \rangle$ be a game. We prove that for any state $(q, t, \phi) \in Q_C$ (where Q_C is defined by Equation 2) such that P has a strategy to ensure that all plays in $out(G, q)$ are in $\mathcal{L}(B, t)$, $(q, t, \phi) \vDash \nu z.\mu s.\pi_1((pre_P(s \vee (z \wedge F_C), \bot) \wedge (Q_{CP} \times Q_{CA})) \vee (pre_A(s \vee (z \wedge F_C), \bot) \wedge (Q_{CA} \times Q_{CP}))$. Since B is trivially determinisable and the pre_P/pre_A operators evaluate Büchi transitions existentially, we can treat B as deterministic by assuming that every existential choice is the one that would be used in trivially determinising B. We prove using fixed-point induction and contradiction: Suppose $(q, t, \phi) \vDash \nu^n z.\mu s.\pi_1((pre_P(s \vee (z \wedge F_C), \bot) \wedge (Q_{CP} \times Q_{CA})) \vee (pre_A(s \vee (z \wedge F_C), \bot) \wedge (Q_{CA} \times Q_{CP}))$ but $(q, t, \phi) \nVDash \nu^{n+1} z.\mu s.\pi_1((pre_P(s \vee (z \wedge F_C), \bot) \wedge (Q_{CP} \times Q_{CA})) \vee (pre_A(s \vee (z \wedge F_C), \bot) \wedge (Q_{CA} \times Q_{CP}))$. Then, for any strategy f_P, there would be a run $\lambda \in out(q, f_P)$ such that the corresponding run $\rho_B \in run(B, t, \lambda)$ visits F_B n times; but on the n-th visit (call this $(\lambda[i], \rho_B[i+1])$), $(\lambda[i], \rho_B[i+1]) \nVDash \nu^0 z.\mu s.\pi_1((pre_P(s \vee (z \wedge F_C), \bot) \wedge (Q_{CP} \times Q_{CA})) \vee (pre_A(s \vee (z \wedge F_C), \bot) \wedge (Q_{CA} \times Q_{CP}))$. This in turn, means that for any strategy f'_P for P there is a run $\lambda' \in out(\lambda[i], f'_P)$ such that the corresponding Büchi run never visits F_B. So whatever the strategy f_P does from $\lambda[i]$, it is guaranteed never to reach F_B. This contradicts the premise that there is a winning strategy for P. □

A.2 Proof of Theorem 2

We begin with a loop invariant on the inner loop to say (roughly) that if $(q, t, \phi) \in s_{j,k,l} - w_{j-1}$ then $\delta_{\mathcal{F}\infty,\infty,\infty}$ is a partial strategy to reach $\tau_{j,k}$ in l steps or less from (q, t, ϕ) (possibly using a shift).

Lemma 1. *For all $l \geq 0$, if $\delta_{\mathcal{F}j-1,\infty,\infty}$ is a partial strategy then $\delta_{\mathcal{F}j,k,l}$ is a partial strategy and if $(q, t, \phi) \in s_{j,k,l} - w_{j-1}$ then $\forall \lambda \in out(\delta_{\mathcal{F}j,k,l}, (q, t, \phi))$ there exists an integer $i \in [1, l]$ such that $\forall m \in [1, i-1]\ \lambda[i] \in s_{j,k,l-m}$ and either $(\lambda[i] \in w_{j-1})$ or $(\lambda[i] \in z_{j,k-1} \cap F_C$ and $\forall m \in [0, i-1]\ \pi_2(\lambda[m+1]) \in \delta_B(\pi_2(\lambda[m]), \pi_1(\lambda[m+1])))$.*

Proof. We prove by induction over l.

- $l = 0$: By line 7 $s_{j,k,0}$ is empty. By line 8 $\delta_{\mathcal{F}j,k,0} = \delta_{\mathcal{F}j-1,\infty,\infty}$, so by hypothesis $\delta_{\mathcal{F}j,k,0}$ is a non-deterministic partial strategy.

– $l > 0$: If $(q, t, \phi) \in s_{j,k,l-1}$ then we simply use the inductive hypothesis. From line 12 of the algorithm, if $\delta_{\mathcal{F}j,k,l-1}(q, t, \phi)$ is already defined then $\delta_{\mathcal{F}j,k,l-1}(q, t, \phi) = \delta_{Fj,k,l}(q, t, \phi)$. Since each move in the partial strategy is in a smaller $s_{j,k,l}$, there is a continuous path of preserved strategies up to each i.

1. Otherwise, $(q, t, \phi) \in s_{j,k,l} - s_{j,k,l-1}$.
2. First, suppose that $q \in Q_{GP}$, then by lines 10, 11, and 13 of the algorithm $(q, t, \phi) \in \pi_1(u_P)$ and $(q, t, \phi) \notin \pi_1(u_A)$. Let $Succ = \{(q', t', \phi') \mid \langle(q, t, \phi), (q', t', \phi')\rangle \in u_P\}$.
3. By line 10 of the algorithm and the definition of pre_A (Equation 6) for all $(q', t', \phi') \in Succ$ $q' \in \delta_G(q)$, $\phi' = \delta_S(\phi, q')$ and either $(t' \in \phi' \wedge (q', t', \phi') \in w_{j-1})$ or $(t' \in \delta_B(t, q') \wedge (q', t', \phi') \in \tau_{j,k} \cup s_{j,k,l-1})$. If $\delta_{\mathcal{F}j,k,l-1}(q, t, \phi)$ is undefined then $\delta_{\mathcal{F}j,k,l}(q, t, \phi) = Succ$; we have dealt with the case where $(q, t, \phi) \in s_{j,k,l-1}$ and by hypothesis $(q, t, \phi) \notin w_{j-1}$ so $\delta_{\mathcal{F}j,k,l-1}(q, t, \phi)$ is undefined. Thus, $\delta_{\mathcal{F}j,k,l}(q, t, \phi) = Succ$ which are all legal moves, so if $\delta_{\mathcal{F}j,k,l-1}$ is a partial strategy (by induction, it is) then $\delta_{\mathcal{F}j,k,l}$ is a partial strategy.
4. We split cases on $(q', t', \phi') \in Succ$:
   - If $(q', t', \phi') \in w_{j-1}$, then we are done for $q \in Q_{GP}$.
   - If $(q', t', \phi') \in \tau_{j,k}$, then $(q', t', \phi') \in F_C \cap z_{j,k-1}$ so we are done for $q \in Q_{GP}$.
   - If $(q', t', \phi') \in s_{j,k,l-1}$ we apply the inductive hypothesis
5. Now suppose that $q \in Q_{GA}$. By lines 10, 11, and 13 of the algorithm $(q, t, \phi) \in \pi_1(u_A)$ and $(q, t, \phi) \notin \pi_1(u_P)$. Let $Succ = \{(q', t', \phi') \mid \langle(q, t, \phi), (q', t', \phi')\rangle \in u_A\}$.
6. We argue, as above, that $\delta_{\mathcal{F}j,k,l-1}(q, t, \phi)$ is undefined, so $\delta_{\mathcal{F}j,k,l}(q, t, \phi) = Succ$.
7. We must show that $\mathcal{F}_{j,k,l}$ is a non-deterministic partial strategy automaton (i.e. all moves are legal and it includes moves for all $q' \in \delta_G(q)$) and that all outcomes are good. By line 11 of the algorithm and the definition of pre_V (7) $\forall q' \in \delta_G(q) \exists (q', t', \phi') \in \pi_2(u_A)$ so if $\delta_{\mathcal{F}j,k,l-1}$ is a partial strategy (by induction, it is) then $\delta_{\mathcal{F}j,k,l}$ is a partial strategy. Furthermore, $\langle(q, t, \phi)(q', t', \phi')\rangle \in u_A$ if and only if either $(t' \in \phi' \wedge (q', t', \phi') \in w_{j-1})$ or $(t' \in \delta_B(t, q') \wedge (q', t', \phi') \in \tau_{j,k} \cup s_{j,k,l-1})$. We can show that all outcomes are good by splitting cases on (q', t', ϕ'), in a similar way as above.

$\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\square$

Now we prove the result of the middle loop: if $(q, t, \phi) \in z_{j,\infty}$, then $\delta_{\mathcal{F}j,\infty,\infty}$ is a partial strategy to reach F_C infinitely often, or reach w_{j-1}.

Lemma 2. *If $\delta_{\mathcal{F}j-1,\infty,\infty}$ is a partial strategy then $\delta_{\mathcal{F}j,\infty,\infty}$ is a partial strategy and if $(q, t, \phi) \in z_{j,\infty}$ then $\forall \lambda \in out(\delta_{\mathcal{F}j,\infty,\infty}, (q, t, \phi))$ either $(\exists i > 0$ such that $\lambda[i] \in w_{j-1})$ or $(\exists^\infty i > 0.\lambda[i] \in F_C$ and $\forall m \geq 0$ $\pi_2(\lambda[m + 1]) \in \delta_B(\pi_2(\lambda[m]), \pi_1(\lambda[m + 1])))$.*

Proof. We know that the middle loop will terminate at a fixed-point for $z_{j,k}$ i.e. at termination $z_{j,k} = z_{j,k-1}$. Let n be the value of k when the middle loop terminates. The loop executes at least once, so $n > 0$. $(q, t, \phi) \in z_{j,n}$, so by line 15 of the algorithm $(q, t, \phi) \in s_{j,k,\infty}$. Thus by Lemma 1 if $\delta_{\mathcal{F}j-1,\infty,\infty}$ is a partial strategy then $\delta_{\mathcal{F}j,n,\infty}$ is a partial strategy and $\forall \lambda \in out(\delta_{\mathcal{F}j,n,\infty}, (q, t, \phi))$ there exists an integer $i > 0$ such that either $(\lambda[i] \in w_{j-1})$ or $(\lambda[i] \in z_{j,n-1} \cap F_C$ and $\forall m \in [0, i-1]\ \pi_2(\lambda[m+1]) \in \delta_B(\pi_2(\lambda[m]), \pi_1(\lambda[m+1])))$. If $\lambda[i] \in w_{j-1}$ then we are done. If $\lambda[i] \in z_{j,n-1} \cap F_C$, we have shown one visit to F_C and since $z_{j,n} = z_{j,n-1}$ we can re-use the same argument to show that $\forall \lambda' \in out(F_{j,\infty,\infty}, C, \lambda[i])$ there exists an integer $i_2 > 0$ such that either $(\lambda'[i_2] \in w_{j-1})$ or $(\lambda'[i_2] \in z_{j,n-1} \cap F_C$ and $\forall m \in [0, i_2 - 1]\ \pi_2(\lambda[m+1]) \in \delta_B(\pi_2(\lambda[m]), \pi_1(\lambda[m+1])))$.

Since λ and λ' meet at $\lambda[i]$ and $\lambda'[0]$, we have shown continuous run which either reaches w_{j-1} or visits F_C twice and obeys δ_B until the second visit. We can re-use this argument infinitely many times to show that there is a run visiting F_C infinitely often and obeying δ_B forever, or there is a point where λ reaches w_{j-1}. $\qed$

For the final lemma we prove an invariant on the outer loop: $(q, t, \phi) \in w_j$ implies that $\delta_{\mathcal{F}j,\infty,\infty}$ is a partial strategy such that the game part of all outcomes can be accepted by B from ϕ.

Lemma 3. *For all $j \geq 0$ $\delta_{\mathcal{F}j,\infty,\infty}$ is a non-deterministic partial strategy automaton and if $(q, t, \phi) \in w_j$ then $\forall \lambda \in out(\delta_{\mathcal{F}j,\infty,\infty}, (q, t, \phi))\ \pi_1(\lambda[1, \infty]) \in \mathcal{L}(B, \phi)$.*

Proof. We prove by induction over j:

- $j = 0$: By line 1 of the algorithm w_0 is empty. By line 2 of the algorithm, $\delta_{\mathcal{F},0,\infty,\infty} = \emptyset$ so $\delta_{\mathcal{F}0,\infty,\infty}$ is a partial strategy.
- $j > 0$: First we show that $\delta_{\mathcal{F}j,\infty,\infty}$ is a partial strategy. By inductive hypothesis, $\delta_{\mathcal{F}j-1,\infty,\infty}$ is a partial strategy; thus it follows from Lemma 2 that $\delta_{\mathcal{F}j,\infty,\infty}$ is a partial strategy.

  If $(q, t, \phi) \in w_{j-1}$, then we use the inductive hypothesis. Moves in $\delta_{\mathcal{F}j-1,\infty,\infty}$ are carried forward to $\delta_{\mathcal{F}j,\infty,\infty}$ and never overwritten, so this is sound. Otherwise, $(q, t, \phi) \in w_j - w_{j-1}$:

  1. By line 17 of the algorithm, $(q, t, \phi) \in z_{j,\infty}$. By Lemma 2, this implies that $\forall \lambda \in out(\delta_{\mathcal{F}j,\infty,\infty}, (q, t, \phi))$ either $(\exists i > 0$ such that $\lambda[i] \in w_{j-1})$ or $(\exists^\infty i > 0.\lambda[i] \in F_C$ and $\forall m \geq 0\ \pi_2(\lambda[m+1]) \in \delta_B(\pi_2(\lambda[m]), \pi_1(\lambda[m+1])))$.

  2. Suppose $\exists^\infty i > 0.\lambda[i] \in F_C$ and $\forall m \geq 0\ \pi_2(\lambda[m+1]) \in \delta_B(\pi_2(\lambda[m]), \pi_1(\lambda[m+1]))$. This means that there exists $\rho_B \in run(B, t, \pi_1(\lambda[1, \infty]))$ such that $\rho_B = \pi_2(\lambda)$. So, by the Equation 3 $\lambda[i] \in F_C \Leftrightarrow \rho_B[i] \in F_B$. Thus, $\pi_1(\lambda[1, \infty]) \in \mathcal{L}(B, t)$.

  3. Otherwise, $\exists i > 0.\lambda[i] \in w_{j-1}$. By the inductive hypothesis, $\forall \lambda' \in out(\delta_{\mathcal{F}j-1,\infty,\infty}, \lambda[i])\ \pi_1(\lambda'[1, \infty]) \in \mathcal{L}(B, \pi_3(\lambda[i]))$.

  4. Since strategies are not overwritten $\lambda[i, \infty] \in out(\delta_{\mathcal{F}j-1,\infty,\infty}, \lambda[i])$, thus $\pi_1(\lambda[i+1, \infty]) \in \mathcal{L}(B, \pi_3(\lambda[i]))$

5. This means that there exists $t' \in \pi_3(\lambda[i])$ such that $\exists \rho' \in run(B, t', \pi_1(\lambda[i+1, \infty])).\exists^\infty j \geq 0.\rho'[j] \in F_B$.

6. Since the strategy obeys the transition relation of S, by the definition of δ_S in Equation 1 there exists $t_B \in \phi, \rho_B \in run(B, t_B, \pi_1(\lambda[1, i]))$ such that $\rho_B[i] = t'$.

7. Thus, there is a path in B, $\rho = \rho_B[0, i]\rho'[1, \infty]$ which is labelled by $\pi_1(\lambda)$ and for which $\exists^\infty j \geq 0.\rho[j] \in F_B$ i.e. $\pi_1(\lambda[1, \infty]) \in \mathcal{L}(B, \phi)$

$\square$

Theorem 2. *Once the algorithm has terminated, for all $(q, t, \phi) \in w_\infty$, $\delta_{\mathcal{F}\infty, \infty, \infty}$ is a partial strategy such that $\forall \lambda \in out(\delta_{\mathcal{F}\infty, \infty, \infty}, (q, t, \phi))$ $\pi_1(\lambda[1, \infty]) \in \mathcal{L}(B, \phi)$.*

Proof. Follows directly from termination and Lemma 3. $\square$

A.3 Proof of Theorem 3

We begin by proving the following unsurprising lemma:

Lemma 4. *A Büchi automaton B and its generalised Rabin expansion R are language-equivalent.*

Proof. First we prove that if $\lambda \in \mathcal{L}(B)$ then $\lambda \in \mathcal{L}(R)$. By the definition of δ_R, if $t' \in \delta_B(t, a)$ then for all $(t, \phi) \in Q_R$ such that $t \in \phi$, $t' \in \pi_1(\delta_R(t, \phi))$. So, since there exists a run $\rho_B \in win(B, \lambda)$ and for any reachable state (t, ϕ) in R, $t \in \phi$, we can conclude that there exists a run $\rho_R \in run(R, \lambda)$ such that $\pi_1(\rho_R) = \rho_B$. Furthermore, if ρ_B visits F_B infinitely often, then ρ_R visits F_R infinitely often and since every transition in ρ_R is in δ_B none of them are in E_R so ρ_R is a winning run and thus $\lambda \in \mathcal{L}(R)$.

Now we prove that if $\lambda \in \mathcal{L}(R)$ then $\lambda \in \mathcal{L}(B)$. There exists a run $\rho_R \in win(R, \lambda)$ so $\exists i \geq 0.\forall j \geq i \ (\rho_R[j], \lambda[j], \rho_R[j+1]) \notin E_R$. So for all $j \geq i$ $\pi_1(\rho_R[j+1]) \in \delta_B(\pi_1(\rho_R[j]), \lambda[j])$ i.e. $\pi_1(\rho_R[i, \infty]) \in run(B, \pi_1(\rho_R[i]), \lambda[i, \infty])$. We must show that B can reach this point and that the complete run is accepting. Since S maintains the set of reachable states in B and for every reachable state (t, ϕ) in R, $t \in \phi$, there is a run ρ_B of B on $\lambda[0, i-1]$ such that $\rho_B[i] = \pi_1(\rho_R[i])$. Thus there is a complete run ρ_B of B on λ such that $\rho_B[i, \infty] = \pi_1(\rho_R[i, \infty])$ and by the acceptance condition of R this run visits F_B infinitely often so $\lambda \in \mathcal{L}(B)$. $\square$

Now we define new predecessor functions to work over the product of a game and a generalised Rabin expansion.

$$pre_{RP}(X, W) = \{\langle(q, (t, \phi)), (q', (t', \phi'))\rangle \mid q' \in \delta_G(q), (t', \phi') \in \delta_R((t, \phi), q')$$
$$(\langle(t, \phi), q'(t', \phi')\rangle \notin E_R \wedge (q', (t', \phi')) \in X) \vee ((q', (t', \phi')) \in W)\}$$
(9)

$$pre_{RA}(X, W) = \{\langle(q, (t, \phi)), (q', (t', \phi'))\rangle \in pre_P(X, W) \mid \forall q'_2 \in \delta_G(q)$$
$$\exists \langle(q, (t, \phi)), (q'_2, (t'_2, \phi'_2))\rangle \in pre_P(X, W)\}$$
(10)

We also define abbreviations over $G \times R$ as follows:

$$F_{GR} = \{(q, (t, \phi)) \in Q_G \times Q_R \mid (t, \phi) \in F_R\} \tag{11}$$

$$Q_{GRP} = \{(q, (t, \phi)) \in Q_G \times Q_R \mid q \in Q_P\} \tag{12}$$

$$Q_{GRA} = \{(q, (t, \phi)) \in Q_G \times Q_R \mid q \in Q_A\} \tag{13}$$

Lemma 5. *For any game G, with a winning condition specified by a Büchi automaton B, if B's generalised Rabin expansion R is trivially determinisable then all winning states for P in $G \times R$ satisfy*

$$\mu w.\nu z.\mu s.\pi_1((pre_{RP}(s \vee (z \wedge F_{GR}), w) \wedge (Q_{GRP} \times Q_{GRA}))$$
$$\vee (pre_{RA}(s \vee (z \wedge F_{GR}), w) \wedge (Q_{GRA} \times Q_{GRP})) \tag{14}$$

Proof. Since R is trivially determinisable, and its transitions are evaluated existentially in pre_{RP} and pre_{RA}, we treat it as deterministic. We prove that (under the assumed determinism) for all states $(q, (t, \phi)) \in Q_G \times Q_R$ if $(q, (t, \phi))$ is winning for P then it satisfies Equation 14. For any winning state $(q, (t, \phi))$ there exists $n \geq 0$ such that if f is the winning strategy for P, $\forall \lambda \in out(f, q)$ the run $\rho = win(R, (t, \phi), \lambda)$ has at most n transitions from E_R: a pumping argument shows that if there are a series of runs in $out(f, q)$ requiring an increasing number of transitions from E_R, then there is a run which requires infinitely many, violating the premise that f is winning. We prove by induction over n that $(q, (t, \phi)) \vDash \mu^n w.\nu z.\mu s.\pi_1((pre_{RP}(s \vee (z \wedge F_{GR}), w) \wedge (Q_{CP} \times Q_{CA})) \vee (pre_{RA}(s \vee (z \wedge F_R), w) \wedge (Q_{GRA} \times Q_{GRP}))$.

- $n = 0$: When $n = 0$, R behaves like a deterministic Büchi automaton so we can prove that $(q, (t, \phi)) \vDash \mu^0 w.\nu z.\mu s.\pi_1((pre_{RP}(s \vee (z \wedge F_R), w) \wedge (Q_{CP} \times Q_{CA})) \vee (pre_{RA}(s \vee (z \wedge F_R), w) \wedge (Q_{CA} \times Q_{CP}))$ in the same way as we proved Theorem 1.
- $n > 0$: We show by fixed-point induction that $(q, (t, \phi))$ satisfies

$$\nu z.\mu s.\pi_1((pre_{RP}(s \vee (z \wedge F_R), W^{n-1}) \wedge (Q_{CP} \times Q_{CA}))$$
$$\vee (pre_{RA}(s \vee (z \wedge F_R), W^{n-1}) \wedge (Q_{CA} \times Q_{CP})) \tag{15}$$

Where W^{n-1} is the $n-1$-th approximant of Equation 14. If $(q, (t, \phi))$ is in the m-th approximant of Equation 15, and not the $m+1$-th then for all strategies f there exists $\lambda \in out(f, q)$ such that on the corresponding run $\rho = run(R, \lambda)$ visits F_R m times, but on the m-th visit (say, at $(\lambda[i], \rho[i+1]))$, $(\lambda[i], \rho[i+1])$ is not in the 0-th approximant of Equation 15. This means that for every strategy f' there is a run $\lambda' \in run(f', \lambda[i])$ such that the corresponding run $\rho' = run(R, \rho[i+1], \lambda')$ either does not visit F_R infinitely often without visiting E_R or it visits E_R and then is not in W^{n-1}. Take the first case, if this is true, then $\lambda' \notin \mathcal{L}(R, \rho[i+1])$ since R is deterministic, this means that f cannot be a winning strategy: a contradiction. In the second case, the target of the E_R transition (say, $(\lambda'[j], \rho'[j+1]))$ is a reachable state and must use strictly less than n E_R transitions in accepting any outcome from $\lambda'[j]$ or else

it would not be true that n was the maximum needed from $(q, (t, \phi))$. This means that by the inductive hypothesis, $(\lambda'[j], \rho'[j+1]) \in W^{n-1}$, another contradiction. Thus, $(q, (t, \phi))$ satisfies Equation 15.

Theorem 3. *For any game G, with a winning condition specified by a Büchi automaton B, if B's generalised Rabin expansion is trivially determinisable then all winning states for P in Q_C satisfy $\mu w.\nu z.\mu s.\pi_1((pre_P(s \vee (z \wedge F_C), w) \wedge (Q_{CP} \times Q_{CA})) \vee (pre_A(s \vee (z \wedge F_C), w) \wedge (Q_{CA} \times Q_{CP})))$.*

This theorem follows from Lemma 5, the definitions of the generalised Rabin expansion of B, pre_{RP} and pre_{RA} and the definition of the main algorithm.

A.4 Proof of Theorem 4

Theorem 4. *Let ψ be an LTL formula over a set, $\mathcal{P}$, of atomic propositions. Let T be the symbolic tableau automaton for ψ constructed as above. For any $\psi_1 \in \mathrm{sub}(\psi)$, any $t \in \{r, o\}$ and any path, π, in T, for all $i \geq 0$ if $\pi[i] \in \mathrm{sat}(\psi_1^t)$ and $\pi[i, \infty]$ satisfies $\mathrm{fsat}(\psi_1^t)$ then $\pi[i, \infty] \vDash \psi_1$.*

Proof. We prove for all $t \in \{r, o\}$ by induction over the structure of ψ_1.

- Case $\psi_1 \in \Pi$: $\pi[i] \in \mathrm{sat}(\psi_1^t) \Leftrightarrow \pi[i] \vDash \psi_1 \Leftrightarrow \pi[i, \infty] \vDash \psi_1$.
- Case $(\psi_1 \vee \psi_2)^t$: We split on the cases for $\pi[i] \in \mathrm{sat}((\psi_1 \vee \psi_2)^t)$. In all cases, $\pi[i, \infty]$ satisfies $\mathrm{fsat}(\psi_1^t) \wedge \mathrm{fsat}(\psi_2^t) \wedge (\mathrm{fsat}(\psi_1^o) \vee \mathrm{fsat}(\psi_2^o))$:
  - $\pi[i] \in \mathrm{sat}(\psi_1^t)$: By induction, $\pi[i, \infty] \vDash \psi_1$ so $\pi[i, \infty] \vDash \psi_1 \vee \psi_2$.
  - $\pi[i] \in \mathrm{sat}(\psi_2^t)$: By induction, $\pi[i, \infty] \vDash \psi_2$ so $\pi[i, \infty] \vDash \psi_1 \vee \psi_2$.
  - $\pi[i] \in \mathrm{sat}(\psi_1^o) \wedge \mathrm{sat}(\psi_2^o)$: By induction, $\pi[i, \infty] \vDash \psi_1$ or $\pi[i, \infty] \vDash \psi_2$ so $\pi[i, \infty] \vDash \psi_1 \vee \psi_2$.
- Case $(\psi_1 \wedge \psi_2)^t$: $\pi[i] \in \mathrm{sat}(\psi_1^t)$ and $\pi[i] \in \mathrm{sat}(\psi_2^t)$ and $\pi[i, \infty]$ satisfies $\mathrm{fsat}(\psi_1^t) \wedge \mathrm{fsat}(\psi_2^t)$ so by induction $\pi[i, \infty] \vDash \psi_1 \wedge \psi_2$.
- Case $(X\psi_1)^t$: By the definition of sat and the transition relation, $\pi[i] \in \mathrm{sat}((X\psi_1)^t)$ implies that $\pi[i+1] \in \mathrm{sat}(\psi_1^t)$. $\pi[i, \infty]$ satisfies $\mathrm{fsat}(\psi_1)$ so, by induction $\pi[i+1] \vDash \psi_1$. Thus, $\pi[i, \infty] \vDash X\psi_1$.
- Case $(\psi_1 \mathcal{U} \psi_2)^t$: $\pi[i] \in \mathrm{sat}(\psi_2^t)$ or $\pi[i] \in \mathrm{sat}(\psi_1^t) \wedge \mathrm{sat}((X(\psi_1 \mathcal{U} \psi_2))^t)$. In either case $\pi[i, \infty]$ satisfies $\mathrm{fsat}(\psi_1^t) \wedge \mathrm{fsat}(\psi_2^t)$ and $\exists^\infty j \geq i.\pi[j] \in \mathrm{sat}(\psi_2^t) \vee \neg(X(\psi_1 \mathcal{U} \psi_2))^t$.
  - If $\pi[i] \in \mathrm{sat}(\psi_2^t)$, then by induction $\pi[i, \infty] \vDash \psi_2$ and thus $\pi[i, \infty] \vDash \psi_1 \mathcal{U} \psi_2$.
  - If $\pi[i] \in \mathrm{sat}(\psi_1^t) \wedge \mathrm{sat}((X(\psi_1 \mathcal{U} \psi_2))^t)$ then, again by induction, $\pi[i, \infty] \vDash \psi_1$. By the definition of sat and the transition relation $\pi[i+1] \in \mathrm{sat}((\psi_1 \mathcal{U} \psi_2)^t)$. We can make the same argument that $\pi[i+1] \in \mathrm{sat}(\psi_2^t)$ or $\pi[i+1] \in \mathrm{sat}(\psi_1^t) \wedge \mathrm{sat}((X(\psi_1 \mathcal{U} \psi_2))^t)$. Since $\exists^\infty j \geq i.\pi[j] \in \mathrm{sat}(\psi_2^t) \vee \neg(X(\psi_1 \mathcal{U} \psi_2))^t$, there must eventually be some $k > i$ such that $\pi[k] \in \mathrm{sat}(\psi_2)$. Take k to be the smallest such integer; since $\forall i_2 \in [i, k-1]$ we can prove that $\pi[i_2, \infty] \vDash \psi_1$, we conclude that $\pi[i, \infty] \vDash \psi_1 \mathcal{U} \psi_2$.

– Case $(\psi_1 \mathcal{R} \psi_2)^t$: By definition of sat, either $\pi[i] \in \text{sat}(\psi_1^t) \wedge \text{sat}(\psi_2^t)$ or $\pi[i] \in \text{sat}(\psi_2^t) \wedge \text{sat}((X(\psi_1 \mathcal{R} \psi_2))^t)$. In both cases $\pi[i, \infty]$ satisfies $\text{fsat}(\psi_1^t) \wedge \text{fsat}(\psi_2^t)$.

- If $\pi[i] \in \text{sat}(\psi_1^t) \wedge \text{sat}(\psi_2^t)$, then by induction $\pi[i, \infty] \vDash \psi_1$ and $\pi[i, \infty] \vDash \psi_2$ so $\pi[i, \infty] \vDash \psi_1 \mathcal{R} \psi_2$.
- If $\pi[i] \in \text{sat}(\psi_2^t) \wedge \text{sat}((X(\psi_1 \mathcal{R} \psi_2))^t)$, then by induction $\pi[i, \infty] \vDash \psi_2$ and $\pi[i + 1] \in \text{sat}((\psi_1 \mathcal{R} \psi_2)^t)$. We can re-use the same argument for $\pi[i + 1]$ and all further cases. Either there exists some $j > i$ such that $\pi[j] \in \text{sat}(\psi_1^t) \wedge \text{sat}(\psi_2^t)$, or there does not. If such a j exists, take it to be the smallest one; we prove by induction that $\pi[j, \infty] \vDash \psi_1 \wedge \psi_2$ and since $\forall k \in [i, j]$ we can prove that $\pi[k, \infty] \vDash \psi_2$ we have proven that $\pi[i, \infty] \vDash \psi_1 \mathcal{R} \psi_2$. If there is no such k, then for all $j \geq i$ we can prove that $\pi[i, \infty] \vDash \psi_2$ so $\pi[i, \infty] \vDash \psi_1 \mathcal{R} \psi_2$.

$\square$

A.5 Proof of Proposition 1

Proposition 1. *The tableau for LTL formulae of the form $G(p \wedge (r_0 \Rightarrow Fs_0) \wedge \ldots \wedge (r_n \Rightarrow Fs_n))$ (where p, r_i, and s_i are propositional formulae) is always trivially determinisable.*

Proof. Let ψ be the entire request-response formula. We prove the proposition by showing that the tableau generated for ψ can be trivially determinised to a language equivalent tableau, T_D. Instead of using $\text{sat}(\psi)$ for the initial states and Equation 8 for the transition relation of the tableau we make a number of changes. First, we assume that optional variables are always false so any variables mentioned are in their required forms. Then, we use $X\psi \wedge p \wedge \bigwedge_{i \in [0,n]} (XFs_i \Leftrightarrow r_i \wedge \neg s_i)$ for the initial states and $X\psi \wedge p' \wedge \bigwedge_{i \in [0,n]} (XFs_i' \Leftrightarrow (XFs_i \vee r_i') \wedge \neg s_i')$ for the transition. The initial formula is clearly deterministic for a given assignment to the labels of the tableau $\{r_0, s_0, r_1, s_1, \ldots\}$ and the transition is deterministic for a given assignment to current-state tableau variables and the next-state labels $\{XFs_0, r_0', s_0', XFs_1, r_1', s_1', \ldots\}$, thus T_D is deterministic. The transition formula is non-blocking and both formulae imply the corresponding definitions in Section 4[2]. We now prove that this deterministic tableau still accepts any path satisfying ψ by showing that for any path $\lambda \vDash \psi$, the run $\rho = run(T_D, \lambda)$ is fair with respect to $\text{fsat}(\psi)$. Specifically, this means that $\forall i \in [0, n] \; \exists^\infty j . XFs_i \notin \rho[j]$. By the definition of the initial states and transition relation, $\forall i \in [0, n], j \geq 0 \; XFs_i \in \rho[j] \Rightarrow \exists k \leq j . r_i \in \rho[k] \wedge \forall m \in [k, j] \; s_i \notin \rho[m]$. Since $\lambda \vDash \psi$, $\lambda[k, \infty] \vDash r_i \Rightarrow Fs_i$ so there exists $n > j$ such that $s_i \in \rho[n]$ and by the definition of the transition relation $XFs_i \notin \rho[n]$. So the fairness constraints are satisfied. Thus the tableau for ψ can be trivially determinised. $\square$

---

[2] This would not be the case had we not weakened the condition on the transition relation (Equation 8) from the $\Leftrightarrow$ definition of [4] to $\Rightarrow$