BRICS

Basic Research in Computer Science

Dynamic Linear Time Temporal Logic

Jesper G. Henriksen
P. S. Thiagarajan

BRICS Report Series **RS-97-8**

ISSN 0909-0878 **April 1997**

See back inner page for a list of recent BRICS Report Series publications.
Copies may be obtained by contacting:

BRICS
Department of Computer Science
University of Aarhus
Ny Munkegade, building 540
DK–8000 Aarhus C
Denmark

Telephone: +45 8942 3360
Telefax: +45 8942 3255
Internet: BRICS@brics.dk

BRICS publications are in general accessible through the World Wide Web and anonymous FTP through these URLs:

`http://www.brics.dk`
`ftp://ftp.brics.dk`
This document in subdirectory `RS/97/8/`

Dynamic Linear Time Temporal Logic

Jesper Gulmann Henriksen
BRICS[*], Department of Computer Science,
University of Aarhus, Denmark
email: gulmann@brics.dk

P. S. Thiagarajan[†]
SPIC Mathematical Institute, Madras, India
email: pst@smi.ernet.in

April 1997

Abstract

A simple extension of the propositional temporal logic of linear time is proposed. The extension consists of strengthening the until operator by indexing it with the regular programs of propositional dynamic logic (PDL). It is shown that DLTL, the resulting logic, is expressively equivalent to S1S, the monadic second-order theory of ω-sequences. In fact a sublogic of DLTL which corresponds to propositional dynamic logic with a linear time semantics is already as expressive as S1S. We pin down in an obvious manner the sublogic of DLTL which correponds to the first order fragment of S1S. We show that DLTL has an exponential time decision procedure. We also obtain an axiomatization of DLTL. Finally, we point to some natural extensions of the approach presented here for bringing together propositional dynamic and temporal logics in a linear time setting.

[*]**B**asic **R**esearch **i**n **C**omputer **S**cience,
Centre of the Danish National Research Foundation.

[†]Part of this work was done while visiting BRICS.
Part of this work has been supported by IFCPAR Project 1502-1.

1 Introduction

We present here a simple extension of the propositional temporal logic of linear time (LTL). The basic idea is to strengthen the until modality by indexing it with the regular programs of propositional dynamic logic (PDL). The resulting logic, called dynamic linear time temporal logic (DLTL), is easy to handle. It has the full expressive power of the monadic second order theory of ω-sequences (S1S). Indeed a sublogic of DLTL is already equal in expressive power to S1S. A pleasant feature of this sublogic is that it is just PDL operating in a linear time framework.

In addition to our expressiveness results we show that DLTL has an exponential time decision procedure. We also extend the well known axiomatization of PDL [9] to obtain an axiomatization of DLTL.

Our work may be viewed from two different perspectives. The first one is from the standpoint of process logics [6, 14, 16] which attempt a rapprochement between dynamic and temporal logics. However the study of process logics is committed to viewing dynamic logic as a restricted kind of a *branching time* temporal logic. One then attempts to bring in some additional mechanisms for talking about the computational paths of a Kripke frame in a more sophisticated fashion. Our point of departure from this line of work consists of merging, in a very simple way, dynamic logic and temporal logic in a *linear time* setting.

The second perspective from which our work may be viewed has to do with attempts to augment the limited expressive power of LTL. One route consists of permitting quantification over atomic propositions. The resulting logic called QPTL [18] has the expressive power of S1S and has a decision procedure of non-elementary complexity. The second route consists of augmenting LTL with the so called automaton connectives. Indeed this additional feature is so powerful that the next and until modalities become derived ones. The resulting logic called ETL [23] is equal in expressive power to S1S while admitting an exponential time decision procedure. To be precise, the version of ETL we have in mind is the one referred to as $\mathrm{ETL}_{\mathrm{f}}$ in [21] but for convenience we will call it just ETL here.

Our logic is, in spirit, inspired by ETL and in fact it may appear to be at first sight just a reformulation of ETL with some cosmetic changes. This however has to do with the instinctive identification one makes between finite state automata and regular expressions. In fact DLTL is quite different in terms of the mechanisms it offers for structuring formulas and we feel that it is more transparent and easier to work with. The results and the proofs we present here are designed to support this claim. Our approach also leads to smooth generalizations in non-sequential settings where similar extensions in terms of ETL will be hard to cope with at least in terms of transparency.

In the next section we start with an action-based version of LTL in order to fix terminology and to introduce a key feature of DLTL. In Section 3 we present DLTL and its semantics. This is then followed by a more detailed assessment of the similarities and the differences between ETL and DLTL.

In Section 4 we prove the decidability of DLTL by reducing it to the emptiness problem for Büchi automata. In Section 5 we show that $DLTL^-$, a sublogic of DLTL, has the same expressive power as S1S. This then easily leads to the conclusion that DLTL, $DLTL^-$ and S1S are all equal in expressive power. We then establish similar results for the first-order fragment of S1S with the help of the "star-free" fragments of DLTL and $DLTL^-$.

In Section 6, we extend the axiomatization of PDL and the completeness proof in [9] to obtain finitary axiomatizations of DLTL and $DLTL^-$. In the final section we point to two natural generalizations in the setting of distributed systems. These generalizations appear to be eminently accessible and offer additional support to our belief that the synthesis of dynamic and temporal logics in a linear time framework as pursued here is a fruitful one.

2 Linear Time Temporal Logic

One key feature of the syntax and semantics of our temporal logic is the treatment of *actions* as first class objects.

The usual treatment of LTL [10, 15] is based on *states* viewed as subsets of a finite set of atomic propositions. We wish to bring in actions explicitly for a number of reasons. Firstly, in a state-based approach there is a mismatch between the models of a formula and the language accepted by the associated ω-automaton. A model will consist of a sequence of states whereas the automaton will accept or reject a sequence of actions. Thus the alphabet of the automaton will correspond to the possible states of the model. However this is not a major hurdle and it leads, at worst, to minor technical overheads. The more important point is that it is difficult to define synchronized products of sequential components in a purely state-based setting. Since this is a very common and useful method of specifying distributed behaviours, we feel that it is handy to work with logics in which both states and actions can be treated on an equal footing.

As a vehicle for introducing some terminology and for later comparisons we shall first introduce an action-based version of LTL denoted LTL(Σ). We begin with some notations.

Through the rest of the paper we fix a finite non-empty alphabet Σ. We let a, b range over Σ and refer to members of Σ as actions. Σ^* is the set of finite words and Σ^ω is the set of infinite words generated by Σ with $\omega = \{0, 1, \ldots\}$. We set $\Sigma^\infty = \Sigma^* \cup \Sigma^\omega$ and denote the null word by ε. We let σ, σ' range over Σ^ω and τ, τ', τ'' range over Σ^*. Finally $\preceq$ is the usual prefix ordering defined over Σ^* and for $u \in \Sigma^\infty$, we let $\mathrm{prf}(u)$ be the set of finite prefixes of u.

Next we fix a countable set of atomic propositions $P = \{p_1, p_2, \ldots\}$ and let p, q range over P. The set of formulas of LTL(Σ) is then given by the syntax:

$$\mathrm{LTL}(\Sigma) ::= p \mid \sim\alpha \mid \alpha \vee \beta \mid \langle a \rangle \alpha \mid \alpha\, \mathcal{U}\, \beta.$$

Through the rest of this section α, β will range over LTL(Σ). The modality $\langle a \rangle$ is an action-indexed version of the next state modality of LTL as we shall soon see.

A model of LTL(Σ) is a pair $M = (\sigma, V)$ where $\sigma \in \Sigma^\omega$ and $V : \mathrm{prf}(\sigma) \longrightarrow 2^P$ is a valuation function. Let $M = (\sigma, V)$ be a model, $\tau \in \mathrm{prf}(\sigma)$ and α be a formula. Then $M, \tau \models \alpha$ will stand for α being satisfied at τ in M. This notion is defined inductively in the expected manner.

- $M, \tau \models p$ iff $p \in V(\tau)$.
- $M, \tau \models \sim\alpha$ iff $M, \tau \not\models \alpha$.
- $M, \tau \models \alpha \vee \beta$ iff $M, \tau \models \alpha$ or $M, \tau \models \beta$.
- $M, \tau \models \langle a \rangle \alpha$ iff $\tau a \in \mathrm{prf}(\sigma)$ and $M, \tau a \models \alpha$.
- $M, \tau \models \alpha\, \mathcal{U}\, \beta$ iff there exists τ' such that $\tau\tau' \in \mathrm{prf}(\sigma)$ and $M, \tau\tau' \models \beta$. Moreover for every τ'' such that $\varepsilon \preceq \tau'' \prec \tau'$, it is the case that $M, \tau\tau'' \models \alpha$.

For further reference we note that the next-state modality of LTL is definable via $O\alpha \overset{\Delta}{\Longleftrightarrow} \bigvee_{a \in \Sigma} \langle a \rangle \alpha$. It is well known [4, 8] that $LTL(\Sigma)$ is expressively equivalent to the first-order theory of sequences. Hence this temporal logic has definite limitations of expressive power. In particular, as pointed out by Wolper in a state-based setting [22], the property "p holds at every even position" is not definable in this logic. The dynamic version of LTL(Σ) that we introduce in the next section is designed to overcome this kind of expressive limitation.

3 Dynamic Linear Time Temporal Logic

Our extension of LTL(Σ) basically consists of indexing the until operator with the programs of PDL (e.g. [3, 5]). We start by defining the set of programs (regular expressions) generated by Σ. This set is denoted by Prg(Σ) and its syntax is given by:

$$\mathrm{Prg}(\Sigma) ::= a \mid \pi_0 + \pi_1 \mid \pi_0 ; \pi_1 \mid \pi^*.$$

Here and elsewhere, π, π' with or without subscripts will range over Prg(Σ). With each program we associate a set of finite words via the map $||\cdot|| : \mathrm{Prg}(\Sigma) \longrightarrow 2^{\Sigma^*}$. This map is defined in the standard fashion.

- $||a|| = \{a\}$.
- $||\pi_0 + \pi_1|| = ||\pi_0|| \cup ||\pi_1||$.
- $||\pi_0 ; \pi_1|| = \{\tau_0\tau_1 \mid \tau_0 \in ||\pi_0|| \text{ and } \tau_1 \in ||\pi_1||\}$.
- $||\pi^*|| = \bigcup_{i\in\omega} ||\pi^i||$, where
  - $||\pi^0|| = \{\varepsilon\}$ and
  - $||\pi^{i+1}|| = \{\tau_0\tau_1 \mid \tau_0 \in ||\pi|| \text{ and } \tau_1 \in ||\pi^i||\}$ for every $i \in \omega$.

Again we fix a countable set of atomic propositions $P = \{p_1, p_2, \ldots\}$ and let p, q range over P. Then the set of formulas of DLTL(Σ) is given by the following syntax.

$$\mathrm{DLTL}(\Sigma) ::= p \mid \sim\alpha \mid \alpha \vee \beta \mid \alpha\,\mathcal{U}^{\pi}\beta.$$

Here and throughout the rest of the paper we take α, β to range over DLTL(Σ). The notion of a model is as in the case of LTL(Σ). So let $M = (\sigma, V)$ be a model, $\tau \in \mathrm{prf}(\sigma)$ and $\alpha \in \mathrm{DLTL}(\Sigma)$. Then $M, \tau \models \alpha$ is defined inductively.

- $M, \tau \models p$ iff $p \in V(\tau)$.

- $M, \tau \models \sim\alpha$ iff $M, \tau \not\models \alpha$.
- $M, \tau \models \alpha \vee \beta$ iff $M, \tau \models \alpha$ or $M, \tau \models \beta$.
- $M, \tau \models \alpha\, \mathcal{U}^{\pi}\beta$ iff there exists $\tau' \in ||\pi||$ such that $\tau\tau' \in \mathrm{prf}(\sigma)$ and $M, \tau\tau' \models \beta$. Moreover, for every τ'' such that $\varepsilon \preceq \tau'' \prec \tau'$, it is the case that $M, \tau\tau'' \models \alpha$.

Thus DLTL(Σ) is obtained form LTL(Σ) by strengthening the until operator. To satisfy $\alpha\ \mathcal{U}^{\pi}\beta$, one must satisfy $\alpha\mathcal{U}\beta$ along some finite stretch of behaviour which is in the (linear time) behaviour of the program π.

As usual, $\alpha \in$ DLTL(Σ) is *satisfiable* iff there exist a model $M = (\sigma, V)$ and $\tau \in \mathrm{prf}(\sigma)$ such that $M, \tau \models \alpha$. The *satisfiability problem* is, given α, to decide whether α is satisfiable.

Apart from the conventional derived propositional connectives such as $\wedge, \supset$ and $\equiv$ the derived modality $\langle\pi\rangle$ and its dual $[\pi]$ will play an important role in the sequel. They are defined as follows:

- $\top \stackrel{\Delta}{\Longleftrightarrow} p_1 \vee \sim p_1$. Recall that $P = \{p_1, p_2, \ldots\}$.
- $\langle\pi\rangle\alpha \stackrel{\Delta}{\Longleftrightarrow} \top\, \mathcal{U}^{\pi}\alpha$.
- $[\pi]\alpha \stackrel{\Delta}{\Longleftrightarrow} \sim\langle\pi\rangle\sim\alpha$.

Suppose $M = (\sigma, V)$ is a model and $\tau \in \mathrm{prf}(\sigma)$. It is easy to see that $\sigma, \tau \models \langle\pi\rangle\alpha$ iff there exists $\tau' \in ||\pi||$ such that $\tau\tau' \in \mathrm{prf}(\sigma)$ and $\sigma, \tau\tau' \models \alpha$. It is also easy to see that $\sigma, \tau \models [\pi]\alpha$ iff for every $\tau' \in ||\pi||$, if $\tau\tau' \in \mathrm{prf}(\sigma)$ then $\sigma, \tau\tau' \models \alpha$. In this sense, the program modalities of PDL acquire a linear time semantics in the present setting.

Note that $a \in \Sigma$ is a member of Prg(Σ) and hence $\langle a\rangle\alpha$ is a derived modality. It is also easy to see that the until operator of LTL(Σ) can be obtained via: $\alpha\mathcal{U}\beta \stackrel{\Delta}{\Longleftrightarrow} \alpha\ \mathcal{U}^{\Sigma^*}\beta$. Thus LTL($\Sigma$) is a fragment of DLTL($\Sigma$). To see that DLTL($\Sigma$) is strictly more expressive than LTL(Σ), let $\Sigma = \{a_1, a_2, \ldots, a_n\}$ and $\pi_{ev} = ((a_1+a_2+\ldots+a_n);(a_1+a_2+\ldots+a_n))^*$.

It is now easy to see that $\alpha_{ev} = [\pi_{ev}]p$ is a specification of the property "p holds at every even position".

We shall close out the section by briefly discussing the key differences between DLTL(Σ) and ETL, the extension of LTL proposed by Wolper [22]. As mentioned earlier, the version of ETL we shall consider here is the one called $\mathrm{ETL_f}$ in [21]. We shall present a simplified form of ETL so as to stay close to DLTL. First we fix an enumeration of $\Sigma = \{a_1, a_2, \ldots, a_n\}$. The syntax of the logic that we shall name as ETL(Σ) is then given by:

$$\mathrm{ETL}(\Sigma) ::= p \mid \sim\phi \mid \phi \vee \phi' \mid \mathcal{A}(\phi_0, \phi_1, \ldots, \phi_n).$$

Here $\phi, \phi', \phi_0, \phi_1, \ldots, \phi_n$ are assumed to be formulas of ETL(Σ). Further, $\mathcal{A}$ is a finite state automaton of the form $\mathcal{A} = (Q, \longrightarrow, Q_{in}, F)$ with $\longrightarrow \subseteq Q \times \Sigma \times Q$ as the transition relation, $Q_{in} \subseteq Q$ as the initial states and $F \subseteq Q$ as the accepting states. In what follows we let $\mathcal{L}(\mathcal{A})$ be the language of finite words accepted by $\mathcal{A}$. Moreover we shall assume for the sake of convenience that $\varepsilon \notin \mathcal{L}(\mathcal{A})$ for each formula of the form $\mathcal{A}(\phi_0, \phi_1, \ldots, \phi_n)$.

A model for ETL(Σ) is, as before, a pair $M = (\sigma, V)$ with $V : \mathrm{prf}(\sigma) \longrightarrow 2^P$. Let $\tau \in \mathrm{prf}(\sigma)$. Then $M, \tau \models \phi$ is defined via:

- $M, \tau \models p$ iff $p \in V(\tau)$.
- $M, \tau \models \sim\phi$ iff $M, \tau \not\models \phi$.
- $M, \tau \models \phi \vee \phi'$ iff $M, \tau \models \phi$ or $M, \tau \models \phi'$.
- $M, \tau \models \mathcal{A}(\phi_0, \phi_1, \ldots, \phi_n)$ iff there exists $a_{i_1} a_{i_2} \ldots a_{i_m} \in \mathcal{L}(\mathcal{A})$ such that the following conditions are satisfied:
  - — $i_1, i_2, \ldots, i_m \in \{1, 2, \ldots, n\}$. (recall that $\Sigma = \{a_1, a_2, \ldots, a_n\}$).
  - — $\tau a_{i_1} a_{i_2} \cdots a_{i_m} \in \mathrm{prf}(\sigma)$.
  - — $M, \tau \models \phi_0$ and $M, \tau a_{i_1} \cdots a_{i_j} \models \phi_{i_j}$ for $1 \leq j \leq m$.

Though the technical details are somewhat different, ETL(Σ) captures the spirit of the logic presented in [21]. The key drawback of ETL(Σ), as we see it, lies in its lack of structuring principles for forming compound formulas. Stated differently, the only mechanism that ETL(Σ) has — apart from the boolean connectives — to form compound formulas is *by nesting* the automaton formulas. Thus a typical compound formula would look like:

$$\mathcal{A}^1(\phi_0^1, \mathcal{A}^2(\phi_0^2, \phi_1^2, \mathcal{A}^3(\phi_0^3, \ldots, \phi_n^3), \phi_3^2, \ldots, \phi_n^2), \phi_2^1, \ldots, \phi_n^1).$$

In contrast, DLTL(Σ) adds to the familiar mechanisms of LTL an *orthogonal* and well-understood component; namely, the language of regular expressions. Equally important, this orthogonal component is formulated purely in terms of Σ and not in terms of arbitrary formulas as is the case of ETL. In fact, ETL, as formulated in [21] has an uncontrolled amount of "external" elements in the sense that the states and the alphabets of the automata which are used to write down the automaton formulas have little to do with the logic under consideration.

4 A Decision Procedure for DLTL

The goal here is to show that the satisfiability problem for DLTL(Σ) can be solved in deterministic exponential time. This will be achieved by effectively constructing for each $\alpha \in$ DLTL(Σ), a Büchi automaton $\mathcal{B}_\alpha$ such that the language of ω-words accepted by $\mathcal{B}_\alpha$ is non-empty iff α is satisfiable.

We shall use the terminology adopted in [20] for dealing with Büchi automata. In particular, a Büchi automaton over Σ is a tuple $\mathcal{B} = (Q, \longrightarrow, Q_{in}, F)$ where:

- Q is a finite non-empty set of states.
- $\longrightarrow \subseteq Q \times \Sigma \times Q$ is a transition relation.
- $Q_{in} \subseteq Q$ is a set of initial states.
- $F \subseteq Q$ is a set of accepting states.

Let $\sigma \in \Sigma^\omega$. Then a run of $\mathcal{B}$ over σ is a map $\rho : \mathrm{prf}(\sigma) \longrightarrow Q$ such that:

- $\rho(\varepsilon) \in Q_{in}$.
- $\rho(\tau) \xrightarrow{a} \rho(\tau a)$ for each $\tau a \in \mathrm{prf}(\sigma)$.

The run ρ is accepting iff $\inf(\rho) \cap F \neq \emptyset$ where $\inf(\rho) \subseteq Q$ is given by $q \in \inf(\rho)$ iff $\rho(\tau) = q$ for infinitely many $\tau \in \mathrm{prf}(\sigma)$. Finally $\mathcal{L}(\mathcal{B})$, the language of ω-words accepted by $\mathcal{B}$, is:

$$\mathcal{L}(\mathcal{B}) = \{\sigma \mid \exists \text{ an accepting run of } \mathcal{B} \text{ over } \sigma\}.$$

Through the rest of the section we fix a formula α_0. To construct $\mathcal{B}_{\alpha_0}$ we first define the (Fischer-Ladner) closure of α_0 as follows. $cl(\alpha_0)$ is the least set of formulas that satisfies:

- $\alpha_0 \in cl(\alpha_0)$.

- If $\sim\beta \in cl(\alpha_0)$ then $\beta \in cl(\alpha_0)$.
- If $\alpha \vee \beta \in cl(\alpha_0)$ then $\alpha, \beta \in cl(\alpha_0)$.
- If $\alpha\, \mathcal{U}^{\pi} \beta \in cl(\alpha_0)$ then $\alpha, \beta \in cl(\alpha_0)$.

Now $CL(\alpha_0)$, the *closure* of α_0, is defined to be:

$$CL(\alpha_0) = cl(\alpha_0) \cup \{\sim\beta \mid \beta \in cl(\alpha_0)\}.$$

In what follows $\sim\sim\beta$ will be identified with β. Moreover, until the end of the section, all the formulas that we encounter — unless stated otherwise — will be assumed to be members of $CL(\alpha_0)$. For convenience, we shall often write CL instead of $CL(\alpha_0)$.

$A \subseteq CL$ is called an *atom* iff it is a subset of CL satisfying:

- $\beta \in A$ iff $\sim\beta \notin A$.
- $\alpha \vee \beta \in A$ iff $\alpha \in A$ or $\beta \in A$.
- If $\beta \in A$ and $\varepsilon \in ||\pi||$ then $\alpha\, \mathcal{U}^{\pi} \beta \in A$.

$AT(\alpha_0)$ is the set of atoms and again we shall often write AT instead of $AT(\alpha_0)$. Next we define $Req(\alpha_0)$, the set of until requirements of α_0, to be the subset of CL given by:

$$Req(\alpha_0) = \{\alpha\, \mathcal{U}^{\pi} \beta \mid \alpha\, \mathcal{U}^{\pi} \beta \in CL\}.$$

We shall write Req instead $Req(\alpha_0)$ and take ξ, ξ' to range over Req. For each $\xi = \alpha\, \mathcal{U}^{\pi} \beta \in Req$ we fix a finite state automaton $\mathcal{A}_\xi$ such that $\mathcal{L}(\mathcal{A}_\xi) = ||\pi||$ where $\mathcal{L}(\mathcal{A}_\xi)$ is the language of finite words accepted by $\mathcal{A}_\xi$. We shall assume each such $\mathcal{A}_\xi$ is of the form $\mathcal{A}_\xi = (Q_\xi, \longrightarrow_\xi, I_\xi, F_\xi)$ where Q_ξ is the set of states, $\longrightarrow_\xi \subseteq Q_\xi \times \Sigma \times Q_\xi$ is the transition relation, $I_\xi \subseteq Q_\xi$ is the set of initial states and $F_\xi \subseteq Q_\xi$ is the set of final states. Without loss of generality, we shall assume that $\xi \neq \xi'$ implies $Q_\xi \cap Q_{\xi'} = \emptyset$ for every $\xi, \xi' \in Req$. We set $Q = \bigcup_{\xi \in Req} Q_\xi$ and $\widehat{Q} = Q \times \{0, 1\}$.

The Büchi automaton $\mathcal{B}_{\alpha_0}$ associated with α_0 (from now on denoted as $\mathcal{B}$) can now be defined as

$$\mathcal{B} = (S, \Longrightarrow, S_{in}, F),$$

where the various components of $\mathcal{B}$ are specified as follows:

(1) $S \subseteq AT \times 2^Q \times 2^{\widehat{Q}} \times \{0,1\} \times \{\downarrow, \checkmark\}$ such that $(A, X, \widehat{X}, x, f) \in S$ iff the following conditions are satisfied for each $\xi = \alpha\, \mathcal{U}^\pi \beta$:

   (i) If $\beta \in A$ then $F_\xi \subseteq X$. (Recall that $\mathcal{A}_\xi = (Q_\xi, \longrightarrow_\xi, I_\xi, F_\xi))$.

   (ii) If $\alpha \in A$ and $q \in X$ for some $q \in I_\xi$ then $\alpha\, \mathcal{U}^\pi \beta \in A$.

   (iii) If $\alpha\, \mathcal{U}^\pi \beta \in A$ then either $\beta \in A$ and $\varepsilon \in ||\pi||$ or $(q, 1-x) \in \widehat{X}$ for some $q \in I_\xi$. (Note that we are considering the candidate $(A, X, \widehat{X}, x, f)$ for membership in S).

   (iv) If $(q, z) \in \widehat{X}$ with $q \notin F_\xi$ or $\beta \notin A$ then $\alpha \in A$.

(2) The transition relation $\Longrightarrow \subseteq S \times \Sigma \times S$ is defined as follows:

$$(A, X, \widehat{X}, x, f) \stackrel{a}{\Longrightarrow} (B, Y, \widehat{Y}, y, g)$$

iff the following conditions are satisfied for each $\xi = \alpha\, \mathcal{U}^\pi \beta$:

   (i) Suppose $q' \in Q_\xi \cap Y$ and $q \stackrel{a}{\longrightarrow}_\xi q'$ and $\alpha \in A$. Then $q \in X$.

   (ii) Suppose $(q, z) \in \widehat{X}$ with $q \in Q_\xi$. Suppose further that $q \notin F_\xi$ or $\beta \notin A$. Then $(q', z) \in \widehat{Y}$ for some q' with $q \stackrel{a}{\longrightarrow}_\xi q'$.

   (iii) If $f = \checkmark$ then $(y, g) = (1 - x, \downarrow)$. If $f = \downarrow$ then,

$$(y, g) = \begin{cases} (x, \downarrow), & \text{if there exists } (q, x) \in \widehat{X} \text{ such that} \\ & q \notin F_\xi \text{ or } \beta \notin A. \\ (x, \checkmark), & \text{otherwise.} \end{cases}$$

(3) $S_{in} = \{(A, X, \widehat{X}, x, f) \mid \alpha_0 \in A \text{ and } (x, f) = (0, \checkmark)\}$.

(4) $F = \{(A, X, \widehat{X}, x, f) \mid f = \checkmark\}$

We wish to prove that α_0 is satisfiable iff $\mathcal{L}(\mathcal{B}) \neq \emptyset$. Afterwards we will argue that the size of $\mathcal{B}$ can be chosen to be at most exponential in the size of α_0.

Lemma 4.1 *Suppose $\mathcal{L}(\mathcal{B}) \neq \emptyset$. Then α_0 is satisfiable.*

Proof: Let $\sigma \in \mathcal{L}(\mathcal{B})$ and $\rho : \mathrm{prf}(\sigma) \longrightarrow S$ be an accepting run. For each $\tau \in \mathrm{prf}(\sigma)$, let $\rho(\tau) = (A_\tau, X_\tau, \widehat{X}_\tau, x_\tau, f_\tau)$. Define the model $M = (\sigma, V)$ via:

$$V(\tau) = A_\tau \cap P \text{ for all } \tau \in \mathrm{prf}(\sigma).$$

We shall prove the following intermediate result.

Claim: For all $\tau \in \mathrm{prf}(\sigma)$ and $\delta \in CL$,

$$M, \tau \models \delta \textit{ iff } \delta \in A_\tau.$$

First note that if the claim is true then Lemma 4.1 follows at once. This is so because ρ is a run of $\mathcal{B}$ and hence $\rho(\varepsilon) \in S_{in}$. But from (3) (in the definition of $\mathcal{B}$), it follows that $\alpha_0 \in A_\varepsilon$. Thus $M, \varepsilon \models \alpha_0$ and α_0 is satisfiable.

In proving the claim we will repeatedly refer to various clauses in the definition of the Büchi automaton $\mathcal{B}$.

We proceed by structural induction on δ. For the base case and the boolean connectives the claim is obvious. Hence assume that $\delta = \alpha\, \mathcal{U}^\pi \beta$.

Suppose that $M, \tau \models \alpha\ \mathcal{U}^\pi \beta$. Since $M, \tau \models \alpha\ \mathcal{U}^\pi \beta$ there exists $\tau' \in ||\pi||$ such that $\tau\tau' \in \mathrm{prf}(\sigma)$ and $M, \tau\tau' \models \beta$. Moreover, $M, \tau\tau'' \models \alpha$ for every $\tau'' \in \Sigma^*$ such that $\varepsilon \preceq \tau'' \prec \tau'$.

Suppose $\tau' = \varepsilon$. Then $\varepsilon \in ||\pi||$ and $M, \tau \models \beta$. By the induction hypothesis $\beta \in A_\tau$. From the definition of an atom it follows that $\alpha\, \mathcal{U}^\pi \beta \in A_\tau$.

So assume that $\tau' \neq \varepsilon$. Let $\xi = \alpha\, \mathcal{U}^\pi \beta$ and R be an accepting run of $\mathcal{A}_\xi$ over $\tau' = a_1 a_2 \ldots a_n$ with $R(\varepsilon) = q_0 \in I_\xi$ and $R(a_1 a_2 \ldots a_i) = q_i$ for $1 \leq i \leq n$ and $q_n \in F_\xi$. Since $M, \tau\tau' \models \beta$ we have from the induction hypothesis that $\beta \in A_{\tau\tau'}$. Hence by (1.i), $F_\xi \subseteq X_{\tau\tau'}$. Now by the definition of R we are assured that $q_{n-1} \xrightarrow{a_n}_\xi q_n$. On the other hand, the fact that $M, \tau \models \alpha\, \mathcal{U}^\pi \beta$ and the choice of τ' guarantee that $M, \tau a_1 \ldots a_{n-1} \models \alpha$ (with the convention that $\varepsilon = a_1 \ldots a_{n-1}$ in case $n = 1$). By the induction hypothesis $\alpha \in A_{\tau a_1 \ldots a_{n-1}}$, so by (2.i) and the fact that $q_n \in X_{\tau a_1 \ldots a_n}$, we have that $q_{n-1} \in X_{\tau a_1 \ldots a_{n-1}}$. In case $n \geq 2$ we repeat the above argument at q_{n-1} to conclude that $q_{n-2} \in X_{\tau a_1 \ldots a_{n-2}}$. Continuing this way we can finally arrive at $q_0 \in X_\tau$ and $\alpha \in A_\tau$. But $q_0 \in I_\xi$ and hence by (1.ii) we are assured that $\alpha\, \mathcal{U}^\pi \beta \in A_\tau$.

For the converse direction assume that $\alpha\ \mathcal{U}^\pi \beta \in A_\tau$. The are four cases to consider depending on the values of x_τ and f_τ. We will only prove the first case. The remaining cases can be resolved by very similar arguments.

So assume that $x_\tau = 0$ and $f_\tau = \downarrow$. Suppose first that $\beta \in A_\tau$ and $\varepsilon \in ||\pi||$. Then by the induction hypothesis $M, \tau \models \beta$ and hence we at once have $M, \tau \models \alpha\, \mathcal{U}^\pi \beta$. So assume that $\beta \notin A_\tau$ or $\varepsilon \notin ||\pi||$. Then by (1.iii), $(q_0, 1) \in \widehat{X}_\tau$ for some $q_0 \in I_\xi$. Suppose $q_0 \in F_\xi$. Then $\varepsilon \in ||\pi||$ and thus $\beta \notin A_\tau$. This implies, by (1.iv), that $\alpha \in A_\tau$, and by the induction hypothesis we have that $M, \tau \models \alpha$.

Now with ρ being an accepting run of $\mathcal{B}$ over σ there must exist τ_1 and τ_2 in Σ^* such that the following conditions are satisfied:

- $\tau_1 \neq \varepsilon$ and $\tau_2 \neq \varepsilon$ and $\tau\tau_1\tau_2 \in \mathrm{prf}(\sigma)$.
- $f_{\tau\tau_1} = \checkmark$ and $f_{\tau\tau_1\tau_2} = \checkmark$. (Recall the notational convention that $\rho(u) = (A_u, X_u, \widehat{X}_u, x_u, f_u)$ for each $u \in \mathrm{prf}(\sigma)$).
- For each τ_1'' and τ_2'' in Σ^*, if $\varepsilon \preceq \tau_1'' \prec \tau_1$ then $f(\tau\tau_1'') \neq \checkmark$ and if $\varepsilon \prec \tau_2'' \prec \tau_2$ then $f(\tau\tau_1\tau_2'') \neq \checkmark$.

Let $\tau_1 = a_1 a_2 \ldots a_n$ and $\tau_2 = b_1 b_2 \ldots b_m$. Now $\rho(\tau) \xRightarrow{a_1} \rho(\tau a_1)$, $\alpha\, \mathcal{U}^\pi \beta \in A_\tau$ and $(q_0, 1) \in \widehat{X}_\tau$. Moreover, we have that $q_0 \notin F_\xi$ (if $\varepsilon \notin ||\pi||$) or $\beta \notin A_\tau$. Thus by (2.ii), there exists $q_1 \in Q_\xi$ such that $q_0 \xrightarrow{a_1}_\xi q_1$ and $(q_1, 1) \in \widehat{X}_{\tau a_1}$.

Now suppose $q_1 \in F_\xi$ and $\beta \in A_{\tau a_1}$. Then $a_1 \in ||\pi||$ and by the induction hypothesis $M, \tau a_1 \models \beta$. Since $M, \tau \models \alpha$ has already been deduced we have $M, \tau \models \alpha\, \mathcal{U}^\pi \beta$. So assume that $q_1 \notin F_\xi$ or $\beta \notin A_{\tau a_1}$. Then by repeating the arguments we had above for q_0 at q_1 we can arrive at $\alpha \in A_{\tau a_1}$, and hence by the induction hypothesis $M, \tau a_1 \models \alpha$. Moreover, we can conclude that there exists $q_2 \in Q_\xi$ such that $q_1 \xrightarrow{a_2}_\xi q_2$ and $(q_2, 1) \in \widehat{X}_{\tau a_1 a_2}$. Marching down τ_1 using this sequence of arguments we will either terminate with the conclusion $M, \tau \models \alpha\, \mathcal{U}^\pi \beta$ or we will exhaust all of τ_1 while being able to conclude that there must exist states $q_0, q_1, \ldots, q_n \in Q_\xi$ such that $q_0 \xrightarrow{a_1}_\xi q_1 \xrightarrow{a_2}_\xi q_2 \ldots q_{n-1} \xrightarrow{a_n}_\xi q_n$. Furthermore, we will be able to conclude that $M, \tau\tau_1'' \models \alpha$ for every τ_1'' such that $\varepsilon \preceq \tau_1'' \prec \tau_1$. Finally, we will also be assured that $(q_n, 1) \in \widehat{X}_{\tau\tau_1}$.

Now suppose $q_n \in F_\xi$ and $\beta \in A_{\tau\tau_1}$. Then $\tau_1 \in ||\pi||$ and $M, \tau\tau_1 \models \beta$ by the induction hypothesis. Consequently $M, \tau \models \alpha\, \mathcal{U}^\pi \beta$. So assume that $q_n \notin F_\xi$ or $\beta \notin A_{\tau\tau_1}$. Then $\alpha \in A_{\tau\tau_1}$ (by (1.iv)) and hence $M, \tau\tau_1 \models \alpha$ by the induction hypothesis. Now by the choice of τ_1, we know that $(x_{\tau\tau_1}, f_{\tau\tau_1}) = (0, \checkmark)$ and hence $(x_{\tau\tau_1 b_1}, f_{\tau\tau_1 b_1}) = (1, \downarrow)$ by (2.iii). On the other hand, $\rho(\tau\tau_1) \xRightarrow{b_1} \rho(\tau\tau_1 b_1)$ implies that there exists $q_1' \in Q_\xi$ such that $q_n \xrightarrow{b_1}_\xi q_1'$ and $(q_1', 1) \in \widehat{X}_{\tau\tau_1 b_1}$. Again $q_1' \in F_\xi$ and $\beta \in A_{\tau\tau_1 b_1}$ will lead to the desired conclusion $M, \tau \models \alpha\, \mathcal{U}^\pi \beta$.

So suppose $q_1' \notin F_\xi$ or $\beta \notin A_{\tau\tau_1 b_1}$. Then as before, $\alpha \in A_{\tau\tau_1 b_1}$ and hence $M, \tau\tau_1 b_1 \models \alpha$ by induction hypothesis. By the choice of τ_2 we are assured that $m \geq 2$ because $f_{\tau\tau_1 b_1} = \downarrow$. So consider $\rho(\tau\tau_1 b_1) \xRightarrow{b_2} \rho(\tau\tau_1 b_1 b_2)$. Then again it follows easily that there must exist $q_2' \in Q_\xi$ such that $q_1' \xrightarrow{b_2}_\xi q_2'$ and $(q_2', 1) \in \widehat{X}_{\tau\tau_1 b_1 b_2}$. If $q_2' \in F_\xi$ and $\beta \in A_{\tau\tau_1 b_1 b_2}$ then we will at once obtain $M, \tau \models \alpha\, \mathcal{U}^\pi \beta$. If not, the facts that $(q_1', 1) \in \widehat{X}_{\tau\tau_1 b_1}$ and that $q_1' \notin F_\xi$ or $\beta \notin A_{\tau\tau_1 b_1}$ holds, guarantee us that $f_{\tau\tau_1 b_1 b_2} = \downarrow$ by (2.iii). Hence $m \geq 3$. Carrying on this way we will eventually exhaust all of τ_2 and while doing so, reach the desired conclusion $M, \tau \models \alpha\, \mathcal{U}^\pi \beta$. $\square$

Lemma 4.2 *Suppose α_0 is satisfiable. Then $\mathcal{L}(\mathcal{B}) \neq \emptyset$.*

Proof: Since our logic has no past modalities it is easy to see that if α_0 is satisfiable then there exists a model $M = (\sigma, V)$ such that $M, \varepsilon \models \alpha_0$. We shall show that $\sigma \in \mathcal{L}(\mathcal{B})$ by constructing a map $\rho : \mathrm{prf}(\sigma) \longrightarrow S$ so that ρ is an accepting run of $\mathcal{B}$ over σ. For each $\tau \in \mathrm{prf}(\sigma)$ we set $\rho(\tau) = (A_\tau, X_\tau, \widehat{X}_\tau, x_\tau, f_\tau)$ and define ρ in a componentwise manner.

For each $\tau \in \mathrm{prf}(\sigma)$ define A_τ via:

$$A_\tau = \{\alpha \mid M, \tau \models \alpha\}.$$

For each $\tau \in \mathrm{prf}(\sigma)$ define X_τ as follows. Suppose $\xi = \alpha\, \mathcal{U}^\pi \beta$ and $q \in Q_\xi$. Then $q \in X_\tau$ iff there exists a pair (τ', R') such that:

- $\tau\tau' \in \mathrm{prf}(\sigma)$ and $M, \tau\tau' \models \beta$.
- For every τ'', if $\varepsilon \preceq \tau'' \prec \tau'$ then $M, \tau\tau'' \models \alpha$.
- $R' : \mathrm{prf}(\tau') \longrightarrow Q_\xi$ such that $R'(\varepsilon) = q$ and $R'(\tau') \in F_\xi$ and $R'(\tau'') \xrightarrow{a}_\xi R'(\tau'' a)$ for every $\tau'' a \in \mathrm{prf}(\tau')$.

To define the remaining three components we will first define the fourth and fifth components by mutual induction. To this end we shall make use of some terminology.

We shall call the pair (τ, ξ) an *obligation in M* if $\tau \in \mathrm{prf}(\sigma)$ and $\xi = \alpha\, \mathcal{U}^\pi \beta \in \mathit{Req}$ such that $M, \tau \models \alpha\, \mathcal{U}^\pi \beta$ but $M, \tau \not\models \beta$ or $\varepsilon \notin ||\pi||$. Let (τ, ξ) be an obligation in M. We shall say that the pair (τ', R') is a *witness* for (τ, ξ) iff the following conditions are satisfied:

- $\tau\tau' \in \mathrm{prf}(\sigma)$ and $M, \tau\tau' \models \beta$ and for every τ'', $\varepsilon \preceq \tau'' \prec \tau'$ implies $M, \tau\tau'' \models \alpha$.
- $\tau' \in ||\pi||$ and $R' : \mathrm{prf}(\tau') \longrightarrow Q_\xi$ such that $R'(\varepsilon) \in I_\xi$, $R'(\tau') \in F_\xi$ and $R'(\tau'') \xrightarrow{a}_\xi R'(\tau'' a)$ for every $\tau'' a \in \mathrm{prf}(\tau')$.

Note that if (τ', R') is a witness for the obligation (τ, ξ) then $\tau' \neq \varepsilon$. We shall fix a *chronicle set* CH for M. It is a set of quadruples which satisfies the following conditions:

- If $(\tau, \xi, \tau', R') \in CH$ then (τ, ξ) is an obligation in M and (τ', R') is witness for (τ, ξ).
- If (τ, ξ) is an obligation in M then $(\tau, \xi, \tau', R') \in CH$ for some witness (τ', R') for (τ, ξ).
- If (τ, ξ, τ', R'), $(\tau, \xi, \tau'', R'') \in CH$ then $(\tau', R') = (\tau'', R'')$.

It is easy to check that CH exists. (In fact it can be chosen in a canonical manner by fixing a lexicographic order on Q_ξ for each $\xi \in Req$).

With these definitions in place, we are now prepared to define the fourth and the fifth components of ρ by induction on τ. For the base case, we set $(x_\varepsilon, f_\varepsilon) = (0, \checkmark)$. Now consider the induction step where $\tau = \tau_0 a$ and assume that $(x_{\tau'}, f_{\tau'})$ is defined for every $\tau' \in \mathrm{prf}(\tau_0)$. If $f_{\tau_0} = \checkmark$ then $(x_\tau, f_\tau) = (1 - x_{\tau_0}, \downarrow)$. Suppose $f_{\tau_0} = \downarrow$. Then $(x_\tau, f_\tau) = (x_{\tau_0}, \downarrow)$ if there exists $(\tau_1, \xi_1, \tau_1', R_1') \in CH$ such that $\tau_1 \preceq \tau_0 \prec \tau_1 \tau_1'$ and $x_{\tau_1} = 1 - x_{\tau_0}$. Otherwise, $f_\tau = \checkmark$ and $x_\tau = x_{\tau_0}$.

Finally, the third component of ρ can now be defined. For each $\tau \in \mathrm{prf}(\sigma)$, we define $\widehat{X}_\tau$ as follows. Suppose $\xi \in Req$ and $q \in Q_\xi$ and $z \in \{0, 1\}$. Then $(q, z) \in \widehat{X}_\tau$ iff there exists $(\tau_1, \xi, \tau_1', R_1') \in CH$ such that for some $\tau_1'' \in \mathrm{prf}(\tau_1')$, $\tau_1 \preceq \tau = \tau_1 \tau_1''$. Moreover, $R_1'(\tau_1'') = q$ and $x_{\tau_1} = 1 - z$.

We now wish to argue that $\rho : \mathrm{prf}(\sigma) \longrightarrow S$ is an accepting run of $\mathcal{B}$ over σ. First we shall show that ρ is well defined. Let $\tau \in \mathrm{prf}(\sigma)$ be given. We must show that $\rho(\tau) \in S$. It is easy to see that A_τ is an atom, $X_\tau \subseteq Q$, $\widehat{X}_\tau \subseteq \widehat{Q}$, $x_\tau \in \{0, 1\}$ and $f_\tau \in \{\downarrow, \checkmark\}$. We will show that $\rho(\tau)$ satisfies all the clauses of the definition of $\mathcal{B}$.

So fix some $\alpha\, \mathcal{U}^\pi \beta = \xi$. Assume initially that $\beta \in A_\tau$ and $q \in F_\xi$. Then $M, \tau \models \beta$ by definition of A_τ. Now consider the pair (τ', R') where $\tau' = \varepsilon$ and $R'(\varepsilon) = q$. From the definition of X_τ it now follows that $q \in X_\tau$. Thus $F_\xi \subseteq X_\tau$ as required by (1.i).

Next assume that $\alpha \in A_\tau$ and $q \in X_\tau$ for some $q \in I_\xi$. From the definition of X_τ it follows that there exists a pair (τ', R') such that $\tau\tau' \in \mathrm{prf}(\sigma)$ and $M, \tau\tau' \models \beta$ and $M, \tau\tau'' \models \alpha$ for every τ'' such that $\varepsilon \preceq \tau'' \prec \tau'$. Furthermore, $R' : \mathrm{prf}(\tau') \longrightarrow Q_\xi$ such that $R'(\varepsilon) = q$ and $R'(\tau') \in F_\xi$ and $R'(\tau'') \xrightarrow{a}_\xi R'(\tau'' a)$ for every $\tau'' a \in \mathrm{prf}(\tau')$. But from the assumption that $q \in I_\xi$ we have that $\tau' \in ||\pi||$, because R' is an accepting run of $\mathcal{A}_\xi$ over τ'. Consequently $M, \tau \models \alpha\, \mathcal{U}^\pi \beta$ and this leads to the conclusion that $\alpha\, \mathcal{U}^\pi \beta \in A_\tau$ as required by (1.ii).

Next assume that $\alpha\, \mathcal{U}^\pi \beta \in A_\tau$ and $\beta \notin A_\tau$ or $\varepsilon \notin ||\pi||$. Then (τ, ξ) is an obligation in M since by the definition of A_τ, $M, \tau \models \alpha\, \mathcal{U}^\pi \beta$

but $M, \tau \not\models \beta$ or $\varepsilon \notin ||\pi||$. Hence there exists $(\tau, \xi, \tau', R') \in CH$. Let $R'(\varepsilon) = q$. From the fact that (τ', R') is a witness for (τ, ξ) we have that $q \in I_\xi$. Moreover, by the definition of $\widehat{X}_\tau$ and from $\tau \preceq \tau = \tau$ (i.e. $\tau_1 = \tau$ and $\tau_1'' = \varepsilon$), it follows that $(q, 1 - x_\tau) \in \widehat{X}_\tau$ as required by (1.iii).

Finally suppose that $(q, z) \in \widehat{X}_\tau$ with $q \in Q_\xi$ such that $q \notin F_\xi$ or $\beta \notin A$. Now $(q, z) \in \widehat{X}_\tau$ implies, by the definition of $\widehat{X}_\tau$, that there exists $(\tau_1, \xi, \tau_1', R_1') \in CH$ such that for some $\tau_1'' \in \mathrm{prf}(\tau_1')$, $\tau_1 \preceq \tau = \tau_1\tau_1''$ and $R_1'(\tau_1'') = q$ and $x_{\tau_1} = 1 - z$. But (τ_1', R_1') is a witness for the obligation (τ_1, ξ) and hence $R_1'(\tau_1') \in F_\xi$ and $M, \tau_1\tau_1' \models \beta$. Since $\beta \notin A_\tau$ or $q \notin F_\xi$ it must be the case that $\tau_1'' \prec \tau_1'$ and hence $M, \tau_1\tau_1'' \models \alpha$. But then $\tau = \tau_1\tau_1''$ now leads to $\alpha \in A_\tau$ as required by (1.iv).

We have now shown that ρ is well defined. Next we wish to show that ρ is a run of $\mathcal{B}$ over σ. Since $M, \varepsilon \models \alpha_0$ we have $\alpha_0 \in A_\varepsilon$. By definition, $(x_\varepsilon, f_\varepsilon) = (0, \checkmark)$. Hence $\rho(\varepsilon) \in S_{in}$.

Now suppose $\tau a \in \mathrm{prf}(\sigma)$. We must show that $\rho(\tau) \stackrel{a}{\Longrightarrow} \rho(\tau a)$. For this purpose we fix $\alpha\, \mathcal{U}^\pi \beta = \xi \in Req$. Suppose $q, q' \in Q_\xi$ with $q' \in X_{\tau a}$ such that $q \xrightarrow{a}_\xi q'$. Further suppose $\alpha \in A_\tau$. Now $q' \in X_{\tau a}$ implies that there exists a pair (τ', R') such that $R'(\varepsilon) = q'$ and $R'(\tau') \in F_\xi$ and $R'(\tau'') \xrightarrow{b}_\xi R'(\tau'' b)$ for every $\tau'' b \in \mathrm{prf}(\tau')$. Furthermore, $M, \tau a \tau' \models \beta$ and $M, \tau a \tau'' \models \alpha$ for every τ'' such that $\varepsilon \preceq \tau'' \prec \tau'$. Now consider the pair $(a\tau', R_a')$ where $R_a' : \mathrm{prf}(a\tau') \longrightarrow Q_\xi$ is given as $R_a'(\varepsilon) = q$ and for every $\tau'' \in \mathrm{prf}(\tau')$, $R_a'(a\tau'') = R'(\tau'')$. From $M, \tau \models \alpha$ (as $\alpha \in A_\tau$ by assumption) it now follows at once that $q \in X_\tau$ as required by (2.i).

Suppose now that $q \in Q_\xi$ and $(q, z) \in \widehat{X}_\tau$ but $q \notin F_\xi$ or $\beta \notin A_\tau$. Since $(q, z) \in \widehat{X}_\tau$ there must exist $(\tau_1, \xi, \tau_1', R_1') \in CH$ and $\tau_1'' \in \mathrm{prf}(\tau_1')$ such that $\tau_1 \preceq \tau = \tau_1\tau_1''$ and $x_{\tau_1} = 1 - z$ and $R_1'(\tau_1'') = q$. But (τ_1', R_1') is a witness for (τ_1, ξ) and hence $R_1'(\tau_1') \in F_\xi$ and $M, \tau_1\tau_1' \models \beta$. Consequently $\tau_1'' \prec \tau_1'$ and thus $\tau_1'' a \in \mathrm{prf}(\tau_1')$ for the unique a. This implies that $R_1'(\tau_1'') \xrightarrow{a}_\xi R_1'(\tau_1'' a)$. Let $R_1'(\tau_1'' a) = q'$. Then $q \xrightarrow{a}_\xi q'$. But then it follows directly from the definition of $X_{\tau a}$, that $(q', 1 - z) \in \widehat{X}_{\tau a}$ as required by (2.ii).

Next suppose that $f_\tau = \checkmark$. Then clearly $(x_{\tau a}, f_{\tau a}) = (1 - x_\tau, \downarrow)$ by the definition of ρ. So assume that $f_\tau = \downarrow$. Supposing there exists $\alpha\, \mathcal{U}^\pi \beta = \xi$ in Req and there exists $q \in Q_\xi$ such that $(q, z) \in \widehat{X}_\tau$ where $z = x_\tau$. Further suppose $q \notin F_\xi$ or $\beta \notin A_\tau$. Now $(q, z) \in \widehat{X}_\tau$ implies that there exists $(\tau_1, \xi, \tau_1', R1') \in CH$ such that $\tau_1 \preceq \tau = \tau_1\tau_1''$ for some $\tau_1'' \in \mathrm{prf}(\tau_1')$ with the further property that $x_{\tau_1} = 1 - z$. From the definitions and the fact that $q \notin F_\xi$ or $\beta \notin A_\tau$ it follows that $\tau_1 \preceq \tau \prec \tau_1\tau_1'$. Hence by the definition of ρ it follows that $(x_{\tau a}, f_{\tau a}) = (x_\tau, \downarrow)$ as required by (2.iii). On the other hand, if such a $(q, z) \in \widehat{X}_\tau$ does not exist, then it follows directly from the definition that $(x_{\tau a}, f_{\tau a}) = (x_\tau, \checkmark)$ as required by (2.iii).

We have now verified that ρ is a run of $\mathcal{B}$ over σ. To show that ρ is accepting it suffices to prove that for any $\tau \in \mathrm{prf}(\sigma)$ there exists τ' such that $\tau\tau' \in \mathrm{prf}(\sigma)$ and $f_{\tau\tau'} = \checkmark$.

Case 1: $\underline{(x_\tau, f_\tau) = (0, \checkmark)}$. By picking $\tau' = \varepsilon$ the desired conclusion follows trivially.

Case 2: $\underline{(x_\tau, f_\tau) = (0, \downarrow)}$. Define the set $\Gamma_\tau \subseteq CH$ as follows. Suppose $(\tau, \xi, \tau', R') \in CH$. Then it must be the case that $(\tau_1, \xi_1, \tau_1', R_1') \in \Gamma_\tau$ iff $\tau_1 \preceq \tau \prec \tau_1\tau_1'$ and $x_{\tau_1} = 1$. If $\Gamma_\tau = \emptyset$ then it is easy to see that with $\tau' = a$ where $\tau a \in \mathrm{prf}(\sigma)$ we must have $f_{\tau\tau'} = \checkmark$ as required.

So suppose $\Gamma_\tau \neq \emptyset$. Define, for each $ch = (\tau_1, \xi_1, \tau_1', R_1') \in \Gamma_\tau$, $k_{ch} = |\tau_1\tau_1'| - |\tau|$ and set $k_\tau = \max(\{k_{ch}\}_{ch \in \Gamma_\tau})$. Let $\tau a \in \mathrm{prf}(\sigma)$. Then it is easy to see that $(x_{\tau a}, f_{\tau a}) = (0, \downarrow)$. But it is also easy to verify $\Gamma_{\tau a} = \emptyset$ or $k_{\tau a} < k_\tau$. Proceeding in this way the required conclusion can be drawn eventually.

The two other cases can be resolved by similar arguments. □

It is now straightforward to establish the main result of this section. To start with we define the size of a formula α, denoted $|\alpha|$, via:

- $|p| = 1$, $|\sim \alpha| = 1 + |\alpha|$ and $|\alpha \vee \beta| = 1 + |\alpha| + |\beta|$.
- $|\alpha\, \mathcal{U}^\pi \beta| = 1 + |\alpha| + |\pi| + |\beta|$,

where $|\pi|$ is given by $|a| = 1$, $|\pi + \pi'| = |\pi; \pi'| = 1 + |\pi| + |\pi'|$ and $|\pi^*| = 1 + |\pi|$.

Theorem 4.3 *For each $\alpha \in \mathrm{DLTL}(\Sigma)$ the question whether or not α is satisfiable can be decided in time $2^{O(|\alpha|)}$.*

Proof: Let $\alpha_0 \in \mathrm{DLTL}(\Sigma)$. Then α_0 is satisfiable iff $\mathcal{L}(\mathcal{B}_{\alpha_0}) \neq \emptyset$ where α_0 is the Büchi automaton constructed above. The emptiness problem for $\mathcal{B}_{\alpha_0}$ can be settled in $O(|S|)$ where S is the set of states of $\mathcal{B}$ [20].

Clearly $CL(\alpha_0)$ is linear in the size of α_0 and hence $|AT| = 2^{O(|\alpha_0|)}$. Let $\alpha\, \mathcal{U}^\pi \beta \in \mathit{Req}$. It is known that for $\pi \in \mathrm{Prg}(\Sigma)$, we can construct in polynomial time a non-deterministic finite state automaton $\mathcal{A}_\xi$ with $\mathcal{L}(\mathcal{A}) = ||\pi||$ such that $|Q_\xi|$ is linear in the size of π (see [7] for a recent account on converting regular expression to small finite state automata).

Let $Req = \{\alpha_1 \,\mathcal{U}^{\pi_1}\beta_1, \ldots, \alpha_m \,\mathcal{U}^{\pi_m}\beta_m\}$. Then $|\pi_1|+|\pi_2|+\ldots+|\pi_m| \leq |\alpha_0|$. Consequently both Q and $\widehat{Q}$ are linear in the size of α_0. It is now easy to see that $|S| = 2^{O(|\alpha_0|)}$. $\square$

As usual, the decision procedure can be applied to solve the associated model checking problem but we will not enter into details here.

5 Some Expressiveness Results

Our main goal here is to show that DLTL(Σ) has the same expressive power as the monadic second-order theory of infinite sequences over Σ. Towards the end of the section we will also establish that a natural sublogic of DLTL(Σ) captures the first order theory of infinite sequences over Σ.

In order to obtain clean formulations of the expressiveness results, we shall banish atomic propositions through the rest of the paper. Instead, we will just work with the constant $\top$ and its negation $\sim\top \stackrel{\Delta}{\Longleftrightarrow} \bot$. To be precise, the syntax of DLTL(Σ) will be from now on assumed to be:

$$\mathrm{DLTL}(\Sigma) ::= \top \mid \sim\alpha \mid \alpha \vee \beta \mid \alpha \,\mathcal{U}^{\pi}\beta,$$

where $\pi \in \mathrm{Prg}(\Sigma)$ with $\mathrm{Prg}(\Sigma)$ defined as before.

A model is now just a ω-sequence $\sigma \in \Sigma^\omega$. For $\tau \in \mathrm{prf}(\sigma)$ we define $\sigma, \tau \models \alpha$ via:

- $\sigma, \tau \models \top$.
- All the other clauses are filled in exactly as in Section 3 while replacing M by σ in the appropriate places.

Each formula α now defines a ω-language $L_\alpha \subseteq \Sigma^\omega$ given by:

$$L_\alpha = \{\sigma \mid \sigma, \varepsilon \models \alpha\}.$$

We say that $L \subseteq \Sigma^\omega$ is DLTL(Σ)-definable iff there exists some $\alpha \in$ DLTL(Σ) such that $L = L_\alpha$.

The monadic second-order theory of infinite sequences over Σ is denoted S1S(Σ). Its vocabulary consists of a family of unary predicates $\{R_a\}_{a\in\Sigma}$, one for each $a \in \Sigma$; a binary predicate $\leq$; a binary predicate ε; a countable supply of individual variables $Var = \{x, y, z, \ldots\}$; a countable supply of set variables (i.e. monadic predicate variables) $SVar = \{X, Y, Z, \ldots\}$. The formulas of S1S($\Sigma$) are then built up by:

- $R_a(x)$, $x \leq y$ and $x \in X$ are atomic formulas.
- If ϕ and ϕ' are formulas then so are $\sim\phi$, $\phi \vee \phi'$, $(\exists x)\phi$ and $(\exists X)\phi$.

A structure for $\mathrm{S1S}(\Sigma)$ is a ω-sequence $\sigma \in \Sigma^\omega$. Let $\mathcal{I}$ be an interpretation of the variables with $\mathcal{I} : Var \longrightarrow \omega$ and $\mathcal{I} : SVar \longrightarrow 2^\omega$. Then the notion of σ being a model of ϕ under the interpretation $\mathcal{I}$, denoted $\sigma \models_\mathcal{I} \phi$, is defined in the expected manner. In particular, $\sigma \models_\mathcal{I} R_a(x)$ iff $\sigma(\mathcal{I}(x)) = a$ (note that $\sigma \in \Sigma^\omega$ is here to be viewed as $\sigma : \omega \longrightarrow \Sigma$); $\sigma \models_\mathcal{I} x \leq y$ iff $\mathcal{I}(x) \leq \mathcal{I}(y)$ (here $\leq$ is the usual ordering over ω); $\sigma \models_\mathcal{I} x \in X$ iff $\mathcal{I}(x) \in \mathcal{I}(X)$.

As usual, a sentence is a formula with no free variables. Each sentence ϕ defines a ω-language denoted L_ϕ where:

$$L_\phi = \{\sigma \mid \sigma \models \phi\}.$$

We say that $L \subseteq \Sigma^\omega$ is $\mathrm{S1S}(\Sigma)$-definable iff there exists a sentence $\phi \in \mathrm{S1S}(\Sigma)$ such that $L = L_\phi$.

Lemma 5.1 *Let $L \subseteq \Sigma^\omega$. If L is* $\mathrm{DLTL}(\Sigma)$*-definable then L is* $\mathrm{S1S}(\Sigma)$*-definable.*

Proof: Consider the construction from the previous section which associates a Büchi automaton $\mathcal{B}_{\alpha_0}$ with each formula $\alpha_0 \in \mathrm{DLTL}(\Sigma)$. Suppose we apply this construction to formulas arising from the restricted syntax assumed in the present section. Then it is easy to see that, in the absence of atomic propositions, $L_{\alpha_0} = \mathcal{L}(\mathcal{B}_{\alpha_0})$. But then the classic result of Büchi [1] asserts that $L \subseteq \Sigma^\omega$ is $\mathrm{S1S}(\Sigma)$-definable iff there exists a Büchi automaton $\mathcal{B}$ operating over Σ such that $L = \mathcal{L}(\mathcal{B})$. $\Box$

Next we wish to show that if $L \subseteq \Sigma^\omega$ is $\mathrm{S1S}(\Sigma)$-definable then L is $\mathrm{DLTL}(\Sigma)$-definable. In fact, it turns out that it suffices to consider a natural fragment of $\mathrm{DLTL}(\Sigma)$ denoted $\mathrm{DLTL}^-(\Sigma)$ whose syntax is given by:

$$\mathrm{DLTL}^-(\Sigma) ::= \top \mid \sim\alpha \mid \alpha \vee \beta \mid \langle\pi\rangle\alpha,$$

where $\pi \in \mathrm{Prg}(\Sigma)$.

Here $\langle\pi\rangle\alpha$ is interpreted as $\top\,\mathcal{U}^\pi\,\alpha$ with the resulting semantics. As before $L \subseteq \Sigma^\omega$ is said to be $\mathrm{DLTL}^-(\Sigma)$-definable iff there exists $\alpha \in \mathrm{DLTL}^-(\Sigma)$ such that $L = L_\alpha$, where L_α is defined as for $\mathrm{DLTL}(\Sigma)$. To get at the result we are after we need to work with Muller automata operating over Σ of the form $\mathcal{A} = (Q, \longrightarrow, Q_{in}, \mathcal{F})$ where:

- Q, $\longrightarrow$ and Q_{in} as in the case of a Büchi automaton.
- $\mathcal{F} \subseteq 2^Q$ is a family of accepting sets of states.

Let $\sigma \in \Sigma^\omega$. Then the notion of a run $\rho : \mathrm{prf}(\sigma) \longrightarrow Q$ of $\mathcal{A}$ over σ is as in the case of a Büchi automaton. The definition of $\inf(\rho)$ is also as before. The run ρ is said to be accepting iff $\inf(\rho) \in \mathcal{F}$. Naturally $\mathcal{L}(\mathcal{A})$, the ω-language accepted by $\mathcal{A}$, is given by : $\sigma \in \mathcal{L}(\mathcal{A})$ iff there exists an accepting run of $\mathcal{A}$ over σ.

The Muller automaton $\mathcal{A} = (Q, \longrightarrow, Q_{in}, \mathcal{F})$ is said to be deterministic iff $|Q_{in}| = 1$ and $\longrightarrow \subseteq Q \times \Sigma \times Q$ is a deterministic transition relation. In other words, whenever $q \xrightarrow{a} q'$ and $q \xrightarrow{a} q''$, we have $q' = q''$. The well-known theorem of McNaughton [12] guarantees that $L \subseteq \Sigma^\omega$ is $\mathrm{S1S}(\Sigma)$-definable iff there exists a deterministic Muller automaton operating over Σ such that $L = \mathcal{L}(\mathcal{A})$. This fact will be the basis for the proof of the next result.

Lemma 5.2 *Let $L \subseteq \Sigma^\omega$. If L is $\mathrm{S1S}(\Sigma)$-definable then L is $\mathrm{DLTL}^-(\Sigma)$-definable.*

Proof: As remarked above L is $\mathrm{S1S}(\Sigma)$-definable implies that there exists a *deterministic* Muller automaton $\mathcal{A} = (Q, \longrightarrow, \{q_{in}\}, \mathcal{F})$ operating over Σ such that $L = \mathcal{L}(\mathcal{A})$. We will exhibit a formula $\alpha_\mathcal{A} \in \mathrm{DLTL}^-(\Sigma)$ such that $L_{\alpha_\mathcal{A}} = \mathcal{L}(\mathcal{A})$.

An easy argument shows that it involves no loss of generality to assume that $\mathcal{A}$ — apart from determinacy — has two additional properties:

(i) $\emptyset \notin \mathcal{F}$.

(ii) $\forall q \in Q \ \forall a \in \Sigma. \quad \exists q'. \ q \xrightarrow{a} q'$.

Determinacy and (ii) ensure that for *every* $\sigma \in \Sigma^\omega$ the Muller automaton $\mathcal{A}$ has a *unique* run over σ. This fact will be crucial in what follows.

Suppose $\mathcal{F} = \emptyset$. In this case we have that $L = \emptyset$, so we set $\alpha_\mathcal{A} = \bot$.

So suppose that $\mathcal{F} \neq \emptyset$. For each $F \in \mathcal{F}$ we shall construct a formula α_F. The required formula $\alpha_\mathcal{A}$ defining L will then be the disjunction of all such α_F.

First we extend $\longrightarrow \subseteq Q \times \Sigma \times Q$ to $\longrightarrow_*$, where $\longrightarrow_*$ is the least subset of $Q \times \Sigma^* \times Q$ satisfying:

- $q \xrightarrow{\varepsilon}_* q$ for every $q \in Q$.

- If $q \xrightarrow{\tau}_* q'$ and $q' \xrightarrow{a} q''$ then $q \xrightarrow{\tau a}_* q''$.

Next define, for each $q, q' \in Q$, the language of finite words $L_{q,q'} \subseteq \Sigma^*$ by:

$$L_{q,q'} = \{\tau \mid q \xrightarrow{\tau}_* q'\}.$$

It is easy to see that each $L_{q,q'}$ is a regular subset of Σ^*. Hence we can fix a regular expression $\pi_{q,q'} \in \mathrm{Prg}(\Sigma)$ such that $L_{q,q'} = ||\pi_{q,q'}||$. Due to the determinacy of $\mathcal{A}$ it follows at once that if $q, q', q'' \in Q$ such that $L_{q,q'} \cap L_{q,q''} \neq \emptyset$ then $q' = q''$.

Now let $F = \{q_0, q_1, \ldots, q_{n-1}\}$ with $n \geq 1$. Then the formula α_F is given by:

$$\alpha_F = \bigvee_{q \in F} \langle \pi_{q_{in},q} \rangle \left(\bigwedge_{q' \notin F} [\pi_{q,q'}]\bot \;\wedge\; \bigwedge_{j=0}^{n-1} [\pi_{q,q_j}]\langle \pi_{q_j,q_{j\oplus 1}} \rangle \top \right),$$

where $\oplus$ denotes addition modulo n. The required formula $\alpha_{\mathcal{A}}$ is then defined as:

$$\alpha_{\mathcal{A}} = \bigvee_{F \in \mathcal{F}} \alpha_F$$

Clearly $\alpha_{\mathcal{A}} \in \mathrm{DLTL}^-(\Sigma)$. We wish to argue that $\mathcal{L}(\mathcal{A}) = L_{\alpha_{\mathcal{A}}}$. So first suppose that $\sigma \in \mathcal{L}(\mathcal{A})$. Let $\rho : \mathrm{prf}(\sigma) \longrightarrow Q$ be the (unique) run of $\mathcal{A}$ over σ. We must have $\inf(\rho) \in \mathcal{F}$ since ρ is an accepting run. Assume that $\inf(\rho) = F = \{q_0, q_1, \ldots, q_{n-1}\}$ with $n \geq 1$. There must exist a $\tau \in \mathrm{prf}(\sigma)$ such that $\rho(\tau) = q \in F$ and for all $\tau' \in \Sigma^*$, $\tau\tau' \in \mathrm{prf}(\sigma)$ implies $\rho(\tau\tau') \in F$. This is so because $\inf(\rho) = F$. If $\rho(\tau) = q$ then $\tau \in ||\pi_{q_{in},q}||$ because $\rho(\varepsilon) = q_{in}$ and (consequently) ρ being a run assures us that $q_{in} \xrightarrow{\tau}_* q$. Now suppose $\sigma, \tau \models \langle \pi_{q,q'} \rangle \top$ for some $q' \notin F$. Then there exists $\tau' \in ||\pi_{q,q'}||$ such that $\tau\tau' \in \mathrm{prf}(\sigma)$. But this violates the assumed property of τ. Consequently, for every $q' \notin F$ we must have $\sigma, \tau \models [\pi_{q,q'}]\bot$. Hence $\sigma, \tau \models \bigwedge_{q' \notin F} [\pi_{q,q'}]\bot$.

Next suppose $j \in \{0, 1, \ldots, n-1\}$ and $\tau' \in ||\pi_{q,q_j}||$ such that $\tau\tau' \in \mathrm{prf}(\sigma)$. Then by the assumed property of τ there must exist $\tau'' \neq \varepsilon$ such that $\tau\tau'\tau'' \in \mathrm{prf}(\sigma)$ and $\rho(\tau\tau'\tau'') = q_{j\oplus 1}$, because $q_{j\oplus 1} \in F = \inf(\rho)$. Clearly $\tau'' \in ||\pi_{q_j,q_{j\oplus 1}}||$ because $\rho(\tau\tau') = q_j$. Thus $\sigma, \tau \models [\pi_{q,q_j}]\langle \pi_{q_j,q_{j\oplus 1}} \rangle \top$ for every $j \in \{0, 1, \ldots, n-1\}$. We have now established that $\sigma, \varepsilon \models \alpha_F$ where $F \in \mathcal{F}$. This in turn implies that $\sigma, \varepsilon \models \alpha_{\mathcal{A}}$ and hence $\sigma \in L_{\alpha_{\mathcal{A}}}$.

Now assume that $\sigma, \varepsilon \models \alpha_{\mathcal{A}}$. Then it must be the case that $\sigma, \varepsilon \models \alpha_F$ for some $F \in \mathcal{F}$ with $F = \{q_0, q_1, \ldots, q_{n-1}\}$, $n \geq 1$. Consequently, there exist $\tau \in \mathrm{prf}(\sigma)$ and $q \in F$ such that $\tau \in ||\pi_{q_{in},q}||$ and

$$\sigma, \tau \models \bigwedge_{q' \notin F} [\pi_{q,q'}]\bot \wedge \bigwedge_{j=0}^{n-1} [\pi_{q,q_j}]\langle \pi_{q_j,q_{j\oplus 1}} \rangle \top.$$

Let ρ be the unique run of $\mathcal{A}$ over σ. From the determinacy of $\mathcal{A}$ it follows that $\rho(\tau) = q$.

Claim: Suppose $\tau' \in \Sigma^*$ is such that $\tau\tau' \in \mathrm{prf}(\sigma)$. Then $\rho(\tau\tau') \in F$.

To see that the claim holds, let $\tau\tau' \in \mathrm{prf}(\sigma)$ and $\rho(\tau\tau') = q'$. If $q' \notin F$ then $\tau' \in ||\pi_{q,q'}||$ would lead to the contradiction that $\sigma, \tau \models \langle \pi_{q,q'} \rangle \top$ for some $q' \notin F$. Thus the claim holds, which at once implies that $\inf(\rho) \subseteq F$.

If we now show that $F \subseteq \inf(\rho)$ we are done. So let $j \in \{0, 1, \ldots, n-1\}$. To show that $q_j \in \inf(\rho)$ assume that $\tau' \in \Sigma^*$ such that $\tau\tau' \in \mathrm{prf}(\sigma)$. Let $\rho(\tau\tau') = q_i$, say. We must show that there exists τ'' such that $\tau\tau'\tau'' \in \mathrm{prf}(\sigma)$ and $\rho(\tau\tau'\tau'') = q_j$.

If $i = j$ then we can set $\tau'' = \varepsilon$ and be done. So suppose $i \neq j$. Then clearly $\tau' \in ||\pi_{q,q_i}||$. But $\sigma, \tau \models [\pi_{q,q_i}]\langle \pi_{q_i,q_{i\oplus 1}} \rangle \top$ and hence there must exist $\tau_i \in ||\pi_{q_i,q_{i\oplus 1}}||$ such that $\tau\tau'\tau_i \in \mathrm{prf}(\sigma)$ so that $\rho(\tau\tau'\tau_i) = q_{i\oplus 1}$. If $i \oplus 1 = j$ then we can set $\tau'' = \tau_i$. Otherwise $\tau'\tau_i \in ||\pi_{q,q_{i\oplus 1}}||$ and hence we can find $\tau_{i\oplus 1}$ such that $\tau\tau'\tau_i\tau_{i\oplus 1} \in \mathrm{prf}(\sigma)$ such that $\rho(\tau\tau'\tau_i\tau_{i\oplus 1}) = q_{i\oplus 2}$. If $i \oplus 2 = j$ we set $\tau'' = \tau_i\tau_{i\oplus 1}$, otherwise we proceed as before. Eventually we will be done. □

Theorem 5.3 *Let $L \subseteq \Sigma^\omega$. Then the following statements are equivalent:*

(i) *L is* S1S(Σ)*-definable.*

(ii) *L is* DLTL(Σ)*-definable.*

(iii) *L is* DLTL$^-$(Σ)*-definable.*

Proof: Follows immediately from Lemmas 5.1, 5.2 and the fact that DLTL$^-$(Σ) is a sublogic of DLTL(Σ). □

At present we do not know of a direct conversion of DLTL(Σ)-formulas into DLTL$^-$(Σ)-formulas. Although these two logics have the same expressive power in the sense of Theorem 5.3, it appears that DLTL(Σ)

will admit more natural specifications. In addition, it is a conservative extension of LTL(Σ) even from a syntactic standpoint and hence conventional LTL specifications can be brought in with no overhead translation costs.

We shall conclude this section by pointing out that star-free programs can be used to capture the first-order definable subsets of Σ^ω. Admittedly this is not a big surprise, but it illustrates once more that our method of augmenting the expressive power of LTL is a natural one.

FO(Σ) will denote the first-order theory of ω-sequences generated by Σ. It is the fragment of S1S(Σ) obtained by eliminating set variables from the syntax. We shall say that $L \subseteq \Sigma^\omega$ is FO(Σ)-definable iff there exists a sentence ϕ in FO(Σ) such that $L = L_\phi$.

The set of *star-free regular programs* over Σ is denoted $\mathrm{Prg}_{\mathbb{SF}}(\Sigma)$ and its syntax is given by:

$$\mathrm{Prg}_{\mathbb{SF}}(\Sigma) ::= 0 \mid a \mid \pi + \pi' \mid \overline{\pi} \mid \pi; \pi'.$$

The set of finite words denoted by each star-free program is obtained via the map $||\cdot|| : \mathrm{Prg}_{\mathbb{SF}}(\Sigma) \longrightarrow 2^{\Sigma^*}$ which is defined as follows: $||\overline{\pi}|| = \Sigma^* - ||\pi||$ and $||0|| = \emptyset$. The remaining cases are handled as before.

Next let $\varphi \in$ FO(Σ), $\tau \in \Sigma^*$ and $\mathcal{I} : Var \longrightarrow \{0, 1, \ldots, |\tau| - 1\}$ where $|\tau|$ is the length of τ. In case $\tau = \varepsilon$ we set $\mathcal{I}$ — viewed as a set of ordered pairs — to the empty set. The notion $\tau \models_{\mathcal{I}} \varphi$ is defined in the obvious way.

With each sentence $\varphi \in$ FO(Σ) we can now associate the language of finite words L^{fin}_φ given by:

$$L^{fin}_\varphi = \{\tau \mid \tau \in \Sigma^* \text{ and } \tau \models \varphi\}.$$

We let $SF^{fin}(\Sigma)$ be the subset of 2^{Σ^*} where $L \subseteq \Sigma^*$ is admitted as a member of $SF^{fin}(\Sigma)$ iff there exists a sentence $\varphi \in$ FO(Σ) such that $L = L^{fin}_\varphi$. A well-known fact [13] is that $SF^{fin}(\Sigma)$ and $\mathrm{Prg}_{\mathbb{SF}}(\Sigma)$ denote each other in the following sense.

Proposition 5.4 *Let $L \subseteq \Sigma^*$. Then $L \in SF^{fin}(\Sigma)$ iff there exists $\pi \in \mathrm{Prg}_{\mathbb{SF}}(\Sigma)$ such that $L = ||\pi||$.*

Next we recall a well-known characterization of FO(Σ)-definable subsets of Σ^ω. Let $SF^\omega(\Sigma)$ be the least subset of 2^{Σ^ω} which satisfies:

- $\emptyset \in SF^\omega(\Sigma)$.

- If $\widehat{L} \in SF^{\omega}(\Sigma)$ then $\Sigma^{\omega} - \widehat{L} \in SF^{\omega}(\Sigma)$.
- If $\widehat{L}_1, \widehat{L}_2 \in SF^{\omega}(\Sigma)$ then $\widehat{L}_1 \cup \widehat{L}_2 \in SF^{\omega}(\Sigma)$.
- If $L \in SF^{fin}(\Sigma)$ and $\widehat{L} \in SF^{\omega}(\Sigma)$ then $\{\tau\sigma \mid \tau \in L \text{ and } \sigma \in \widehat{L}\} \in SF^{\omega}(\Sigma)$.

We then have the following characterization (see [20]).

Proposition 5.5 *Let $L \subseteq \Sigma^{\omega}$. Then $L \in SF^{\omega}(\Sigma)$ iff L is* FO(Σ)-*definable.*

The star-free version of DLTL(Σ) will be denoted — for want of a better notation — by DLTL$_{\mathbb{SF}}$(Σ) and its syntax is given by:

$$\mathrm{DLTL}_{\mathbb{SF}}(\Sigma) ::= \top \mid \sim\alpha \mid \alpha \vee \beta \mid \alpha\, \mathcal{U}^{\pi} \beta \quad (\pi \in \mathrm{Prg}_{\mathbb{SF}}(\Sigma)).$$

Thus the only difference is that the programs that are used to build up the until-formulas are required to be star-free programs. All the semantic notions concerning this logic are transported from DLTL(Σ) in the natural manner. The fragment of DLTL$_{\mathbb{SF}}$(Σ) which corresponds to DLTL^{-}(Σ) has the syntax:

$$\mathrm{DLTL}^{-}_{\mathbb{SF}}(\Sigma) ::= \top \mid \sim\alpha \mid \alpha \vee \beta \mid \langle\pi\rangle\alpha \quad (\pi \in \mathrm{Prg}_{\mathbb{SF}}(\Sigma)).$$

Once again, all the semantic notions concerning DLTL$^{-}_{\mathbb{SF}}$(Σ) are transported from DLTL^{-}(Σ) in the natural way.

We now wish to show that FO(Σ), DLTL$_{\mathbb{SF}}$(Σ) and DLTL$^{-}_{\mathbb{SF}}$(Σ) all have the same expressive power.

Lemma 5.6 *Let $\widehat{L} \subseteq \Sigma^{\omega}$. If $\widehat{L}$ is* FO(Σ)-*definable then $\widehat{L}$ is* DLTL$^{-}_{\mathbb{SF}}$(Σ)-*definable.*

Proof: Using Proposition 5.5 and following the definition of $SF^{\omega}(\Sigma)$, we now inductively define the map $g : SF^{\omega}(\Sigma) \longrightarrow \mathrm{DLTL}^{-}_{\mathbb{SF}}(\Sigma)$ as follows.

- If $\widehat{L} = \emptyset$ then $g(\widehat{L}) = \perp$.
- Suppose $\widehat{L} = \Sigma^{\omega} - \widehat{L}_1$ for some $\widehat{L}_1 \in SF^{\omega}(\Sigma)$ with $g(\widehat{L}_1)$ defined. Then $g(\widehat{L}) = \sim g(\widehat{L}_1)$.
- Suppose $\widehat{L} = \widehat{L}_1 \cup \widehat{L}_2$ with $\widehat{L}_1, \widehat{L}_2 \in SF^{\omega}(\Sigma)$ so that both $g(\widehat{L}_1)$ and $g(\widehat{L}_2)$ are defined. Then $g(\widehat{L}) = g(\widehat{L}_1) \vee g(\widehat{L}_2)$.

- Suppose $L_1 \in SF^{fin}(\Sigma)$ and $\widehat{L}_1 \in SF^{\omega}(\Sigma)$ such that $L = \{\tau\sigma \mid \tau \in L_1 \text{ and } \sigma \in \widehat{L}_1\}$ and $g(\widehat{L}_1)$ is defined. Let $\pi \in \mathrm{Prg}_{\mathbb{SF}}(\Sigma)$ such that $L = ||\pi||$. The existence of π is guaranteed by Proposition 5.4. Then $g(\widehat{L}) = \langle\pi\rangle g(\widehat{L}_1)$.

It is now routine to verify that for each $\widehat{L} \in SF^{\omega}(\Sigma)$

$$\widehat{L} = \{\sigma \mid \sigma, \varepsilon \models g(\widehat{L})\}.$$

Hence $\widehat{L}$ is $\mathrm{DLTL}^{-}_{\mathbb{SF}}(\Sigma)$-definable. $\square$

The last step is to show that $\mathrm{DLTL}_{\mathbb{SF}}(\Sigma)$-definability implies $\mathrm{FO}(\Sigma)$-definability. The proof will be syntactic and it will be convenient to go through some preliminaries.

First let $\varphi \in \mathrm{FO}(\Sigma)$ and x, y be variables that do not appear in φ (free or bound). Then the formulas $\mathrm{Rel}(\varphi, 0, x)$ and $\mathrm{Rel}(\varphi, x, y)$ — describing the *relativizations* of φ — are defined inductively as follows:

- $\mathrm{Rel}(R_a(z), 0, x) = z \leq x \wedge R_a(z)$.
  $\mathrm{Rel}(R_a(z), x, y) = x \leq z \wedge z \leq y \wedge R_a(z)$.
- $\mathrm{Rel}(z_1 \leq z_2, 0, x) = z_2 \leq x \wedge z_1 \leq z_2$.
  $\mathrm{Rel}(z_1 \leq z_2, x, y) = x \leq z_1 \wedge z_2 \leq y \wedge z_1 \leq z_2$.
- $\mathrm{Rel}(\sim\varphi, 0, x) = \sim\mathrm{Rel}(\varphi, 0, x)$.
  $\mathrm{Rel}(\sim\varphi, x, y) = \sim\mathrm{Rel}(\varphi, x, y)$.
- $\mathrm{Rel}(\varphi_1 \vee \varphi_2, 0, x) = \mathrm{Rel}(\varphi_1, 0, x) \vee \mathrm{Rel}(\varphi_2, 0, x)$.
  $\mathrm{Rel}(\varphi_1 \vee \varphi_2, x, y) = \mathrm{Rel}(\varphi_1, x, y) \vee \mathrm{Rel}(\varphi_2, x, y)$.
- $\mathrm{Rel}((\exists z)\varphi, 0, x) = (\exists z)(z \leq x \wedge \mathrm{Rel}(\varphi, 0, x))$.
  $\mathrm{Rel}((\exists z)\varphi, x, y) = (\exists z)(x \leq z \wedge z \leq y \wedge \mathrm{Rel}(\varphi, x, y))$.

It is easy to see that both $\mathrm{Rel}(\varphi, 0, x)$ and $\mathrm{Rel}(\varphi, x, y)$ are members of $\mathrm{FO}(\Sigma)$ for any $\varphi \in \mathrm{FO}(\Sigma)$.

Next we fix for each $\pi \in \mathrm{Prg}_{\mathbb{SF}}(\Sigma)$ a sentence $\varphi_\pi \in \mathrm{FO}(\Sigma)$ such that $||\pi|| = L^{fin}_{\varphi}$. The existence of φ_π is guaranteed by Proposition 5.4 and the definition of $SF^{fin}(\Sigma)$.

Finally we inductively define, for each $\alpha \in \mathrm{DLTL}_{\mathbb{SF}}(\Sigma)$, the formulas $\mathrm{SAT}(\alpha, 0)$ and $\mathrm{SAT}(\alpha, x)$ as follows:

- $\mathrm{SAT}(\top, 0) = \mathrm{SAT}(\top, x) = (\forall z)(z \leq z)$.

- $\mathrm{SAT}(\sim\alpha, 0) = \sim\mathrm{SAT}(\alpha, 0)$.
  $\mathrm{SAT}(\sim\alpha, x) = \sim\mathrm{SAT}(\alpha, x)$.

- $\mathrm{SAT}(\alpha \vee \beta, 0) = \mathrm{SAT}(\alpha, 0) \vee \mathrm{SAT}(\beta, 0)$.
  $\mathrm{SAT}(\alpha \vee \beta, x) = \mathrm{SAT}(\alpha, x) \vee \mathrm{SAT}(\beta, x)$.

- Let $\alpha = \beta_1 \, \mathcal{U}^\pi \beta_2$. Suppose first that $\varepsilon \in ||\pi||$. Then

  $$\mathrm{SAT}(\alpha, 0) = \mathrm{SAT}(\beta_2, 0) \vee (\mathrm{SAT}(\beta_1, 0) \wedge (\exists x)(\varphi_1 \wedge \varphi_2 \wedge \varphi_3)),$$

  where

  $$\begin{aligned} \varphi_1 &= \mathrm{SAT}(\beta_2, x), \\ \varphi_2 &= (\forall y)(y < x \supset \mathrm{SAT}(\beta_1, y)), \\ \varphi_3 &= \mathrm{Rel}(\varphi_\pi, 0, x). \end{aligned}$$

  Moreover,

  $$\mathrm{SAT}(\alpha, x) = \mathrm{SAT}(\beta_2, x) \vee (\mathrm{SAT}(\beta_1, x) \wedge (\exists y)(x < y \wedge \varphi_1' \wedge \varphi_2' \wedge \varphi_3')),$$

  where

  $$\begin{aligned} \varphi_1' &= \mathrm{SAT}(\beta_2, y), \\ \varphi_2' &= (\forall z)(x \leq z \wedge z < y \supset \mathrm{SAT}(\beta_1, z)), \\ \varphi_3' &= \mathrm{Rel}(\varphi_\pi, x, y). \end{aligned}$$

  Now, if it is the case that $\varepsilon \notin ||\pi||$ we disregard $\mathrm{SAT}(\beta_2, 0)$ and $\mathrm{SAT}(\beta_2, x)$, respectively, above. In other words,

  $$\mathrm{SAT}(\alpha, 0) = \mathrm{SAT}(\beta_1, 0) \wedge (\exists x)(\varphi_1 \wedge \varphi_2 \wedge \varphi_3),$$

  and

  $$\mathrm{SAT}(\alpha, x) = \mathrm{SAT}(\beta_1, x) \wedge (\exists y)(x < y \wedge \varphi_1' \wedge \varphi_2' \wedge \varphi_3'),$$

We can now prove the following result.

Lemma 5.7 *Let $L \subseteq \Sigma^\omega$. If L is $\mathrm{DLTL}_{\mathbb{SF}}(\Sigma)$-definable then L is $\mathrm{FO}(\Sigma)$-definable.*

Proof: Let $L \subseteq \Sigma^\omega$ be given such that $L = L_\alpha$ with $\alpha \in \mathrm{DLTL}_{\mathbb{SF}}(\Sigma)$, i.e. $L = \{\sigma \mid \sigma, \varepsilon \models \alpha\}$. It is now routine to check that $\mathrm{SAT}(\alpha, 0)$ is a sentence of $\mathrm{FO}(\Sigma)$ and moreover that

$$L = \{\sigma \mid \sigma \models \mathrm{SAT}(\alpha, 0)\}.$$

Hence, L is $\mathrm{FO}(\Sigma)$-definable. $\square$

Theorem 5.8 *Let $L \subseteq \Sigma^\omega$. Then the following statements are equivalent:*

(i) *L is $\mathrm{FO}(\Sigma)$-definable.*

(ii) *L is $\mathrm{DLTL}_{\mathbb{SF}}(\Sigma)$-definable.*

(iii) *L is $\mathrm{DLTL}^-_{\mathbb{SF}}(\Sigma)$-definable.*

Proof: Follows easily from Lemmas 5.6 and 5.7. $\square$

6 Axiomatizations

Our axiomatization of the set of valid formulas of DLTL is an extension of Segerberg's axiomatization of PDL [17]. Moreover, our completeness argument is based on the elegant proof of completeness of Segerberg's axioms due to Kozen and Parikh [9]. It will be convenient to first axiomatize DLTL^-.

We begin by augmenting the set of regular programs with the atomic program 1. We set $||1|| = \{\varepsilon\}$. By abuse of notation this augmented set of programs will also be denoted as $\mathrm{Prg}(\Sigma)$. Next we define the transition relation $\longrightarrow_{\mathrm{Prg}(\Sigma)}$ (from now on written as just $\longrightarrow$) to be the least subset of $\mathrm{Prg}(\Sigma) \times \Sigma \times \mathrm{Prg}(\Sigma)$ yielded by the following rules:

- $\dfrac{}{a \xrightarrow{a} 1}$

- $\dfrac{\pi \xrightarrow{a} \pi_1}{\pi + \pi' \xrightarrow{a} \pi_1} \qquad \dfrac{\pi \xrightarrow{a} \pi_1}{\pi' + \pi \xrightarrow{a} \pi_1}$

- $\dfrac{\pi \xrightarrow{a} \pi_1}{\pi ; \pi' \xrightarrow{a} \pi_1 ; \pi'}$ if $\pi_1 \neq 1$

- $$\frac{\pi \xrightarrow{a} 1}{\pi; \pi' \xrightarrow{a} \pi'}$$

- $$\frac{\pi' \xrightarrow{a} \pi''}{\pi; \pi' \xrightarrow{a} \pi''} \text{ if } \varepsilon \in ||\pi||$$

- $$\frac{\pi \xrightarrow{a} \pi'}{\pi^* \xrightarrow{a} \pi'; \pi^*}.$$

This transition relation is extended to the relation $\longrightarrow_* \subseteq \mathrm{Prg}(\Sigma) \times \Sigma^* \times \mathrm{Prg}(\Sigma)$ via:

- $\pi \xrightarrow{\varepsilon}_* \pi$
- If $\pi \xrightarrow{\tau}_* \pi'$ and $\pi' \xrightarrow{a} \pi''$ then $\pi \xrightarrow{\tau a}_* \pi''$.

Finally the sets of programs $\delta_a(\pi)$ and $\delta_*(\pi)$ for each π and each a are defined as follows:

- $\delta_a(\pi) = \{\pi' \mid \pi \xrightarrow{a} \pi'\}$.
- $\delta_*(\pi) = \{\pi' \mid \exists\tau.\ \pi \xrightarrow{\tau}_* \pi'\}$.

Proposition 6.1 *For each π and each a, both $\delta_a(\pi)$ and $\delta_*(\pi)$ are finite sets.*

Proof: The proof follows easily by structural induction on π. □

We are now ready to present an axiomatization of DLTL$^-$ (Recall that $O\alpha \stackrel{\Delta}{\Longleftrightarrow} \bigvee_{a\in\Sigma}\langle a\rangle\alpha$). The logical system $\mathcal{DLTL}^-$ is given as follows.

Axiom schemes:

(A0) All the tautologies of propositional calculus.

(A1) $[\pi]\ (\alpha \supset \beta) \supset ([\pi]\alpha \supset [\pi]\beta)$.

(A2) $\langle \pi + \pi'\rangle\alpha \equiv \langle\pi\rangle\alpha \vee \langle\pi'\rangle\alpha$.

(A3) $\langle \pi; \pi'\rangle\alpha \equiv \langle\pi\rangle\langle\pi'\rangle\alpha$.

(A4) $\langle \pi^*\rangle\alpha \equiv \alpha \vee \langle\pi\rangle\langle\pi^*\rangle\alpha$.

(A5) $[\pi^*](\alpha \supset [\pi]\alpha) \supset (\alpha \supset [\pi^*]\alpha)$.

(A6) $\alpha \equiv \langle 1\rangle\alpha$.

(A7) $O\top$.

(A8) $\langle a\rangle\top \supset \bigwedge_{b\neq a}[b]\bot$.

(A9) $\langle a\rangle\alpha \supset [a]\alpha$.

(A10) $\langle\pi\rangle\alpha \equiv \alpha \vee \left(\bigvee_{a\in\Sigma}\langle a\rangle \bigvee_{\pi'\in\delta_a(\pi)}\langle\pi'\rangle\alpha\right)$, $(\varepsilon \in ||\pi||)$.

(A11) $\langle\pi\rangle\alpha \equiv \bigvee_{a\in\Sigma}\langle a\rangle \bigvee_{\pi'\in\delta_a(\pi)}\langle\pi'\rangle\alpha$, $(\varepsilon \notin ||\pi||)$.

Inference rules:

(MP) $\dfrac{\alpha \quad \alpha \supset \beta}{\beta}$.

(TG) $\dfrac{\alpha}{[\pi]\alpha}$.

(A0) through (A5) and the inference rules together constitute an axiomatization of PDL. The behaviour of 1 is captured by (A6). The remaining axiom schemes describe the linear time semantics provided for regular programs in the setting of DLTL$^-$. Due to Proposition 6.1 both (A10) and (A11) are well-defined. It is easy to see that the axioms are valid and that the inference rules preserve validity.

We shall say, as usual, that a formula α is ($\mathcal{DLTL}^-$ $-$) consistent in case $\sim\alpha$ is not a thesis derivable from the system $\mathcal{DLTL}^-$. We shall prove that every consistent formula is satisfiable. To this end, fix a consistent formula α_0. Define $\widehat{cl}(\alpha_0)$ to be the least set of formulas containing $cl(\alpha_0)$ (recall the definition of $cl(\alpha_0)$ in Section 4) and satisfying:

- If $\langle\pi\rangle\alpha \in \widehat{cl}(\alpha_0)$ and $\pi' \in \delta_a(\pi)$ then$\langle\pi'\rangle\alpha, \langle a\rangle\langle\pi'\rangle\alpha \in \widehat{cl}(\alpha_0)$.
- If $\langle 1\rangle\alpha \in \widehat{cl}(\alpha_0)$ then $\alpha \in \widehat{cl}(\alpha_0)$.
- $\langle a\rangle\top \in \widehat{cl}(\alpha_0)$ for every $a \in \Sigma$.

Next define $\widehat{CL}(\alpha_0)$ as $\widehat{CL}(\alpha_0) = \widehat{cl}(\alpha_0) \cup \{\sim\beta \mid \beta \in \widehat{cl}(\alpha_0)\}$. As usual, we will identify $\sim\sim\beta$ with β in what follows.

Proposition 6.2 *$\widehat{CL}(\alpha_0)$ is a finite set.*

Proof: Follows at once from Proposition 6.1. □

In this section, an atom is a maximal consistent subset of $\widehat{CL}(\alpha_0)$. If A is an atom then $\widehat{A}$ will be the conjunction of all the formulas in A. Let AT_0 be the set of all atoms. We now define the transition system $TS_0 = (AT_0, \Longrightarrow)$ where $\Longrightarrow \subseteq AT_0 \times \Sigma \times AT_0$ is given by $A \overset{a}{\Longrightarrow} B$ iff $\widehat{A} \wedge \langle a\rangle\widehat{B}$ is consistent. As before, the transition relation $\Longrightarrow$ is extended to $\Longrightarrow_* \subseteq AT_0 \times \Sigma^* \times AT_0$ in the obvious way.

Lemma 6.3

(i) *Suppose $A, B \in AT_0$ and $\pi \in \mathrm{Prg}(\Sigma)$ such that $\widehat{A} \wedge \langle\pi\rangle\widehat{B}$ is consistent. Then there exists $\tau \in ||\pi||$ such that $A \overset{\tau}{\Longrightarrow}_* B$.*

(ii) *Suppose* $\langle\pi\rangle\alpha \in A \in AT_0$. *Then there exists* $B \in AT_0$ *and* $\tau \in ||\pi||$ *such that* $\alpha \in B$ *and* $A \overset{\tau}{\Longrightarrow}_* B$.

Proof: Part (i) can be established by just repeating the proof of [9, Lemma 1]. Now part (ii) follows easily from part (i) with the help of a few tautologies of propositional calculus. □

We are now ready to extract a model of α_0 from TS_0. We shall do so by inductively defining a map $\widehat{\rho} : \omega \longrightarrow AT_0$ and an ascending chain of sequences $\tau_0 \prec \tau_1 \prec \ldots$ where each τ_i is in Σ^*. In what follows we will denote $\widehat{\rho}(i)$ by A_i for each $i \in \omega$. We shall also assume that we have fixed an enumeration of the countable set $\widehat{CL}(\alpha_0) \times \Sigma^*$.

- $\widehat{\rho}(0) = A_0$ where $A_0 \in AT_0$ such that $\alpha_0 \in A_0$. Further, $\tau_0 = \varepsilon$.
- Assume $\widehat{\rho}(i)$ and τ_i are defined. We say that the pair $(\langle\pi\rangle\alpha, \tau)$ is a requirement at stage i provided the following conditions are satisfied:
  - $\tau \preceq \tau_i$ and $\langle\pi\rangle\alpha \in A_j$ where $|\tau| = j$.
  - For every $\tau' \in \Sigma^*$, if $\tau\tau' \preceq \tau_i$ then $\tau' \notin ||\pi||$ or $\alpha \notin A_k$ where $|\tau\tau'| = k$.

Let RQ_i be the set of requirements at stage i. Assume that $RQ_i = \emptyset$. Let $a \in \Sigma$ such that $\langle a\rangle\top \in A_i$. The fact that such an a exists and is unique is guaranteed by (A7) and (A8). Since $\bigvee_{A \in AT_0} \widehat{A}$ is a thesis, it follows from simple propositional reasoning that $\widehat{A} \wedge \langle a\rangle\widehat{B}$ is consistent for some $B \in AT_0$. Consequently $A \overset{a}{\Longrightarrow} B$. Now let $\widehat{\rho}(i+1) = B$ and $\tau_{i+1} = \tau_i a$. The construction now proceeds from stage $i+1$.

Assume now that $RQ_i \neq \emptyset$. Let $(\langle\pi\rangle\alpha, \tau)$ be the least member of RQ_i in the enumeration we have fixed for $\widehat{CL}(\alpha_0) \times \Sigma^*$. Let $j = |\tau|$ and $\tau\tau' = \tau_i$. Then using (A10) and (A11) it is easy to show that there exists π' such that $\pi \overset{\tau'}{\longrightarrow}_* \pi'$ and $\langle\pi'\rangle\alpha \in A_i$. Moreover $\alpha \notin A_i$ or $\varepsilon \notin ||\pi'||$. By part (ii) of Lemma 6.3, there exists $B \in AT_0$ and $\tau'' \in ||\pi'||$ such that $A_i \overset{\tau''}{\Longrightarrow}_* B$ and $\alpha \in B$. Let $\tau'' = b_1 b_2 \ldots b_m$. Then we can find $B_0, B_1, \cdots, B_m \in AT_0$ such that $A_i = B_0$ and $B_m = B$ and $B_k \overset{b_k}{\Longrightarrow} B_{k+1}$ for $0 \leq k < m$. We now extend $\widehat{\rho}$ by:

$$\widehat{\rho}(i+k) = B_k \quad \text{for} \quad 1 \leq k \leq m.$$

Further we define $\tau_{i+k} = \tau_i b_1 b_2 \ldots b_k$ for $1 \leq k \leq m$. The construction now proceeds from stage $i+m$.

Now consider the model $M_0 = (\sigma, V_0)$ where $\sigma \in \Sigma^\omega$ is the sequence satisfying that $\tau_i \preceq \sigma$ for every $i \in \omega$. Further, $V_0(\tau) = A_{|\tau|} \cap P$ for each $\tau \in \mathrm{prf}(\sigma)$. It is a routine exercise to establish that for all $\tau \in \mathrm{prf}(\sigma)$ and $\alpha \in \widehat{CL}(\alpha_0)$, $M_0, \tau \models \alpha$ iff $\alpha \in A_{|\tau|}$. Hence $M_0, \varepsilon \models \alpha_0$ as required.

The system $\mathcal{DLTL}$ is obtained by *replacing* (A10) and (A11) with the following axiom schemes:

(A12) $\alpha \, \mathcal{U}^\pi \beta \supset \langle \pi \rangle \beta$.

(A13) $\alpha \, \mathcal{U}^\pi \beta \equiv \beta \vee \left(\alpha \wedge \bigvee_{a \in \Sigma} \langle a \rangle \bigvee_{\pi' \in \delta_a(\pi)} \alpha \, \mathcal{U}^{\pi'} \beta\right), \quad (\varepsilon \in ||\pi||)$.

(A14) $\alpha \, \mathcal{U}^\pi \beta \equiv \alpha \wedge \bigvee_{a \in \Sigma} \langle a \rangle \bigvee_{\pi' \in \delta_a(\pi)} \alpha \, \mathcal{U}^{\pi'} \beta, \quad (\varepsilon \notin ||\pi||)$.

It is an easy exercise to extend the completeness argument for $\mathcal{DLTL}^-$ to $\mathcal{DLTL}$. Thus we have:

Theorem 6.4

(i) *$\mathcal{DLTL}^-$ is a sound and complete axiomatization of the set of valid formulas of* $\mathrm{DLTL}^-(\Sigma)$.

(ii) *$\mathcal{DLTL}$ is a sound and complete axiomatization of the set of valid formulas of* $\mathrm{DLTL}(\Sigma)$.

7 Conclusion

We have presented here an enriched version of LTL called DLTL. The extension is obtained by indexing the until operator of LTL with regular programs. We have shown that in terms of the complexity of the decision procedure and expressiveness, DLTL compares very favourably with ETL. It is worth pointing out here that the decision procedure for DLTL is carried out directly in terms of Büchi automata whereas for ETL it is carried out in terms of the so called set-subword automata, which are then translated to Büchi automata [21]. Two additional results that are available for DLTL are: A characterization of the first order fragment of S1S in terms of the sublogics $\mathrm{DLTL}^-_{\mathbb{SF}}$ and $\mathrm{DLTL}_{\mathbb{SF}}$; and a relatively clean axiomatization of DLTL^- and DLTL. All these results demonstrate that our means of bringing together propositional dynamic and temporal logics in a linear time setting is natural.

It turns out that our idea extends smoothly to richer domains. In particular, we can obtain similar results concerning the so called ω-regular product languages [19] in terms of the product version of DLTL. Roughly speaking, a ω-regular product language is a ω-regular language $L \subseteq \Sigma^\omega$ generated by a distributed alphabet $\{\Sigma_i\}_{i=1}^K$ with $\Sigma = \bigcup_{i=1}^K \Sigma_i$. The language L is a product language in the sense it is a finite union languages of the form $L_1 \otimes L_2 \otimes \cdots \otimes L_K$ with each L_i a regular subset of finite and infinite strings over Σ_i and $\otimes$ standing for the synchronized product operation. In other words $\sigma \in \Sigma^\omega$ is in $L_1 \otimes L_2 \otimes \cdots \otimes L_K$ iff $\sigma \upharpoonright \Sigma_i$ (i.e. the sequence obtained by erasing all symbols from σ that are not in Σ_i) is in L_i for each i. The interesting distributed alphabets are of course those in which the component alphabets are *not* pairwise disjoint. The ω-regular product languages can be used to capture the linear time behaviour of a widely used model of distributed programs. These programs consist of a fixed set of finite state sequential programs that coordinate their behaviours by performing common actions together. Our logical characterization of ω-regular product languages is obtained by taking boolean combinations of formulas in $\bigcup_{i=1}^K \mathrm{DLTL}(\Sigma_i)$.

Yet another extension is obtained by considering ω-regular trace languages over a trace alphabet (Σ, I) with $I \subseteq \Sigma \times \Sigma$ being an irreflexive and symmetric independence relation. Here by traces we mean the class of Σ-labelled posets known as Mazurkiewicz traces. These objects play a fundamental role in the theory of distributed systems [2, 11]. It turns out that ω-regular trace languages can be captured by a natural extension of $\mathrm{DLTL}(\Sigma)$ denoted $\mathrm{DLTL}(\Sigma, I)$. The main idea is that one must use regular "parallel" programs instead of regular programs to index the until operator. What is interesting about this characterization is that so far there has been no comparable result concerning ω-regular trace languages in terms of a modal logic with a fixed point operator or in terms of a linear time temporal logic. These results and related generalizations will be reported in a forthcoming paper.

References

[1] J. R. Büchi, On a decision method in restricted second order arithmetic, Proceedings of the International Congress on Logic, Methodology and Philosophy of Science (Stanford University Press, 1960) 1–11.

[2] V. Diekert and G. Rozenberg, eds., The Book of Traces (World Scientific Singapore, 1995).

[3] M. J. Fischer and R. E. Ladner, Propositional dynamic logic of regular programs, Journal of Computer and System Sciences 18(2) (1979) 194–211.

[4] A. Gabbay, A. Pnueli, S. Shelah, and J. Stavi, On the temporal analysis of fairness, Proceedings of the 7th Annual Symposium on Principles of Programming Languages (ACM, 1980) 163–173.

[5] D. Harel, Dynamic logic, in D. Gabbay and F. Guenthner, eds., "Handbook of Philosophical Logic - Vol. II" (Reidel Dordrecht, 1984) 497–604.

[6] D. Harel, D. Kozen, and R. Parikh, Process logic: expressiveness, decidability, completeness, Journal of Computer and System Sciences 25 (1982) 144–170.

[7] J. Hromkovič, S. Seibert, and T. Wilke, Translating regular expressions into small ε-free nondeterministic automata, Proceedings of the 12th Annual Symposium on Theoretical Aspects of Computer Science, Lecture Notes in Computer Science 1200 (Springer-Verlag, 1997) 55–66.

[8] H. R. Kamp, Tense Logic and the Theory of Linear Order, Ph.D. thesis, University of California (1968)

[9] D. Kozen and R. Parikh, An elementary proof of the completeness of PDL, Theoretical Computer Science 14 (1981) 113–118.

[10] Z. Manna and A. Pnueli, The Temporal Logic of Reactive and Concurrent Systems (Specification) (Springer-Verlag, 1992).

[11] A. Mazurkiewicz, Concurrent program schemes and their interpretations, Technical report DAIMI PB-78, Computer Science Department, University of Aarhus (1977).

[12] R. McNaughton, Testing and generating infinite sequences by a finite automaton, Information and Control 9 (1966) 521–530.

[13] R. McNaughton and S. Papert, Counter-Free Automata (MIT Press, 1971).

[14] H. Nishimura, Descriptively complete process logic, Acta Informatica 14(4) (1980) 359–369.

[15] A. Pnueli, The temporal logic of programs, Proceedings of the 18th Annual Symposium on Foundations of Computer Science (IEEE, 1977) 46–57.

[16] V. R. Pratt, Process logic, Proceedings of the 6th Symposium on Principles of Programming Languages (ACM, 1979) 93–100.

[17] K. Segerberg, A completeness theorem in the modal logic of programs, Notices AMS 24(6) (1977) A–522.

[18] A. P. Sistla, Theoretical issues in the design and verification of distributed systems, Ph.D. Thesis, Harvard University (1983).

[19] P. S. Thiagarajan, PTL over product state spaces, Report TCS-95-4, School of Mathematics, SPIC Science Foundation, Madras (1995)

[20] W. Thomas, Automata over infinite objects, in: J. van Leeuwen, ed., Handbook of Theoretical Computer Science, Vol. B: Formal Models and Semantics (Elsevier/MIT Press, 1990) 133–191.

[21] M. Vardi and P. Wolper, Reasoning about infinite computations, Information and Computation 115(1) (1994) 1–35.

[22] P. Wolper, Temporal logic can be more expressive, Proceedings of the 22nd Annual Symposium on Foundations of Computer Science (IEEE, 1981) 340–348.

[23] P. Wolper, M. Y. Vardi, and A. P. Sistla, Reasoning about infinite computation paths, Proceedings of the 24nd Annual Symposium on Foundations of Computer Science (IEEE, 1983) 185–194.

Recent BRICS Report Series Publications

RS-97-8 **Jesper G. Henriksen and P. S. Thiagarajan.** ***Dynamic Linear Time Temporal Logic*****. April 1997. 33 pp.**

RS-97-7 **John Hatcliff and Olivier Danvy.** ***Thunks and the λ-Calculus (Extended Version)*****. March 1997. 55 pp. Extended version of article to appear in the** ***Journal of Functional Programming*****.**

RS-97-6 **Olivier Danvy and Ulrik P. Schultz.** ***Lambda-Dropping: Transforming Recursive Equations into Programs with Block Structure*****. March 1997. 53 pp. Extended version of an article to appear in the 1997 ACM SIGPLAN Symposium on Partial Evaluation and Semantics-Based Program Manipulation (PEPM '97), Amsterdam, The Netherlands, June 1997.**

RS-97-5 **Kousha Etessami, Moshe Y. Vardi, and Thomas Wilke.** ***First-Order Logic with Two Variables and Unary Temporal Logic*****. March 1997. 18 pp. To appear in** ***Twelfth Annual IEEE Symposium on Logic in Computer Science*****, LICS '97 Proceedings.**

RS-97-4 **Richard Blute, Josée Desharnais, Abbas Edalat, and Prakash Panangaden.** ***Bisimulation for Labelled Markov Processes*****. March 1997. 48 pp. To appear in** ***Twelfth Annual IEEE Symposium on Logic in Computer Science*****, LICS '97 Proceedings.**

RS-97-3 **Carsten Butz and Ieke Moerdijk.** ***A Definability Theorem for First Order Logic*****. March 1997. 10 pp.**

RS-97-2 **David A. Schmidt.** ***Abstract Interpretation in the Operational Semantics Hierarchy*****. March 1997. 33 pp.**

RS-97-1 **Olivier Danvy and Mayer Goldberg.** ***Partial Evaluation of the Euclidian Algorithm (Extended Version)*****. January 1997. 16 pp. To appear in the journal** ***Lisp and Symbolic Computation*****.**

RS-96-62 **P. S. Thiagarajan and Igor Walukiewicz.** ***An Expressively Complete Linear Time Temporal Logic for Mazurkiewicz Traces*****. December 1996. i+13 pp. To appear in** ***Twelfth Annual IEEE Symposium on Logic in Computer Science*****, LICS '97 Proceedings.**

RS-96-61 **Sergei Soloviev.** ***Proof of a Conjecture of S. Mac Lane*****. December 1996. 53 pp. Extended abstract appears in Pitt, Rydeheard and Johnstone, editors,** ***Category Theory and Computer Science: 6th International Conference*****, CTCS '95 Proceedings, LNCS 953, 1995, pages 59–80.**