LFCS

Laboratory for Foundations of Computer Science
Department of Computer Science - University of Edinburgh

Translating CTL* into the modal μ-calculus

by

Mads Dam

Translating CTL* into the modal μ-calculus

LFCS Report Series

ECS-LFCS-90-123

LFCS
Department of Computer Science
University of Edinburgh
The King's Buildings
Edinburgh EH9 3JZ

November 1990

Translating CTL* into the modal μ-calculus

Mads Dam
Department of Computer Science
University of Edinburgh

1 Introduction

The modal μ-calculus L_μ [10] is a very general and expressive temporal logic encompassing a wide range of logics such as PDL [6], PDLΔ [15], Process Logic (PL) [8], linear-time temporal logic [7], the branching-time CTL [2], and CTL* [4] which captures both of the latter two. It arises by the addition to Hennessy-Milner logic [9] of least and greatest fixpoints of syntactically monotone operators. This extension preserves the characterisation of bisimulation equivalence and thus provides a natural temporal logic for process calculi such as CCS [12, 11, 13]. A local model checker for checking L_μ-properties against finite-state (CCS) processes due to Stirling and Walker [14] has been implemented in the Edinburgh Concurrency Workbench [3].

A pragmatic point strongly counting against the practical use of L_μ, however, is its lack of transparency: already at the second level of nesting of fixpoints formulas can become quite unintelligible. It is thus a matter of great interest to provide direct translations into L_μ of the logics mentioned above—not only is this the most illuminating way of proving relative expressiveness results, but translations can also provide "macros" which may be much easier in use than the L_μ primitives themselves. For some of these logics (PDL, PDLΔ, CTL) such a translation os easily found. Here we provide a direct translation into L_μ of CTL*. Previously, only an indirect translation through PDLΔ was known (unpublished work by Wolper [17]).

2 The modal μ-calculus

We consider a slight extension of L_μ introduced by Bradfield and Stirling [1]. Formulas $\phi, \psi, \gamma \in \mathcal{F}_\mu$ of this language are generated by

$$\phi ::= Y \mid \neg\phi \mid \phi_1 \vee \phi_2 \mid [K]\phi \mid \mu Y.\phi$$

where Y (and Z) ranges over propositional variables (of which we have a countable supply), K over subsets of a label set $\mathcal{L}$, and where $\mu Y.\phi$ is subject to the syntactic monotonicity condition that all free occurrences of Y in ϕ lie in the scope of an even number of negations. Other connectives are derived in the usual way—in particular: $<K>\phi \triangleq \neg[K]\neg\phi$, $\nu Y.\phi \triangleq \neg\mu Y.\neg\phi[\neg Y/Y]$ (using a standard substitution notation), $\Box\phi \triangleq [\mathcal{L}]\phi$, and $\Diamond\phi \triangleq <\mathcal{L}>\phi$. Also we can introduce the abbreviations $\bot \triangleq Y \wedge \neg Y$ and $\top \triangleq \neg\bot$ for a distinguished Y. We use σ as a metavariable ranging over $\{\mu, \nu\}$.

The semantics of formulas relative to a fixed transition system $T = (S, \{\xrightarrow{a}\}_{a\in\mathcal{L}})$ and a valuation $\mathcal{V} : Y \mapsto A \subseteq S$ is determined by the mapping $||\cdot||$:

$$\begin{aligned}
||Y||\mathcal{V} &= \mathcal{V}(Y) \\
||\neg\phi||\mathcal{V} &= \overline{||\phi||\mathcal{V}} \\
||\phi \wedge \psi||\mathcal{V} &= ||\phi||\mathcal{V} \cap ||\psi||\mathcal{V} \\
||[K]\phi||\mathcal{V} &= \{s \in S \mid \forall s' \in S, a \in K. s \xrightarrow{a} s' \text{ implies } s' \in ||\phi||\mathcal{V}\} \\
||\mu Y.\phi||\mathcal{V} &= \bigcap\{A \subseteq S \mid ||\phi||\mathcal{V}[Y \mapsto A] \subseteq A\}
\end{aligned}$$

Here $\mathcal{V}[Y \mapsto A]$ is the obvious update of $\mathcal{V}$. Two formulas ϕ and ψ are *equivalent*, if for all T and $\mathcal{V}$, $||\phi||\mathcal{V} = ||\psi||\mathcal{V}$. A formula ϕ is *well-guarded*, if whenever $\sigma Y.\psi$ is a subformula of ϕ then each occurrence of Y in ψ is within the scope of a modal operator. We generate only well-guarded formulas. Moreover it may be checked that for each formula ψ there is an equivalent well-guarded ψ'. We consequently restrict attention to well-guarded formulas.

We give a tableau system generalising Stirling and Walker's [14] to arbitrary models, characterising the relation $s \in ||\phi||\mathcal{V}$. Briefly, we show that in order to check $s \in ||\phi||\mathcal{V}$ it suffices to check that atomic propositions hold as appropriately at successors of s and that no μ-constant need be unfolded infinitely often along a path from s. Streett and Emerson [16] prove a closely related predecessor of this result. A key ingredient is the use of *constants* $V, W, \ldots$ and *definition lists*: finite sequences $\Delta = (V_1, \phi_1), \ldots, (V_n, \phi_n)$ of definition pairs for which each V_i is unique and each ϕ_i only mention constants among $\{V_1, \ldots, V_{i-1}\}$. Then $\mathrm{dom}\Delta = \{V_1, \ldots, V_n\}$ and $\Delta(V_i) = \phi_i$. For $V \notin \mathrm{dom}\Delta$, $\Delta \cdot (V = \phi)$ extends Δ to the right by the pair (V, ϕ), and Δ^* is the obvious extension of Δ to arbitrary formulas which replace all constants in $\mathrm{dom}\Delta$ by their definitions. The tableau system is presented in terms of a derivation relation $\rightarrow$ on sequents of the form $s \vdash_\Delta \phi$. This is a minimal relation satisfying the following properties:

$s \vdash_\Delta \neg\neg\phi \rightarrow s \vdash_\Delta \phi$

$s \vdash_\Delta \phi_1 \wedge \phi_2 \rightarrow s \vdash_\Delta \phi_i$ for $i = 1$ and $i = 2$

$s \vdash_\Delta \phi_1 \vee \phi_2 \rightarrow s \vdash_\Delta \phi_i$ for $i = 1$ or $i = 2$

$s \vdash_\Delta [K]\phi \rightarrow s' \vdash_\Delta \phi$ whenever $s \xrightarrow{a} s'$ and $a \in K$

$s \vdash_\Delta <K>\phi \rightarrow s' \vdash_\Delta \phi$ for some s' and $a \in K$ s.t. $s \xrightarrow{a} s'$

$$s \vdash_\Delta \sigma Y.\phi \rightarrow s \vdash_{\Delta \cdot V = \sigma Y.\phi} V \text{ for some } V \text{ s.t. } V \notin \mathrm{dom}\Delta$$

$$s \vdash_\Delta V \rightarrow s \vdash_\Delta \phi(V/Y) \text{ if } \Delta(V) = \sigma Y.\phi$$

A tableau is *partially successful* if all its terminal nodes are true, i.e. whenever $s \vdash_\Delta \phi \not\rightarrow$ then $s \in \|\Delta^*(\phi)\|\mathcal{V}$. Note that a node $s \vdash_\Delta \phi$ is terminal just in case ϕ has one of the forms Y or $\neg Y$. A tableau is *(totally) successful* if moreover there is no constant V and formula $\mu Y.\phi$ s.t. for some infinite path

$$\pi = s_1 \vdash_{\Delta_1} \phi_1 \rightarrow \ldots \rightarrow s_i \vdash_{\Delta_i} \phi_i \rightarrow \ldots$$

for infinitely many i, $\phi_i = V$ and $\Delta_i(V) = \mu Y.\phi$. The proof of soundness and completeness is closely related to the corresponding proof in [14].

Theorem 2.1 *$s \vdash_\Delta \phi$ has a successful tableau iff $s \in \|\Delta^*(\phi)\|\mathcal{V}$.*

PROOF: See appendix. □

3 CTL*

Corresponding to the extension of L_μ by K-indexed modalities we consider a slight extension of CTL* by K-indexed nexttime operators. Moreover, as in PL (c.f. [4, 8]) we do not distinguish between state- and path-formulas—this gives a slightly more succinct account of syntax and semantics. Formulas $\phi \in \mathcal{F}_*$ are generated by

$$\phi ::= Y \mid \neg\phi \mid \phi_1 \wedge \phi_2 \mid (K)\phi \mid \phi_1 \mathrm{U} \phi_2 \mid \mathrm{A}\phi.$$

The dual of the universal path quantifier A is the existential one defined by $\mathrm{E}\phi \triangleq \neg\mathrm{A}\neg\phi$, and Q ranges over $\{\mathrm{A},\mathrm{E}\}$. Other connectives are derived as usual. In particular $\mathrm{F}\phi \triangleq \top\mathrm{U}\phi$, $\mathrm{G}\phi \triangleq \neg\mathrm{F}\neg\phi$, $\phi\neg\mathrm{U}\psi \triangleq \neg(\phi\mathrm{U}\psi)$ and $\mathrm{O}\phi \triangleq (\mathcal{L})\phi$. Moreover we admit the notations $\mathrm{A}\Phi$ and $\mathrm{E}\Phi$ where Φ is a finite subset of $\mathcal{F}_*$. A formula of one of these forms is called a state-formula. In the first case Φ is to be understood disjunctively and in the second conjunctively, i.e. $\mathrm{A}\Phi \triangleq \mathrm{A}\bigvee\Phi$ and $\mathrm{E}\Phi \triangleq \mathrm{E}\bigwedge\Phi$.

As in [5] formulas are interpreted over R-generable models, i.e. models for which paths are generated by the transition relations. Let $T = (S, \{\xrightarrow{a}\}_{a\in\mathcal{L}})$ and $\mathcal{V}$ be given. A *path*, or derivation sequence, through T is a maximal sequence $\sigma = s_1 \xrightarrow{a_1} \cdots \xrightarrow{a_{i-1}} s_i \xrightarrow{a_i} \cdots$. Then $\sigma(i) \triangleq s_i$ and $l(\sigma,i) \triangleq a_i$. We assume for technical reasons that for any state s an a and s' can be found s.t. $s \xrightarrow{a} s'$, hence derivation sequences are infinite. With this proviso models for CTL* and L_μ are the same. Relative to T the semantics of formulas is now given by the relation $\sigma, i \models_\mathcal{V} \phi$ defined by

$$\sigma, i \models_\mathcal{V} Y \text{ iff } \sigma(i) \in \mathcal{V}(Y)$$

$\sigma, i \models_V \neg\phi$ iff $\sigma, i \not\models_V \phi$

$\sigma, i \models_V \phi_1 \wedge \phi_2$ iff $\sigma, i \models_V \phi_1$ and $\sigma, i \models_V \phi_2$

$\sigma, i \models_V (K)\phi$ iff $l(\sigma, i) \in K$ and $\sigma, i+1 \models_V \phi$

$\sigma, i \models_V \phi_1 \mathrm{U} \phi_2$ iff $\exists j \geq i. \sigma, j \models_V \phi_2$ and $\forall k : i \leq k < j. \sigma, k \models_V \phi_1$

$\sigma, i \models_V \mathrm{A}\phi$ iff $\forall \sigma', j.$ if $\sigma(i) = \sigma'(j)$ then $\sigma', j \models_V \phi$

This logic is capable of expressing succinctly not only safety and liveness related connectives such as the CTL-expressible AF and AG, but beyond those also fairness-related connectives such as AGF, and in general arbitrary nestings and Boolean combinations of linear and branching-time connectives for which the task of finding equivalent L_μ-formulations may present considerable difficulties.

4 Syntax trees

The translation proceed in two stages. Given a CTL* formula we first build a *syntax tree* for it, to expose its recursive properties. Such trees are then used by the main translation function of section 5 below to generate the resulting L_μ formula. Syntax trees are built using the following annotated rules (where I in particular is the identity operator):

(i) A-rules:

$$I\colon\ \frac{\mathrm{A}(\Phi, \neg\neg\phi)}{\mathrm{A}(\Phi, \phi)} \qquad \vee\colon\ \frac{\mathrm{A}(\Phi, Y)}{\mathrm{A}\Phi \quad Y} \qquad \vee\colon\ \frac{\mathrm{A}(\Phi, \neg Y)}{\mathrm{A}\Phi \quad \neg Y}$$

$$I\colon\ \frac{\mathrm{A}(\Phi, \phi \vee \psi)}{\mathrm{A}(\Phi, \phi, \psi)} \qquad \wedge\colon\ \frac{\mathrm{A}(\Phi, \phi \wedge \psi)}{\mathrm{A}(\Phi, \phi) \quad \mathrm{A}(\Phi, \psi)}$$

$$\vee\colon\ \frac{\mathrm{A}(\Phi, \mathrm{A}\Psi)}{\mathrm{A}\Phi \quad \mathrm{A}\Psi} \qquad \vee\colon\ \frac{\mathrm{A}(\Phi, \mathrm{E}\Psi)}{\mathrm{A}\Phi \quad \mathrm{E}\Psi}$$

$$\wedge\colon\ \frac{\mathrm{A}(\Phi, \phi_1 \mathrm{U} \phi_2)}{\mathrm{A}(\Phi, \phi_2, \phi_1) \quad \mathrm{A}(\Phi, \phi_2, \mathrm{O}(\phi_1 \mathrm{U} \phi_2))}$$

$$\wedge\colon\ \frac{\mathrm{A}(\Phi, \neg(\phi_1 \mathrm{U} \phi_2))}{\mathrm{A}(\Phi, \neg\phi_2) \quad \mathrm{A}(\Phi, \neg\phi_1, \mathrm{O}\neg(\phi_1 \mathrm{U} \phi_2))}$$

$$I\colon\ \frac{\mathrm{A}(\Phi, \neg(K)\phi)}{\mathrm{A}(\Phi, (K)\neg\phi, \neg(K)\bot)} \ (\phi \neq \bot) \qquad I\colon\ \frac{\mathrm{A}(\Phi, \neg(K_1)\bot, \neg(K_2)\bot)}{\mathrm{A}(\Phi, \neg(K_1 \cup K_2)\bot)}$$

$$I\colon\ \frac{\mathrm{A}(\Phi, (K_1)\phi_1, (K_2)\phi_2)}{\mathrm{A}(\Phi, (K_1 - K_2)\phi_1, (K_2 - K_1)\phi_2, (K_1 \cap K_2)\phi_1 \vee \phi_2)} \quad (K_1 \cap K_2 \neq \emptyset)$$

$$\Omega\colon\ \frac{\mathrm{A}((K_1)\phi_1,\ldots,(K_n)\phi_n,\neg(K)\bot)}{\mathrm{A}(\phi_1)\quad\cdots\quad\mathrm{A}(\phi_n)}\quad(\forall i,j.i\neq j\supset K_i\cap K_j=\emptyset)$$

where

$$\Omega=\lambda Y_1,\ldots,Y_n.[K_1\cap K]Y_1\wedge\cdots\wedge[K_n\cap K]Y_n\wedge[K-(K_1\cup\cdots\cup K_n)]\bot$$

(ii) E-rules:

$$I\colon\ \frac{\mathrm{E}(\Phi,\neg\neg\phi)}{\mathrm{E}(\Phi,\phi)}\qquad\wedge\colon\ \frac{\mathrm{E}(\Phi,Y)}{\mathrm{E}\Phi\quad Y}\qquad\wedge\colon\ \frac{\mathrm{E}(\Phi,\neg Y)}{\mathrm{E}\Phi\quad\neg Y}$$

$$\vee\colon\ \frac{\mathrm{E}(\Phi,\phi\vee\psi)}{\mathrm{E}(\Phi,\phi)\quad\mathrm{E}(\Phi,\psi)}\qquad I\colon\ \frac{\mathrm{E}(\Phi,\phi\wedge\psi)}{\mathrm{E}(\Phi,\phi,\psi)}$$

$$\wedge\colon\ \frac{\mathrm{E}(\Phi,\mathrm{A}\Psi)}{\mathrm{E}\Phi\quad\mathrm{A}\Psi}\qquad\wedge\colon\ \frac{\mathrm{E}(\Phi,\mathrm{E}\Psi)}{\mathrm{E}\Phi\quad\mathrm{E}\Psi}$$

$$\vee\colon\ \frac{\mathrm{E}(\Phi,\phi_1\mathrm{U}\phi_2)}{\mathrm{E}(\Phi,\phi_2)\quad\mathrm{E}(\Phi,\phi_1,\mathrm{O}(\phi_1\mathrm{U}\phi_2))}$$

$$\vee\colon\ \frac{\mathrm{E}(\Phi,\neg(\phi_1\mathrm{U}\phi_2))}{\mathrm{E}(\Phi,\neg\phi_2,\neg\phi_1)\quad\mathrm{E}(\Phi,\neg\phi_2,\mathrm{O}\neg(\phi_1\mathrm{U}\phi_2))}$$

$$I\colon\ \frac{\mathrm{E}(\Phi,(K)\phi)}{\mathrm{E}(\Phi,\neg(K)\neg\phi,(K)\top)}\quad(\phi\neq\top)\qquad I\colon\ \frac{\mathrm{E}(\Phi,(K_1)\top,(K_2)\top)}{\mathrm{E}(\Phi,(K_1\cap K_2)\top)}$$

$$I\colon\ \frac{\mathrm{E}(\Phi,\neg(K_1)\phi_1,\neg(K_2)\phi_2)}{\mathrm{E}(\Phi,\neg(K_1-K_2)\phi_1,\neg(K_2-K_1)\phi_2,\neg(K_1\cap K_2)\phi_1\vee\phi_2)}\quad(K_1\cap K_2\neq\emptyset)$$

$$\Omega\colon\ \frac{\mathrm{E}(\neg(K_1)\neg\phi_1,\ldots,\neg(K_n)\neg\phi_n,(K)\top)}{\mathrm{E}(\phi_1)\quad\cdots\quad\mathrm{E}(\phi_n)}\quad(\forall i,j.i\neq j\supset K_i\cap K_j=\emptyset)$$

where

$$\Omega=\lambda Y_1,\ldots,Y_n.\langle K_1\cap K\rangle Y_1\vee\cdots\vee\langle K_n\cap K\rangle Y_n\vee\langle K-(K_1\cup\cdots\cup K_n)\rangle\top$$

We refer to the two rules labelled Ω as the *transition rules.* If $|\mathcal{L}|=1$ the nexttime-related rules can be replaced by the following two:

$$\Box\colon\ \frac{\mathrm{A}(\mathrm{O}\phi_1,\ldots,\mathrm{O}\phi_n)}{\mathrm{A}(\phi_1,\ldots,\phi_n)}\qquad\Diamond\colon\ \frac{\mathrm{E}(\mathrm{O}\phi_1,\ldots,\mathrm{O}\phi_n)}{\mathrm{E}(\phi_1,\ldots,\phi_n)}$$

The following proposition expresses the correctness of the syntax tree rules.

Proposition 4.1 *If* $\Omega\colon\ \dfrac{\phi}{\phi_1\ \cdots\ \phi_n}$ *is an instance of any of the above rules then* $s\models_{\mathcal{V}}\Omega(\phi_1,\ldots,\phi_n)$ *iff* $s\models_{\mathcal{V}}\phi$, *where* $[K]$ *is interpreted as* $\mathrm{A}\neg(K)\neg$ *and* $\langle K\rangle$ *as* $\mathrm{E}(K)$.

PROOF: By inspection of the rules. □

A *syntax tree*, t, for a state-formula $\phi_0 = \mathrm{Q}\Phi_0$ is a tree with root $\underline{n}_0$ labelled ϕ_0, generated by the above rules, and such that a node $\underline{n}$ labelled by ϕ (written $\underline{n} : \phi$) has a successor in the tree just in case $\underline{n}$ is not *terminal.* This means that either

(i) no rule is applicable to $\underline{n}$, or

(ii) some node $\underline{n}'$ strictly above $\underline{n}$ on the path from $\underline{n}_0$ to $\underline{n}$ is also labelled by ϕ.

In the latter case $\underline{n}$ is a *preterminal* and $\underline{n}'$ is its *companion.* Let $\rightarrow$ denote the successor relation on nodes given by the syntax tree rules and $\rightarrowtriangle$ the elementwise descendancy relation defined in the obvious way such that for instance in a transition $(\underline{n} : \mathrm{A}(\Phi, \phi \vee \psi)) \rightarrow (\underline{n}' : \mathrm{A}(\Phi, \phi, \psi))$, $\phi \vee \psi \rightarrowtriangle \phi$ and $\phi \vee \psi \rightarrowtriangle \psi$. Some care must be taken such that e.g. in a transition $(\underline{n} : \mathrm{A}(\neg(\phi_1 \mathrm{U} \phi_2), \mathrm{O}\neg(\phi_1 \mathrm{U} \phi_2))) \rightarrow (\underline{n}' : \mathrm{A}(\neg\phi_1, \mathrm{O}\neg(\phi_1 \mathrm{U} \phi_2)))$, $\neg(\phi_1 \mathrm{U} \phi_2) \rightarrowtriangle \neg\phi_1$ but *not* $\neg(\phi_1 \mathrm{U} \phi_2) \rightarrowtriangle \mathrm{O}\neg(\phi_1 \mathrm{U} \phi_2)$. We use Π (π) for $\rightarrow$- ($\rightarrowtriangle$-) derivations, or paths, and write $\pi \in \Pi$ if every consecutive descendancy transition in π is derived from a corresponding transition in Π.

Proposition 4.2 *Syntax trees are finite.*

PROOF: Let $\mathrm{cl}_{\mathrm{A}}(\Phi)$ be the least set containing all $\phi \in \Phi$ s.t.

(i) If $\neg\neg\phi \in \mathrm{cl}_{\mathrm{A}}(\Phi)$ or $(K)\phi \in \mathrm{cl}_{\mathrm{A}}(\Phi)$ then $\phi \in \mathrm{cl}_{\mathrm{A}}(\Phi)$

(ii) If $\Omega(\phi, \psi) \in \mathrm{cl}_{\mathrm{A}}(\Phi)$ for $\Omega \in \{\wedge, \vee, \mathrm{U}, \neg\mathrm{U}\}$ then $\phi, \psi \in \mathrm{cl}_{\mathrm{A}}(\Phi)$

(iii) If $\Omega(\phi, \psi) \in \mathrm{cl}_{\mathrm{A}}(\Phi)$ for $\Omega \in \{\mathrm{U}, \neg\mathrm{U}\}$ then $\mathrm{O}\Omega(\phi, \psi) \in \mathrm{cl}_{\mathrm{A}}(\Phi)$

(iv) If $\neg(K)\phi \in \mathrm{cl}_{\mathrm{A}}(\Phi)$ and $\phi \neq \bot$ then $(K)\neg\phi, \neg(K)\bot \in \mathrm{cl}_{\mathrm{A}}(\Phi)$

(v) If $\neg(K_1)\bot, \neg(K_2)\bot \in \mathrm{cl}_{\mathrm{A}}(\Phi)$ then $\neg(K_1 \cup K_2)\bot \in \mathrm{cl}_{\mathrm{A}}(\Phi)$

(vi) If $(K_1)\phi_1, (K_2)\phi_2 \in \mathrm{cl}_{\mathrm{A}}(\Phi)$ and $K_1 \cap K_2 \neq \emptyset$ then $(K_1 - K_2)\phi_1$, $(K_2 - K_1)\phi_2$, $(K_1 \cap K_2)(\phi_1 \vee \phi_2) \in \mathrm{cl}_{\mathrm{A}}(\Phi)$

The operator cl_{E} is defined similarly to match the E-syntax tree rules. It is not hard to show that for Φ finite both $\mathrm{cl}_{\mathrm{A}}(\Phi)$ and $\mathrm{cl}_{\mathrm{E}}(\Phi)$ are finite sets. Suppose now that $\Pi : \mathrm{Q}_1\Phi_1 \rightarrow \cdots$ is an infinite syntax tree path. Then $\mathrm{Q}_i\Phi_i \neq \mathrm{Q}_j\Phi_j$ whenever $i \neq j$, and then there must be an infinite sequence $\mathrm{Q}_{i_1}\Phi_{i_1}, \ldots$ s.t. for all $j \geq 1$, $\mathrm{Q}_{i_{j+1}}\Phi_{i_{j+1}} \in \mathrm{cl}_{\mathrm{Q}_{i_j}}(\Phi_{i_j})$. Notice now that whenever $\phi \in \mathrm{cl}_{\mathrm{Q}}(\Phi)$ then the depth of nesting of path quantifiers of ϕ is strictly smaller than that of $\mathrm{Q}\Phi$, and we have obtained a contradiction. □

Of particular interest are *circular* paths—paths $\Pi : \mathrm{Q}_1\Phi_1 \rightarrow \cdots \rightarrow \mathrm{Q}_m\Phi_m$ for which $\mathrm{Q}_1\Phi_1 = \mathrm{Q}_m\Phi_m$. In this case $\mathrm{Q}_i = \mathrm{Q}_1$ whenever $1 \leq i \leq m$. Moreover nontrivial Π must involve the application of a transition rule—this is what ensures well-guardedness when translating. Let $\pi : \phi_1 \dashrightarrow \cdots \dashrightarrow \phi_m \in \Pi$ and $\phi_1 = \phi_m$. For such paths each $\pi(i)$ will have the form either $\pi(i) = \mathrm{L}(\phi_1, \phi_2)$ or $\pi(i) = \mathrm{OL}(\phi_1, \phi_2)$ for $\mathrm{L} \in \{\mathrm{U}, \neg\mathrm{U}\}$. Then, if $\mathrm{L} = \mathrm{U}$ we call π a *μ-path*, and otherwise a *ν-path.* Note that only "simple loops" are needed, as if $\pi_1 : \phi_1 \dashrightarrow \cdots \dashrightarrow \phi_2$, $\pi_2 : \phi_2 \dashrightarrow \cdots \dashrightarrow \phi_1 \in \Pi$ then $\phi_1 = \phi_2$.

5 The translation

We start by explaining the intuitive idea. Consider the following syntax-tree t rooted in $\underline{n}_0$ and with $\underline{n}_1$ and $\underline{n}_2$ the sole preterminals having $\underline{n}_0$ as companion:

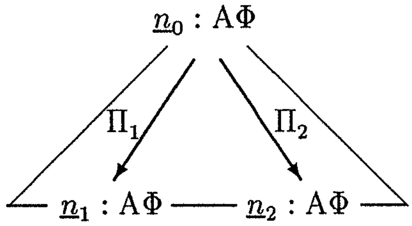

The intention is to translate $\underline{n}_1$ and $\underline{n}_2$ as variables bound during the translation of $\underline{n}_0$. The key is to consider the infinite tree t^ω obtained by repeatedly substituting t for preterminals labelled $\mathrm{A}\Phi$. Any infinite path through t^ω is identical to a composition of copies of Π_1 and Π_2. Such a path is *admissible* w.r.t. $\mathrm{A}\Phi$, if it satisfies a member of Φ.

Two extremal cases are easily identified. If there is a $\phi \in \Phi$ such that for both $i = 1$ and $i = 2$ there is a ν-path $\pi_i \in \Pi_i$ from ϕ to ϕ then $\underline{n}_1$ and $\underline{n}_2$ can be translated as the same ν-variable, as in this situation any infinite path through t^ω will be admissible. Next if neither Π_1 or Π_2 contains a ν-path, both preterminals can be translated as the same μ-variable, as in this case no infinite path as above will be admissible. These cases cover CTL.

The general situation is inbetween. For illustration suppose Π_1 contains just one ν-path from ϕ_1 to ϕ_1 and that Π_2 contains just one ν-path from ϕ_2 to ϕ_2 with $\phi_1 \neq \phi_2$. Then any infinite admissible path through t^ω must be equal to Π_i^ω for $i = 1$ or $i = 2$. The appropriate translation of $\underline{n}_0$ must in this case introduce first a μ-variable to handle the inadmissible paths and next a disjunction of two formulas, each introducing a ν-variable Y_i, $i \in \{1, 2\}$, to handle the infinite path Π_i^ω.

The situation for the existential path-quantifier is dual. For the general definition the translation function tr takes each syntax-tree node $\underline{n} : \phi$ into an L_μ-formula $\mathrm{tr}(\underline{n} : \phi)S$, where S is an environment taking state-formulas $\mathrm{Q}\Phi$ into members of Φ. First for tr applied to a terminal $\underline{n}$ we let $\mathrm{tr}(\underline{n} : \mathrm{A}\emptyset)S = \bot$, $\mathrm{tr}(\underline{n} : \mathrm{E}\emptyset)S = \top$, and $\mathrm{tr}(\underline{n} : (\neg)Y)S = (\neg)Y$.

To account for preterminals and nonterminals we assume for each Φ unique variables $Y_{\mathrm{A}\Phi}$ and $Y_{\mathrm{E}\Phi}$, and for each pair Φ, ϕ with $\phi \in \Phi$ unique variables $Y_{\phi,\mathrm{A}\Phi}$

and $Y_{\phi,\mathrm{E}\Phi}$. Now for $\underline{n} : \phi$ a preterminal with $\underline{n}'$ its companion and ϕ of the form $\mathrm{A}\Phi$ we let $\mathrm{tr}(\underline{n})S = Y_\phi$ if no ν-path π from $S(\phi)$ to $S(\phi)$ in the path from $\underline{n}'$ to $\underline{n}$ exists, and $\mathrm{tr}(\underline{n})S = Y_{S(\Phi),\phi}$ otherwise. Dually, if ϕ has the form $\mathrm{E}\Phi$, $\mathrm{tr}(\underline{n})S = Y_\phi$ if no μ-path as above exists, and $\mathrm{tr}(\underline{n})S = Y_{S(\Phi),\phi}$ otherwise. Finally, let $\underline{n} : \phi$ be a nonterminal, Ω label a transition from $\underline{n}$ to $\underline{n}_1, \ldots, \underline{n}_m$, and let ϕ have the form $\mathrm{A}\Phi$. Then

$$\mathrm{tr}(\underline{n})S = \mu Y_\phi. \bigvee_{\phi' \in \Phi} \nu Y_{\phi',\phi}.\Omega(\mathrm{tr}(\underline{n}_1)S[\phi \mapsto \phi'], \ldots, \mathrm{tr}(\underline{n}_m)S[\phi \mapsto \phi'])$$

where $S[\phi \mapsto \phi']$ is the obvious update of S. Dually, if ϕ has the form $\mathrm{E}\Phi$,

$$\mathrm{tr}(\underline{n})S = \nu Y_\phi. \bigwedge_{\phi' \in \Phi} \mu Y_{\phi',\phi}.\Omega(\mathrm{tr}(\underline{n}_1)S[\phi \mapsto \phi'], \ldots, \mathrm{tr}(\underline{n}_m)S[\phi \mapsto \phi'])$$

Then the root node $\underline{n}_0$ is translated by $\mathrm{tr}(\underline{n}_0) \triangleq \mathrm{tr}(\underline{n}_0)()$, and for $\phi \in \mathcal{F}_*$ and $\psi \in \mathcal{F}_\mu$ we let $\phi \rightsquigarrow \psi$ iff for some syntax-tree with $\underline{n}_0 : \phi$, $\psi = \mathrm{tr}(\underline{n}_0)$.

Example 5.1 Fig. 1 and 2 shows examples of syntax-trees. They use the simplified rules for the case $|\mathcal{L}| = 1$ as well as rules for F and G modified in the obvious way to avoid cluttering up the trees with occurrences of $\top$ and $\bot$. In fig. 1, there is no ν-path from the root to the preterminal labelled (1) whereas there is one to the preterminal labelled (2). Up to equivalence the root $\underline{n}_0$ is translated in the following way:

$$\mathrm{tr}(\underline{n}_0) = \mu Y'.(\nu Z'.\Box((Y \vee Y') \wedge Y')) \vee (\nu Z'.\Box((Y \vee Y') \wedge Z')).$$

This may be simplified further to $\mu Y'.\nu Z'.\Box(Y \vee Y') \wedge Z'$. In fig. 2 there is no ν-path from the root to the preterminal labelled (1), there is one unrolling $\mathrm{G}Z$ to (2), one unrolling $\mathrm{G}Y$ to (3) and one of both to (4). Again up to equivalence the translation of the root $\underline{n}_0$ is

$$\begin{aligned}\mathrm{tr}(\underline{n}_0) \quad = \quad & \mu Y'.(\nu Z'.\Box(Y' \vee Y \vee Z) \wedge (Y' \vee Y) \wedge (Y' \vee Z) \wedge Y') \\ & \vee(\nu Z'.\Box(Y' \vee Y \vee Z) \wedge (Z' \vee Y) \wedge (Y' \vee Z) \wedge Z') \\ & \vee(\nu Z'.\Box(Y' \vee Y \vee Z) \wedge (Y' \vee Y) \wedge (Z' \vee Z) \wedge Z').\end{aligned}$$

As for the previous example the disjunct $\nu Z'.\Box(Y' \vee Y \vee Z) \wedge (Y' \vee Y) \wedge (Y' \vee Z) \wedge Y'$ may be discarded. $\Box$

Note that both the notion of syntax-tree and the translation are self-dual. That is, from any syntax-tree with $\underline{n} : \phi$ a syntax-tree with $\underline{n} : \neg\phi$ can be derived simply by dualising the labelling and annotation of each node. Furthermore $\neg\mathrm{tr}(\underline{n})$ with respect to the original syntax tree and $\mathrm{tr}(\underline{n})$ with respect to the dualised tree are identical.

$$\wedge\ \dfrac{\dfrac{\Box\ \dfrac{\mathrm{A}(\mathrm{OG}Y,\mathrm{OFOG}Y)}{\mathrm{A}(\mathrm{G}Y,\mathrm{FOG}Y)}}{\mathrm{A}(\mathrm{G}Y,\mathrm{OG}Y,\mathrm{OFOG}Y)}}{\vee\ \dfrac{\mathrm{A}(Y,\mathrm{OG}Y,\mathrm{OFOG}Y)}{Y \quad \mathrm{A}(\mathrm{OG}Y,\mathrm{OFOG}Y)^{(1)}} \qquad \mathrm{A}(\mathrm{OG}Y,\mathrm{OFOG}Y)^{(2)}}$$

Figure 1: Example of syntax-tree

$$\wedge\ \dfrac{\dfrac{\Box\ \dfrac{\mathrm{A}\Phi}{\mathrm{A}(\mathrm{FOG}Y,\mathrm{FOG}Z,\mathrm{G}Y,\mathrm{G}Z)}}{\mathrm{A}(\Phi,\mathrm{G}Y,\mathrm{G}Z)}}{\wedge\ \dfrac{\mathrm{A}(\Phi,Y,\mathrm{G}Z)}{\vee\ \dfrac{\mathrm{A}(\Phi,Y,Z)}{\mathrm{A}\Phi^{(1)} \quad Y \quad Z} \qquad \dfrac{\mathrm{A}(\Phi,Y)}{\mathrm{A}\Phi^{(2)} \quad Y}} \qquad \wedge\ \dfrac{\mathrm{A}(\Phi,\mathrm{G}Z)}{\vee\ \dfrac{\mathrm{A}(\Phi,Z)}{\mathrm{A}\Phi^{(3)} \quad Z} \qquad \mathrm{A}\Phi^{(4)}}}$$

Figure 2: Example of syntax tree with $\Phi = \{\mathrm{OFOG}Y, \mathrm{OFOG}Z, \mathrm{OG}Y, \mathrm{OG}Z\}$.

6 The correctness proof

We want to show that if $\underline{n}_0 : \phi_0$ and $s_0 \models_{\mathcal{V}} \phi_0$ then $s_0 \vdash_{()} \mathrm{tr}(\underline{n}_0)$ has a successful tableau. From this the correctness of tr follows. For by theorem 2.1 if follows that if $s_0 \models_{\mathcal{V}} \phi_0$ then $s_0 \in \|\mathrm{tr}(\underline{n}_0)\|\mathcal{V}$, and for the converse if $s_0 \in \|\mathrm{tr}(\underline{n}_0)\|\mathcal{V}$ then $s_0 \models_{\mathcal{V}} \phi_0$ as otherwise by the above observation, $s_0 \vdash_{()} \neg\mathrm{tr}(\underline{n})$ would have a successful tableau—a contradiction by theorem 2.1.

We give a procedure for building a tableau, τ, guided by a syntax-tree and a model. The construction starts with the root syntax-tree node $\underline{n}_0$ and the root tableau-node labelled by $s_0 \vdash_{()} \mathrm{tr}(\underline{n}_0)$. Assume we have reached a tableau-node labelled $s \vdash_{\Delta} \psi$ and a syntax-tree node $\underline{n}$ with $\underline{n} : \phi$ and $\Pi = \underline{n}_0 \rightarrow \cdots \rightarrow \underline{n}_k$ the path from root to $\underline{n}_k = \underline{n}$. Let $P(s \vdash_{\Delta} \psi, \underline{n})$ hold just in case

i) $s \models_{\mathcal{V}} \phi$, and

ii) $\Delta^*(\psi) = (\mathrm{tr}(\underline{n})S)\rho_k \cdots \rho_0$ for some S,

where each substitution ρ_i is determined from $\underline{n}_i : \phi_i = \mathrm{Q}_i\{\phi_{i,1}, \ldots, \phi_{i,l_i}\}$, say, by

$$\rho_i = [\sigma Y_{\phi_{i,1},\phi_i}/Y_{\phi_{i,1},\phi_i}] \cdots [\sigma Y_{\phi_{i,l_i},\phi_i}/Y_{\phi_{i,l_i},\phi_i}][\sigma Y_{\phi_i}/Y_{\phi_i}]$$

where (as in [10]) we use σY to denote the uniquely determined subformula of $\mathrm{tr}(\underline{n}_0)$ of the form $\sigma Y.\phi'$.

If $\underline{n}$ is a "proper" terminal then the construction of τ is either complete or trivial. Next if $\underline{n}$ is a preterminal and $\underline{n}_i$ its companion we continue from $s \vdash_{\Delta}$

ψ and $\underline{n}_i$. Note that $P(s \vdash_\Delta \psi, \underline{n}_i)$, as $(\mathrm{tr}(\underline{n})S)\rho_k \cdots \rho_0 = (\mathrm{tr}(\underline{n}_i)S)\rho_{i-1} \cdots \rho_0$. Assume consequently that $\underline{n}$ is an Ω-annotated nonterminal with descendants $\underline{n}_1' : \phi_1', \ldots, \underline{n}_m' : \phi_m'$. We show that $s \vdash_\Delta \psi$ can be extended according to the tableau-rules s.t. for every successor $s' \vdash_{\Delta'} \psi'$ of $s \vdash_\Delta \psi$ there is an $\underline{n}_l'$ s.t. $P(s' \vdash_{\Delta'} \psi', \underline{n}_l')$. It follows that τ is partially successful.

Case i): ϕ has the form $\mathrm{A}\Phi$. Then

$$\mathrm{tr}(\underline{n}) = \mu Y_\phi. \bigvee_{\phi' \in \Phi} \nu Y_{\phi',\phi}.\Omega(\mathrm{tr}(\underline{n}_1')[\phi \mapsto \phi'], \ldots, \mathrm{tr}(\underline{n}_m')[\phi \mapsto \phi'])$$

Either ψ has the form $\psi = \mu Y_\phi. \bigvee_{\phi' \in \Phi} \nu Y_{\phi',\phi}.\Omega(\psi_1, \ldots, \psi_m)$, or $\psi = V$ for some V with $\Delta(V)$ of this form, or ψ has the form $\psi = \nu Y_{\phi',\phi}.\Omega(\psi_1, \ldots, \psi_m)$ for some $\phi' \in \Phi$, or ψ is a constant V with $\Delta(V)$ of this latter form. In the first case we introduce a (fresh) μ-constant V, and equip it with a scheduler, f_V, picking out a member $f_V(\Phi)$ of Φ in a round-robin fashion. Then $s \vdash_\Delta \psi$ is extended in the following way:

$$\begin{array}{c}
s \vdash_\Delta \psi \\ \hline
s \vdash_{\Delta\cdot(V=\psi)} V \\ \hline
s \vdash_{\Delta\cdot(V=\psi)} \bigvee_{\phi' \in \Phi} \nu Y_{\phi',\phi}\Omega(\psi_1, \ldots, \psi_m)[V/Y_\phi] \\ \hline
s \vdash_{\Delta\cdot(V=\psi)} \nu Y_{f_V(\Phi),\phi}\Omega(\psi_1, \ldots, \psi_m)[V/Y_\phi] \\ \hline
s \vdash_{\Delta\cdot(V=\psi)\cdot(W=\nu Y_{f_V(\Phi),\phi}\Omega(\psi_1,\ldots,\psi_m)[V/Y_\phi])} W \\ \hline
s \vdash_{\Delta\cdot(V=\psi)\cdot(W=\nu Y_{f_V(\Phi),\phi}\Omega(\psi_1,\ldots,\psi_m)[V/Y_\phi])} \Omega(\psi_1, \ldots, \psi_m)[V/Y_\phi][W/Y_{f_V(\Phi),\phi}]
\end{array}$$

In the second case f_V is already defined and all we have to do is to update f_V and expand V. Similarly the third and fourth cases are just subconstructions of these two. Let $s \vdash_{\Delta'} \psi'$ be the sequent obtained at the end of this construction.

We now need to consider Ω. In all cases expect when $\Omega = \vee$ is the construction well-determined. Now $\Omega = \vee$ only when one branch is a strict subformula of Φ of one of the forms Y, $\neg Y$, $\mathrm{A}\Psi$ or $\mathrm{E}\Psi$. In these cases we always choose that "strictly decreasing" branch when doing so is possible—i.e. when $s \models_{\mathcal{V}} Y$, $s \models_{\mathcal{V}} \neg Y$, $s \models_{\mathcal{V}} \mathrm{A}\Psi$ or $s \models_{\mathcal{V}} \mathrm{E}\Psi$ whatever the case may be.

We proceed by cases on Ω. If $\Omega = I$ we just go on to the successor of $\underline{n}$, $\underline{n}_1'$, and note that $P(s \vdash_{\Delta'} \psi', \underline{n}_1')$. Assume instead that $\Omega = \lambda Y_1, \ldots, Y_m.[K_1 \cap K]Y_1 \wedge \cdots \wedge [K_m \cap K]Y_m \wedge [K - (K_1 \cup \cdots \cup K_m)]\bot$ with all K_l pairwise disjoint. Then whenever $s \xrightarrow{a} s'$ and $a \in K_l \cap K$ then $s' \models_{\mathcal{V}} \psi_l$. Also

$$\Delta'^*(\psi_l) = (\mathrm{tr}(\underline{n}_l')S[\phi \mapsto f_V(\Phi)])\rho\rho_k \cdots \rho_1$$

where

$$\rho = [\nu Y_{\phi_{k+1,1},\phi_{k+1}}/Y_{\phi_{k+1,1},\phi_{k+1}}] \cdots [\nu Y_{\phi_{k+1,l_{k+1}},\phi_{k+1}}/Y_{\phi_{k+1,l_{k+1}},\phi_{k+1}}][\mu Y_{\phi_{k+1}}/Y_{\phi_{k+1}}]$$

showing that $P(s' \vdash_{\Delta'} \psi_l, \underline{n}_l')$ as desired. The other cases for $\Omega = \wedge$ and $\Omega = \vee$ are simple exercises given the adopted strategy.

Case ii): Suppose then ϕ has the form $\mathrm{E}\Phi$. The tableau construction is a straightforward variation on case i), once we give a suitable strategy for resolving choices. For this purpose let an occurrence of a subformula $\phi_1 \mathrm{U} \phi_2$ of some $\phi' \in \Phi$ be a *toplevel eventuality* of $\mathrm{E}\Phi$, if $\phi_1 \mathrm{U} \phi_2$ is not within the scope of any operator among U, A, E in ϕ'. An *index* of $\mathrm{E}\Phi$ is a map ι assigning a natural number $\iota(\phi_1 \mathrm{U} \phi_2)$ to each toplevel eventuality $\phi_1 \mathrm{U} \phi_2$ of $\mathrm{E}\Phi$, and then $\phi_1 \mathrm{U}^{\iota(\phi_1 \mathrm{U} \phi_2)} \phi_2$ is the obvious approximation (i.e. $\phi_1 \mathrm{U}^0 \phi_2 = \phi_2$, $\phi_1 \mathrm{U}^{n+1} \phi_2 = \phi_2 \vee (\phi_1 \wedge \mathrm{O}(\phi_1 \mathrm{U}^n \phi_2))$). The *successor* of ι, $\mathrm{succ}(\iota)$, is defined by $\mathrm{succ}(\iota)(\phi_1 \mathrm{U} \phi_2) = \iota(\phi_1 \mathrm{U} \phi_2) - 1$ when $\iota(\phi_1 \mathrm{U} \phi_2) > 0$ and $\mathrm{succ}(\iota)(\phi_1 \mathrm{U} \phi_2) = 0$ otherwise, and the indexing of a state-formula $\mathrm{E}\Phi$ is $\mathrm{E}\Phi[\iota] = \mathrm{E}\{\phi'(\iota) \mid \phi' \in \Phi\}$, where

$$(\mathrm{Q}\Phi)(\iota) = \mathrm{Q}\Phi,\ (\neg(\phi \mathrm{U} \psi))(\iota) = \neg(\phi \mathrm{U} \psi),\ (\neg)Y(\iota) = (\neg)Y$$

$$(\phi \wedge \psi)(\iota) = (\phi(\iota)) \wedge (\psi(\iota)),\ (\phi \vee \psi)(\iota) = (\phi(\iota)) \vee (\psi(\iota)),$$

$$(\mathrm{O}\phi)(\iota) = \mathrm{O}(\phi(\mathrm{succ}(\iota))),\ (\phi \mathrm{U} \psi)(\iota) = \phi \mathrm{U}^{\iota(\phi \mathrm{U} \psi)} \psi$$

It is clear that if $s \models_{\mathcal{V}} \mathrm{E}\Phi$ then there is some index ι appropriate for $\mathrm{E}\Phi$ at s—i.e. such that $s \models_{\mathcal{V}} \mathrm{E}\Phi[\iota]$.

We take indices into account in the tableau-building procedure. If $\underline{n} : \mathrm{E}\Phi$ is the root or the right child of a node $\underline{n}'$ labelled by $\mathrm{A}(\Psi, \mathrm{E}\Phi)$ or $\mathrm{E}(\Psi, \mathrm{E}\Phi)$ then $\underline{n}$ is indexed by some arbitrary ι appropriate for $\mathrm{E}\Phi$ at the given s. For most of the E-rules indexing is obvious—e.g. if $s \models_{\mathcal{V}} \mathrm{E}(\Phi, \phi \vee \psi)[\iota]$ then either $s \models_{\mathcal{V}} \mathrm{E}(\Phi, \phi)[\iota]$ or $s \models_{\mathcal{V}} \mathrm{E}(\Phi, \psi)[\iota]$. As another example, if $s \models_{\mathcal{V}} \mathrm{E}(\Phi, \mathrm{E}\Psi)[\iota]$ then $s \models_{\mathcal{V}} \mathrm{E}\Phi[\iota]$. The only nontrivial cases are the U- and (K)-rules:

i) If $s \models_{\mathcal{V}} \mathrm{E}(\Phi, \phi_1 \mathrm{U} \phi_2)[\iota]$ then we find an ι' which agrees with ι on all toplevel eventualities of $\mathrm{E}(\Phi, \phi_1 \mathrm{U} \phi_2)$ s.t. either $s \models_{\mathcal{V}} \mathrm{E}(\Phi, \phi_2)[\iota']$ or else $s \models_{\mathcal{V}} \mathrm{E}(\Phi, \phi_1)[\iota']$ and $s \models_{\mathcal{V}} \mathrm{E}(\Phi, \mathrm{O}(\phi_1 \mathrm{U} \phi_2))[\iota']$.

ii) If $s \models_{\mathcal{V}} \mathrm{E}(\Phi, \neg(\phi_1 \mathrm{U} \phi_2)[\iota]$ then we again find an ι' which agrees with ι on all toplevel eventualities of $\mathrm{E}(\Phi, \neg(\phi_1 \mathrm{U} \phi_2))$ s.t. $s \models_{\mathcal{V}} \mathrm{E}(\Phi, \neg\phi_2)[\iota']$ and either $s \models_{\mathcal{V}} \mathrm{E}(\Phi, \neg\phi_1)[\iota']$ or $s \models_{\mathcal{V}} \mathrm{E}(\Phi, \mathrm{O}\neg(\phi_1 \mathrm{U} \phi_2))[\iota']$.

iii) If $s \models_{\mathcal{V}} \mathrm{E}(\neg(K_1)\phi_1, \ldots, \neg(K_m)\phi_m, (K)\top)[\iota]$ with all K_i pairwise disjoint then for some s' and $a \in K$, $s \xrightarrow{a} s'$, and if $a \in K_l$ then $s' \models_{\mathcal{V}} \mathrm{E}(\phi_l)[\mathrm{succ}(\iota)]$.

We have thus shown

Lemma 6.1 *τ is partially successful.* □

Suppose then that τ is not totally successful—i.e. we have an infinite derivation

$$s_1 \vdash_{\Delta_1} \psi_1 \to s_2 \vdash_{\Delta_2} \psi_2 \to \ldots \to s_i \vdash_{\Delta_i} \psi_i \to \ldots$$

for which $\psi_i = V$ for infinitely many i with V a μ-constant. Assume in particular $\psi_1 = V$. Correspondingly there will be an infinite path

$$\Pi = \underline{n}_1 : \phi_1 \rightarrow \underline{n}_2 : \phi_2 \rightarrow \ldots \rightarrow \underline{n}_j : \phi_j \rightarrow \ldots$$

and an infinite derivation sequence

$$\sigma = s_1' \rightarrow s_2' \rightarrow \ldots \rightarrow s_k' \rightarrow \ldots$$

s.t. there are monotone maps $g, h : \omega \rightarrow \omega$ s.t. for all $i \geq 1$, $P(s_i \vdash_{\Delta_i} \psi_i, \underline{n}_{g(i)})$ and and for all $i_1, i_2 \geq 1$, if $h(g(i_1)) = h(g(i_2))$ then $s_{i_1} = s_{i_2} = s'_{h(g(i_1))}$.

There are two possibilities. Either for some $\underline{n} : \mathrm{A}\Phi = \phi$,

$$\Delta_1^*(V) = \mathrm{tr}(\underline{n})S = \mu Y_\phi. \bigvee_{\phi' \in \Phi} \nu Y_{\phi',\phi}.\Omega(\mathrm{tr}(\underline{n}_1')S[\phi \mapsto \phi'], \ldots, \mathrm{tr}(\underline{n}_m')S[\phi \mapsto \phi'])$$

or for some $\underline{n} : \mathrm{E}\Phi = \phi$, $\phi' \in \Phi$,

$$\Delta_1^*(V) = \mu Y_{\phi',\phi}.\Omega(\mathrm{tr}(\underline{n}_1')S[\phi \mapsto \phi'], \ldots, \mathrm{tr}(\underline{n}_m')S[\phi \mapsto \phi'])[Y_\phi \mapsto \mathrm{tr}(\underline{n})S]$$

Case i): $\underline{n} : \mathrm{A}\Phi = \phi$. Then for all $j \geq 1$, ϕ_j has the form $\phi_j = \mathrm{A}\Phi_j$. If there is a $\phi' \in \Phi$ s.t. Y_ϕ does not occur freely in $\nu Y_{\phi',\phi}$ then by the scheduling mechanism adopted we have a contradiction. We show by induction on the structure of $\phi_j' \in \Phi_j$ that $\sigma, h(j) \not\models_V \phi_j'$.

$\phi_j' = \gamma_1 \mathrm{U} \gamma_2$. It suffices to show that for all $k \geq h(j)$, $\sigma, k \not\models_V \gamma_2$. Now $\gamma_2 \in \Phi_{j_1}$ for some $j_1 \geq j$ s.t. $h(j) = h(j_1)$, so by the induction hypothesis $\sigma, h(j) \not\models_V \gamma_2$. Also $\mathrm{O}\phi_j' \in \Phi_{j_1}$. Let j_2 be least s.t. $h(j_2) = h(j) + 1$. Note that j_2 exists. We show that $\phi_j' \in \Phi_{j_2}$. The only reason why this could fail is that some strictly decreasing branch has been taken. But this is impossible by the construction. But then $\gamma_2 \in \Phi_{j_3}$ for some $j_3 \geq j_2$ s.t. $j_3 = h(j) + 1$, and thus $\sigma, h(j) + 1 \not\models_V \gamma_2$. Repeating this argument ad infinitum gives the desired result.

$\phi_j' = \neg(\gamma_1 \mathrm{U} \gamma_2)$. It suffices to find a $k \geq h(j)$ s.t. $\sigma, k \models_V \gamma_2$ and $\sigma, k' \models_V \gamma_1$ whenever $h(j) \leq k' < k$. Let $h(g(i)) = h(j)$. The scheduling mechanism ensures some $i' \geq i$ s.t. whenever $\phi' \in \Phi$ some $i_{\phi'}$ can be found with $i \leq i_{\phi'} \leq i'$ and

$$\Delta_{i_{\phi'}}^*(\psi_{i_{\phi'}}) = \nu Y_{\phi',\phi}\Omega(\mathrm{tr}(\underline{n}_1')S[\phi \mapsto \phi'], \ldots, \mathrm{tr}(\underline{n}_m')S[\phi \mapsto \phi'])[\mathrm{tr}(\underline{n})S/Y_\phi]$$

But then for some $i'' \geq i$, $\neg\gamma_2 \in \Phi_{g(i'')}$ as otherwise there would be a ν-path from $\underline{n}_{g(i)}$ to $\underline{n}_{g(i')}$, a contradiction. Let i'' be minimal with this property. Let $k = h(g(i''))$. By the induction hypothesis $\sigma, k \models_V \gamma_2$. Let $h(j) \leq k' < k$. Then by the minimality of i'' we can find j' s.t. $h(j') = k'$ and $\neg\gamma_1 \in \Phi_{j'}$, and then we are done by the induction hypothesis.

The remaining cases are straightforward.

Case ii): $\underline{n} = \mathrm{E}\Phi = \phi$. Again for all $j \geq 1$, ϕ_j has the form $\phi_j = \mathrm{E}\Phi_j$. Let ι_j be the index assigned to ϕ_j in the construction of τ. We show by induction on the structure of $\phi'_j \in \Phi_j$ that $\sigma, h(j) \models_{\mathcal{V}} \phi'_j(\iota_j)$.

$\phi'_j = \gamma_1 \mathrm{U} \gamma_2$. By the assumption, $s'_{h(j)} \models_{\mathcal{V}} \mathrm{E}\Phi_j[\iota_j]$. By induction on $\iota_j(\phi'_j)$ we obtain a smallest $j' \geq j$ s.t. $\gamma_2 \in \Phi_{j'}$ whence by the induction hypothesis, $\sigma, j' \models_{\mathcal{V}} \gamma_2$. For all j'' s.t. $h(j) \leq h(j'') < h(j')$, $\gamma_1 \in \Phi_{j''}$, so $\sigma, j'' \models_{\mathcal{V}} \gamma_1$ and we are done.

$\phi'_j = \neg(\gamma_1 \mathrm{U} \gamma_2)$. The only reason why we can avoid having either $\neg\gamma_1, \neg\gamma_2 \in \Phi_{j_1}$ or $\neg\gamma_2, \mathrm{O}\phi'_j \in \Phi_{j_1}$ for some j_1 with $h(j_1) = h(j_2)$ is if we take some strictly decreasing branch, but this is impossible. By the induction hypothesis, $\sigma, h(j) \models_{\mathcal{V}} \neg\gamma_2$. If $\neg\gamma_1 \in \Phi_{j_1}$ then also $\sigma, h(j) \models_{\mathcal{V}} \neg\gamma_1$ whence $\sigma, h(j) \models_{\mathcal{V}} \phi'_j$. Otherwise let j_2 be least s.t. $h(j_2) = h(j) + 1$. Then $\phi'_j \in \Phi_{j_2}$ and we can repeat the argument. Thus we can conclude that $\sigma, h(j) \models_{\mathcal{V}} \phi'_j$.

The remaining cases are again straightforward. But now we are almost done, for $\psi_i = V$ infinitely often only if there is some infinite μ-path π through Π. Thus for all j, $\pi(j)$ has the form either $\pi(j) = \gamma_1 \mathrm{U} \gamma_2$ or $\pi(j) = \mathrm{O}(\gamma_1 \mathrm{U} \gamma_2)$. By the tableau-construction for some index ι, $s'_1 \models_{\mathcal{V}} \mathrm{E}\Phi_1[\iota]$, so $\sigma, 1 \models_{\mathcal{V}} \Phi_1(\iota)$ as we have just proved. Now either $\gamma_1 \mathrm{U} \gamma_2 \in \Phi_{j_1}$ or $\mathrm{O}(\gamma_1 \mathrm{U} \gamma_2) \in \Phi_{j_1}$, where j_1 is least s.t. $h(j_1) = 1 + \iota(\gamma_1 \mathrm{U} \gamma_2)$, so by the above result again, $\sigma, h(j_1) \models_{\mathcal{V}} (\mathrm{O})\gamma_1 \mathrm{U}^0 \gamma_2$ which is impossible. The proof is thus complete and we have shown

Theorem 6.2 (Correctness of tr) *If $\phi \rightsquigarrow \psi$ then $\{s \mid s \models_{\mathcal{V}} \phi\} = \|\psi\|\mathcal{V}$.* □

7 Concluding remarks

The translation can be optimised in several respects. We can instead of syntax-trees use graphs. In the special case of $|\mathcal{L}| = 1$ we can make do with graphs that are of size exponential in the length n of the input formula, and thus obtain a doubly exponential size-bound for the complete translation. This is true also for the case where we restrict attention to nexttime operators of the forms either $(\{a\})$ or O. By suitably representing the disjuncts and conjuncts introduced in the translation of nonterminals, the size-bound can in these cases be reduced further to $\mathcal{O}(n2^n)$, but then the resulting formula is no longer in L_μ. More fine-tuning can be obtained by noting

(i) not every syntax-tree node can give rise to loop,

(ii) not every formula can be member of a σ-path,

(iii) sufficient syntactic criteria for classifying nonterminals can easily be found: Suppose Φ contains a formula ϕ of one of the forms $\neg(\psi \mathrm{U} \gamma)$ or $\mathrm{O}\neg(\psi \mathrm{U} \gamma)$, and ϕ is not a subformula of any other $\phi' \in \Phi$. Then any $\Pi : A\Phi \rightarrow \cdots \rightarrow A\Phi$ will contain a ν-path from ϕ to ϕ (and dually for E).

It is very likely that by discriminating use of such tuning the complexity will turn out to be manageable in practice.

The reported results may be extended to non-total CTL*-models by taking appropriate account of finite paths. Note finally that a very similar translation strategy can be applied to the full branching-time μ-calculus (with the linear next-time operator and branching quantifiers). Is this translation correct?

Acknowledgements: Thanks to C. Stirling for many valuable discussions. The work was supported by SERC Grant GR/F 32219.

References

[1] J.C. Bradfield and C.P. Stirling. "Local model checking for infinite state spaces," *Lecture Notes in Computer Science* **458** (Springer,1990) pp. 115–125.

[2] E. Clarke and E.A. Emerson. " Design and synthesis of synchronisation skeletons using branching time temporal logic," *Lecture Notes in Computer Science* **131** (Springer,1981) pp. 52–71.

[3] R. Cleaveland, J. Parrow and B. Steffen. "A semantics based verification tool for finite state systems," *Proc. 9th IFIP Symp. on Protocol Specification, Testing, and Verification* North-Holland, 1989.

[4] E.A. Emerson and J. Halpern. ""Sometimes" and "not never" revisited: On branching versus linear time." *Journal of the ACM* **33** (1986) pp. 151–178.

[5] E.A. Emerson and A.P. Sistla. "Deciding full branching time logic," *Information and Control* **61** (Academic Press, 1984) pp. 175–201.

[6] M.J. Fischer and R.E. Ladner. " Propositional dynamic logic of regular programs," *Journal of Computer and System Science* **18** pp. 194–211.

[7] D. Gabbay, A. Pnueli, S. Shelah and J. Stavi. "On the temporal analysis of fairness," in *Proc. of the 7th Annual ACM Symp. on Principles of Programming Languages* (Las Vegas, Nevada, Jan. 1980) (ACM,1980) pp. 163–173.

[8] D. Harel, D. Kozen and R. Parikh. "Process logic: Expressiveness, decidability, and completeness," *Journal of Computer and System Science* **25** (1982) pp. 144–170.

[9] M. Hennessy and R. Milner. "Algebraic laws for nondeterminism and concurrency," *Journal of the ACM* **32** (1985), 137–162.

[10] D. Kozen. "Results on the propositional μ-calculus," *Theoretical Computer Science* **27** (North-Holland, 1983) 333–354.

[11] K.G. Larsen. "Proof systems for Hennessy-Milner logic with recursion," *Lecture Notes in Computer Science* **299** (Springer, 1988).

[12] R. Milner. "*Communication and concurrency*," Prentice Hall International, 1989.

[13] C.P. Stirling. "Temporal logics for CCS," *Lecture Notes in Computer Science* **351** (Springer, 1989).

[14] C.P. Stirling and D.J. Walker. "Local model checking in the modal mu-calculus," *Lecture Notes in Computer Science* **351** (Springer,1989) pp. 369–383.

[15] R.S. Streett. "Propositional dynamic logic of looping and converse is elementarily decidable," *Information and Control* **54** (Academic Press, 1982) pp. 121–141.

[16] R.S. Streett and E.A. Emerson. "An automata theoretic decision procedure for the propositional mu-calculus," *Information and Computation* **81** (Academic Press, 1989) 249–264.

[17] P. Wolper. "A translation from full branching time temporal logic to one letter propositional dynamic logic with looping," unpublished manuscript.

8 Appendix. Proof of theorem 2.1

Theorem 2.1 follows from the soundness and completeness lemmas 8.2 and 8.3 below. We first generalise the finiteness lemma of [14]:

Lemma 8.1 *Let π be an infinite path through a tableau τ. Then there is a unique constant V s.t. $\phi_i = V$ for infinitely many i.*

PROOF: This follows the termination proof of [14], from where we use the degree function $d(\phi)$. Briefly, $d(\phi)$ is the "height" of ϕ with constants having height 0. Then $d(s \vdash_\Delta \phi)$ is $d(\phi)$ if ϕ is not a constant and $d(\Delta(\phi))$ otherwise. Suppose $\pi = s_1 \vdash_{\Delta_1} \phi_1 \rightarrow \cdots \rightarrow s_i \vdash_{\Delta_i} \phi_i \rightarrow \cdots$. The subsequence $\pi' = s'_1 \vdash_{\Delta'_1} V_1, \ldots, s'_i \vdash_{\Delta'_i} V_i, \ldots$ consisting of the constant-sequents of π is infinite. Assume that no constant occurs infinitely often among $V_1, \ldots, V_i, \ldots$ Let then i_0 be maximal s.t. $V_{i_0} = V_0$. Then $d(s'_{i_0+1} \vdash_{\Delta'_{i_0+1}} V_{i_0+1}) < d(s'_0 \vdash_{\Delta'_0} V_0)$, for $\Delta'_{i_0+1}(V_{i_0+1})$ is a strict subformula of $\Delta'_0(V_0)$. Let i_1 be maximal s.t. $V_{i_1} = V_{i_0+1}$. Then by a similar argument $d(s'_{i_1+1} \vdash_{\Delta'_{i_1+1}} V_{i_1+1}) < d(s'_{i_0+1} \vdash_{\Delta'_{i_0+1}} V_{i_0+1})$, and as this construction can be continued ad infinitum some V (which moreover will be unique) must occur infinitely often along π. $\square$

We go on to prove first soundness and then completeness.

Lemma 8.2 *If $s \vdash_\Delta \phi$ has a successful tableau then $s \in \|\Delta^*(\phi)\|\mathcal{V}$.*

PROOF: Suppose τ is a successful tableau for $s \vdash_\Delta \phi$, and suppose for a contradiction that $s \vdash_\Delta \phi$ is false—i.e. that $s \notin \|\Delta^*(\phi)\|\mathcal{V}$. Using only false nodes We trace an infinite $\rightarrow$-derivation unfolding no constant infinitely often, contradicting lemma 8.1. Starting from $s \vdash_\Delta \phi$ pick a path using only false nodes to some false $s_1 \vdash_{\Delta_1} V_1$ s.t. V_1 is introduced as early as possible. If V_1 is a μ-constant then V_1 will eventually fail to occur. If V_1 is a ν-constant then (using ordinal approximations in the standard way) there is some minimal α_1 s.t. $s_1 \notin \|\Delta_1^*(\nu^{\alpha_1} Y_1.\phi_1)\|\mathcal{V}$ where $\Delta_1(V_1) = \nu Y_1.\phi_1$. Consequently along all paths V_1 will eventually become true. In either case we can find a new false $s_2 \vdash_{\Delta_2} V_2$ minimal in the above sense and repeat ad infinitum. $\square$

Lemma 8.3 *If $s \in \|\Delta^*(\phi)\|\mathcal{V}$ then $s \vdash_\Delta \phi$ has a successful tableau.*

PROOF: The proof is related to the corresponding proof in [16]. Suppose $s \in \|\Delta^*(\phi)\|\mathcal{V}$. Let $V_1, \ldots, V_n$ be the sequence of μ-constants of Δ in the order of declaration. The sequence of ordinals $\overline{\alpha}(s \vdash_\Delta \phi) = (\alpha_1, \ldots, \alpha_n, 0, 0, \ldots)$ is the *signature* of $s \vdash_\Delta \phi$, if $\overline{\alpha}(s \vdash_\Delta \phi)$ is lexicographically least s. t. $s \in \|\Delta^*_{\overline{\alpha}(s \vdash_\Delta \phi)}(\phi)\|\mathcal{V}$, where $\Delta_{\overline{\alpha}(s \vdash_\Delta \phi)}$ is Δ with each entry $V_i = \mu Y_i.\phi_i$ changed to $V_i = \mu^{\alpha_i} Y_i.\phi_i$. By always selecting the choice with least signature $s \vdash_\Delta \phi$ can be extended to a partially successful tableau τ. To see it is totally successful suppose there is a path from $s' \vdash_{\Delta'} V$ to $s'' \vdash_{\Delta''} V$ in τ with V a μ-constant. Then $\overline{\alpha}(s' \vdash_{\Delta'} V) > \overline{\alpha}(s'' \vdash_{\Delta''} V)$. Only the introduction of new μ-constants can increase signature from $s' \vdash_{\Delta'} V$ to $s'' \vdash_{\Delta''} V$, but this will not affect the decrease in signature obtained by unfolding V itself. Moreover, the non-zero prefix of $\overline{\alpha}(s' \vdash_{\Delta'} V)$ is of constant length, so no infinitely $>$-decreasing chain can exist. $\square$