A Decidable Intuitionistic Temporal Logic*

Joseph Boudou[1], Martín Diéguez[2], and David Fernández-Duque[3]

1 **Institut de Recherche en Informatique de Toulouse, Toulouse University, Toulouse, France**
2 **Institut de Recherche en Informatique de Toulouse, Toulouse University, Toulouse, France**
3 **Institut de Recherche en Informatique de Toulouse, Toulouse University, Toulouse, France**

Abstract

We introduce the logic $\mathsf{ITL^e}$, an intuitionistic temporal logic based on structures $(W, \preccurlyeq, S)$, where $\preccurlyeq$ is used to interpret intuitionistic implication and S is a $\preccurlyeq$-monotone function used to interpret temporal modalities. Our main result is that the satisfiability and validity problems for $\mathsf{ITL^e}$ are decidable. We prove this by showing that the logic enjoys the strong finite model property. In contrast, we also consider a 'persistent' version of the logic, $\mathsf{ITL^p}$, whose models are similar to Cartesian products. We prove that, unlike $\mathsf{ITL^e}$, $\mathsf{ITL^p}$ does not have the finite model property.

1998 ACM Subject Classification F.4.1 Mathematical Logic

Keywords and phrases intuitionistic logic, temporal logic, products of modal logics

Digital Object Identifier 10.4230/LIPIcs.CSL.2017.14

1 Introduction

Intuitionistic logic [6, 22] and its modal extensions [9, 27, 28] play a crucial role in the area of computer science and artificial intelligence. For instance, Pearce's Equilibrium Logic [26], which characterises the Answer Set semantics [21, 23] of logic programs (ASP), is defined in terms of the intermediate logic of Here and There [15], together with a minimisation criterion. Extensions of Here and There logic allowed the ASP paradigm, already used in a wide range of domains [1, 3, 14, 16, 25], to be applied to reasoning about temporal or epistemic scenarios [5, 10] while satisfying the theorem of strong equivalence [4, 20, 10], central to logic programming and nonmonotonic reasoning.

Such modal extensions of Here and There logic are simple cases of a *modal intuitionistic logic;* in general, the study of such logics can be a challenging enterprise [28]. In particular, there is a huge gap that must be filled regarding combinations of intuitionistic and linear-time temporal logic. Nevertheless, there have been several efforts in this direction, including logics with 'past' and 'future' tenses [9] or with 'next' $\bigcirc$, 'eventually' $\Diamond$ and/or 'henceforth' $\Box$ modalities. The main contributions to the field include the following:

- Davies' intuitionistic temporal logic with $\bigcirc$ [7] was provided Kripke semantics and a complete deductive system by Kojima and Igarashi [18].
- Logics with $\bigcirc, \Box$ were axiomatized by Kamide and Wansing [17], where $\Box$ was interpreted over bounded time.
- Nishimura [24] provided a sound and complete axiomatization for an intuitionistic variant of the propositional dynamic logic PDL.

* This research was partially supported by ANR-11-LABX-0040-CIMI within the program ANR-11-IDEX-0002-02.

26th EACSL Annual Conference on Computer Science Logic (CSL 2017).
Editors: Valentin Goranko and Mads Dam; Article No. 14; pp. 14:1–14:17
Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

- Balbiani and Diéguez [2] axiomatized the Here and There variant of LTL with $\bigcirc, \Diamond, \Box$.
- Davoren [8] introduced topological semantics for temporal logics and Fernández-Duque [11] proved the decidability of a logic with $\bigcirc, \Diamond$ and a universal modality based on topological semantics.

With the exception of [8, 11], semantics for intuitionistic LTL use frames of the form $(W, \preccurlyeq, S)$, where $\preccurlyeq$ is a partial order used to interpret the intuitionistic implication and S is a binary relation used to interpret temporal operators. Since we are interested in linear time, we will restrict our attention to the case where S is a function. Thus, for example, $\bigcirc p$ is true on some world $w \in W$ whenever p is true on $S(w)$. Note, however, that S cannot be an arbitrary function. Intuitionistic semantics have the feature that, for any formula φ and worlds $w \preccurlyeq v \in W$, if φ is true on w then it must also be true on v; that is, truth is *monotone.* If we want this property to be preserved by formulas involving $\bigcirc$, we need for $\preccurlyeq$ and S to satisfy certain confluence properties. In the literature, one generally considers frames satisfying

1. $w \preccurlyeq v$ implies $S(w) \preccurlyeq S(v)$ (*forward confluence,* or simply *confluence*), and
2. if $u \succcurlyeq S(w)$, there is $v \succcurlyeq w$ such that $S(v) = u$ (*backward confluence*).

We will call frames satisfying these conditions *persistent frames* (see Sec. 3), mainly due to the fact that they are closely related to (persistent) products of modal logics [12]. Persistent frames for intuitionistic LTL are the frames of the modal logic $\mathsf{S4} \times \mathsf{LTL}$, which is non-axiomatizable. For this reason, it may not be surprising that it is unknown whether the intuitionistic temporal logic of persistent frames, which we denote $\mathsf{ITL^p}$, is decidable.

However, as we will see in Proposition 1, only forward confluence is needed for truth of all formulas to be monotone, even in the presence of $\Diamond$ and $\Box$. The frames satisfying this condition are, instead, related to *expanding* products of modal logics [13], which are often decidable even when the corresponding product is non-axiomatizable. This suggests that dropping the backwards confluence could also lead to a more manageable intuitionistic temporal logic. This logic, which we denote $\mathsf{ITL^e}$, is the focus of the present paper and, as we will prove in this paper, it enjoys a crucial advantage over $\mathsf{ITL^p}$: $\mathsf{ITL^e}$ has the strong finite model property[1] (hence it is decidable), but $\mathsf{ITL^p}$ does not. In fact, to the best of our knowledge, $\mathsf{ITL^e}$ is the first known decidable intuitionistic temporal logic that

1. is conservative over propositional intuitionistic logic,
2. includes (or can define) the three modalities $\bigcirc, \Diamond, \Box$, and
3. is interpreted over infinite time.

2 Syntax and semantics

We will work in the language $\mathcal{L}$ of LTL given by the following grammar:

$$\varphi, \psi \;\coloneqq\; p \;|\; \bot \;|\; \varphi \wedge \psi \;|\; \varphi \vee \psi \;|\; \varphi \to \psi \;|\; \bigcirc\varphi \;|\; \Diamond\varphi \;|\; \Box\varphi,$$

where p is an element of a countable set of propositional variables $\mathbb{P}$. Given any formula φ, we write $\mathrm{SF}(\varphi)$ for the set of subformulas of φ and $|\varphi|$ for the cardinality of $\mathrm{SF}(\varphi)$.

A *dynamic poset* is a tuple $(W, \preccurlyeq, S)$, where W is a non-empty set of states, $\preccurlyeq$ is a partial order, and S is a function from W to W that satisfies the following *(forward) confluence* condition:

$$\text{for all } w, v \in W, \text{ if } w \preccurlyeq v \text{ then } S(w) \preccurlyeq S(v). \tag{1}$$

[1] That is, there is a computable function bounding the size of the smallest model, if there is one.

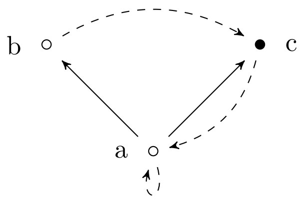

Figure 1 Example of an $\mathsf{ITL^e}$ model $\mathcal{M} = (W, \preccurlyeq, S, V)$, where $\preccurlyeq$ is the reflexive closure of the (already transitive) relation indicated by the solid arrows, S is the relation indicated by the dashed arrows, and a black dot indicates that the variable p is true, so that we only have $p \in V(c)$. Then, the reader may readily verify that $\mathcal{M}, b \models \bigcirc p$ but $\mathcal{M}, b \not\models p$, while $\mathcal{M}, c \models p$ but $\mathcal{M}, c \not\models \bigcirc p$. From this it follows that $\mathcal{M}, a \not\models (\bigcirc p \to p) \vee (p \to \bigcirc p)$.

An *intuitionistic dynamic model,* or simply a *model,* is a tuple $\mathcal{M} = (W, \preccurlyeq, S, V)$ consisting of a dynamic poset equipped with a valuation function V from W to sets of propositional variables satifying the *monotonicity* condition:

$$\text{for all } w, v \in W, \text{ if } w \preccurlyeq v \text{ then } V(w) \subseteq V(v). \tag{2}$$

In the standard way, we define $S^0(w) = w$ and, for all $k > 0$, $S^k(w) = S\left(S^{k-1}(w)\right)$. Then we define the satisfaction relation $\models$ inductively by:

$$\begin{array}{ll}
\mathcal{M}, w \models p & \text{iff } p \in V(w) \\
\mathcal{M}, w \models \bigcirc\varphi & \text{iff } \mathcal{M}, S(w) \models \varphi \\
\mathcal{M}, w \models \bot & \text{never} \\
\mathcal{M}, w \models \Diamond\varphi & \text{iff } \exists k \text{ s.t. } \mathcal{M}, S^k(w) \models \varphi \\
\mathcal{M}, w \models \varphi \wedge \psi & \text{iff } \mathcal{M}, w \models \varphi \text{ and } \mathcal{M}, w \models \psi \\
\mathcal{M}, w \models \Box\varphi & \text{iff } \forall k,\ \mathcal{M}, S^k(w) \models \varphi \\
\mathcal{M}, w \models \varphi \vee \psi & \text{iff } \mathcal{M}, w \models \varphi \text{ or } \mathcal{M}, w \models \psi \\
\mathcal{M}, w \models \varphi \to \psi & \text{iff } \forall v \succcurlyeq w,\ \text{if } \mathcal{M}, v \models \varphi \text{ then } \mathcal{M}, v \models \psi
\end{array}$$

Given a model $\mathcal{M} = (W, \preccurlyeq, S, V)$, a set Σ of formulas, and $w \in W$, we write $\Sigma_{\mathcal{M}}(w)$ for the set $\{\psi \in \Sigma \mid \mathcal{M}, w \models \psi\}$; the subscript '$\mathcal{M}$' is omitted when it is clear from the context. An *eventuality* in $\mathcal{M}$ is a pair (w, φ), where $w \in W$ and φ is a formula such that either $\varphi = \Diamond\psi$ for some formula ψ and $\mathcal{M}, w \models \varphi$, or $\varphi = \Box\psi$ for some formula ψ and $\mathcal{M}, w \not\models \varphi$. The *fulfillment* of an eventuality (w, φ) is the finite sequence $v_0 \dots v_n$ of states of the model such that

1. for all $k \leq n$, $v_0 = S^k(w)$,
2. if $\varphi = \Diamond\psi$ then $\mathcal{M}, v_n \models \psi$ and for all $k < n$, $\mathcal{M}, v_k \not\models \psi$, and
3. if $\varphi = \Box\psi$ then $\mathcal{M}, v_n \not\models \psi$ and for all $k < n$, $\mathcal{M}, v_k \models \psi$.

A formula φ is *satisfiable over a class* Ω *of models* if there is a model $\mathcal{M} \in \Omega$ and a world w so that $\mathcal{M}, w \models \varphi$, and *valid over* Ω if, for every world w of every model $\mathcal{M} \in \Omega$, $\mathcal{M}, w \models \varphi$. Satisfiability (resp. validity) over the class of all intuitionisitic dynamic models is called *satisfiability* (resp. *validity*) for the *expanding domain intuitionisitic temporal logic* $\mathsf{ITL^e}$. We will justify this terminology in the next section. First, we remark that dynamic posets impose the minimal conditions on S and $\preccurlyeq$ in order to preserve the upwards-closure of valuations of formulas. Below, we will use the notation $[\![\varphi]\!] = \{w \in W \mid \mathcal{M}, w \models \varphi\}$.

▶ **Proposition 1.** *Let $\mathcal{D} = (W, \preccurlyeq, S)$, where $(W, \preccurlyeq)$ is a poset and $S \colon W \to W$ is any function. Then, the following are equivalent:*

1. *S satisfies the confluence property* (1)*;*
2. *for every valuation V on W and every formula φ, $[\![\varphi]\!]$ is upwards-closed under $\preccurlyeq$.*

Proof. That 1 implies 2 follows by a standard structural induction on φ. The case where $\varphi \in \mathbb{P}$ follows from the condition on V and most inductive steps are routine. Consider the case where $\varphi = \Box\psi$, and suppose that $w \preccurlyeq v$ and $w \in [\![\varphi]\!]$. Then, for all $i \in \mathbb{N}$, $\mathcal{M}, S^i(w) \models \psi$. Since S is confluent, an easy induction shows that, for all $i \in \mathbb{N}$, $S^i(w) \preccurlyeq S^i(v)$. Therefore, from the induction hypothesis we obtain that $\mathcal{M}, S^i(v) \models \psi$ for all i, hence $v \in [\![\varphi]\!]$. Other cases are similar or easier.

Now we prove that 2 implies 1 by contrapositive. Suppose that $(W, \preccurlyeq, S)$ does *not* satisfy (1), so that there are $w \preccurlyeq v$ such that $S(w) \not\preccurlyeq S(v)$. Choose $p \in \mathbb{P}$ and define $V(u) = \{p\}$ if $S(w) \preccurlyeq u$, $V(u) = \varnothing$ otherwise. It is easy to see that V satisfies the monotonicity condition (2). But, $p \notin V(S(v))$, from which it follows that $(\mathcal{D}, V), w \models \bigcirc p$ but $(\mathcal{D}, V), v \not\models \bigcirc p$. ◀

We are concerned with the satisfiability and validity problems for $\mathsf{ITL}^{\mathsf{e}}$. Observe that satisfiability in propositional intuitionistic logic is equivalent to satisfiability in classical propositional logic. This is because, if φ is classically satisfiable, it is trivially intuitionistically satisfiable in a one-world model; conversely, if φ is intuitionistically satisfiable, it is satisfiable in a finite model, hence in a maximal world of that finite model, and the generated submodel of a maximal world is a classical model. Thus it may be surprising that the same is not the case for intuitionistic temporal logic:

▶ **Lemma 2.** *Any formula φ of the temporal language that is classically satisfiable is satisfiable in a dynamic poset. However, there is a formula satisfiable on a dynamic poset that is not classically satisfiable.*

Proof. If φ is satisfied on a classical LTL model $\mathcal{M}$, then we may regard $\mathcal{M}$ as an intuitionistic model by letting $\preccurlyeq$ be the identity. On the other hand, consider the formula $\neg\bigcirc p \wedge \neg\bigcirc\neg p$ (recall that $\neg\theta$ is a shorthand for $\theta \to \bot$). Classically, this formula is equivalent to $\neg\bigcirc p \wedge \bigcirc p$, and hence unsatisfiable. Define a model $\mathcal{M} = (W, \preccurlyeq, S, V)$, where $W = \{w, v, u\}$, $x \preccurlyeq y$ if $x = y$ or $x = v$, $y = u$, $S(w) = v$ and $S(x) = x$ otherwise, $V(u) = \{p\}$ and $V(v) = V(w) = \varnothing$. Then, one can check that $\mathcal{M}, w \models \neg\bigcirc p \wedge \neg\bigcirc\neg p$. ◀

Hence the decidability of the intuitionistic satisfiability problem is not a corollary of the classical case. In Section 5, we will prove that both the satisfiability and the validity problems are decidable.

3 Expanding and persistent frames

In this section, we discuss expanding and persistent models, and compare them to dynamic models as we have defined above.

3.1 Expanding model property

The logic $\mathsf{ITL}^{\mathsf{e}}$ is closely related to *expanding products* of modal logics [13]. In this subsection, we introduce stratified and expanding frames, and show that satisfiability and validity on arbitrary models is equivalent to satisfiability and validity on expanding models. To do this, it is convenient to represent posets using acyclic graphs.

▶ **Definition 3.** A *directed acyclic graph* is a tuple $(W, \uparrow)$, where W is a set of vertices, $\uparrow \subseteq W \times W$ is a set of edges whose reflexive, transitive closure $\uparrow^*$ is antisymmetric. We will tacitly identify $(W, \uparrow)$ with the poset $(W, \uparrow^*)$. A *path* from w_1 to w_2 is a finite sequence $v_0 \ldots v_n \in W$ such that $v_0 = w_1$, $v_n = w_2$ and for all $k < n$, $v_k \uparrow v_{k+1}$. A *tree* is an acyclic graph $(W, \uparrow)$ with an element $r \in W$, called the root, such that for all $w \in W$ there is a unique path from r to w. A poset $(W, \preccurlyeq)$ is also a tree if there is a relation $\uparrow$ on $W \times W$ such that $(W, \uparrow)$ is a tree and $\preccurlyeq = \uparrow^*$.

▶ **Definition 4.** A model $\mathcal{M} = (W, \preccurlyeq, S, V)$ is *stratified* if there is a partition $\{W_n\}_{n<\omega}$ of W such that

1. each W_n is closed under $\preccurlyeq$,
2. for all n, there is relation $\uparrow_n$ such that $(W_n, \preccurlyeq\restriction_{W_n})$ is a tree, and
3. if $w \in W_n$ then $S(w) \in W_{n+1}$.

If $\mathcal{M}$ is stratified, we write $\preccurlyeq_n, S_n$, and V_n instead of $\preccurlyeq\restriction_{W_n}, S\restriction_{W_n}$, and $V\restriction_{W_n}$ and write $\mathcal{M}_n = (W_n, \preccurlyeq_n, V_n)$. If moreover we have that $S(w) \preccurlyeq S(v)$ implies $w \preccurlyeq v$, then we say that $\mathcal{M}$ is an *expanding model.*

Given a finite, non-empty set of formulas Σ closed under subformulas, a model $\mathcal{M} = (W, \preccurlyeq, S, V)$, and a state $w \in W$, we will construct a stratified model $\mathcal{M}^{\mathrm{e}} = (W^{\mathrm{e}}, \preccurlyeq^{\mathrm{e}}, S^{\mathrm{e}}, V^{\mathrm{e}})$ such that for the root w^{e} of W_0^{e}, $\Sigma(w^{\mathrm{e}}) = \Sigma(w)$. To this end, we first define the set $\mathcal{D} = \mathbb{N} \times \mathbb{N} \times 2^{\Sigma}$ of possible defects. Since Σ is finite and not empty, we assume that $\mathcal{D}$ is ordered such that for each $k \in \mathbb{N}$, the k^{th} element (x, y, H) of $\mathcal{D}$ is such that $x \leq k$. Then, for each $k \in \mathbb{N}$, we construct inductively a tuple $(U_k, \uparrow_k, h_k)$ where $U_k \subseteq \mathbb{N}\times\mathbb{N}$, $\uparrow_k \subseteq U_k \times U_k$ and $h_k : U_k \longrightarrow W$. The model $\mathcal{M}^{\mathrm{e}}$ is defined from all these tuples and the whole construction proceeds as follows:

Base case. Let $U_0 = \{0\} \times \mathbb{N}$, $\uparrow_0 = \varnothing$ and h_0 be such that for all $(0, y) \in U_0$, $h_0(0, y) = S^y(w)$.

Inductive case. Let $k > 0$ and suppose that $(U_k, \uparrow_k, h_k)$ has already been constructed. Let (x, y, H) be the k^{th} element of $\mathcal{D}$. If

(D1) $(x, y) \in U_k$,
(D2) $\Sigma(h_k(x, y)) \neq H$, and
(D3) there is $v \in W$ such that $h_k(x, y) \preccurlyeq v$ and $\Sigma(v) = H$,

then we construct $(U_{k+1}, \uparrow_{k+1}, h_{k+1})$ such that:

$$U_{k+1} = U_k \cup \{(c, d) \in \mathbb{N} \times \mathbb{N} \mid c = k + 1 \text{ and } d \geq y\}$$
$$\uparrow_{k+1} = \uparrow_k \cup \{((a, b), (c, d)) \mid a = x, c = k + 1, d \geq y \text{ and } b = d\}$$
$$h_{k+1} = h_k \cup \left\{((c, d), w) \mid c = k + 1, d \geq y \text{ and } w = S^{d-y}(v)\right\}$$

Otherwise $(U_{k+1}, \uparrow_{k+1}, h_{k+1}) = (U_k, \uparrow_k, h_k)$.

Final step. We construct $\mathcal{M}^{\mathrm{e}} = (W^{\mathrm{e}}, \preccurlyeq^{\mathrm{e}}, S^{\mathrm{e}}, V^{\mathrm{e}})$ such that $W^{\mathrm{e}} = \bigcup_{k\in\mathbb{N}} U_k$, $\preccurlyeq^{\mathrm{e}} = (\uparrow^{\mathrm{e}})^*$, where $\uparrow^{\mathrm{e}} = \bigcup_{k\in\mathbb{N}} \uparrow_k$, $S^{\mathrm{e}} = \{((a, b), (c, d)) \in W^{\mathrm{e}} \times W^{\mathrm{e}} \mid a = c \text{ and } d = b + 1\}$, and $V^{\mathrm{e}}(x, y) = V(h_x(x, y))$.

▶ **Lemma 5.** *For all $(x, y), (x', y') \in W^{\mathrm{e}}$, if $(x, y) \preccurlyeq^{\mathrm{e}} (x', y')$, then $x \leq x'$, $y = y'$ and $h_x(x, y) \preccurlyeq h_{x'}(x', y')$.*

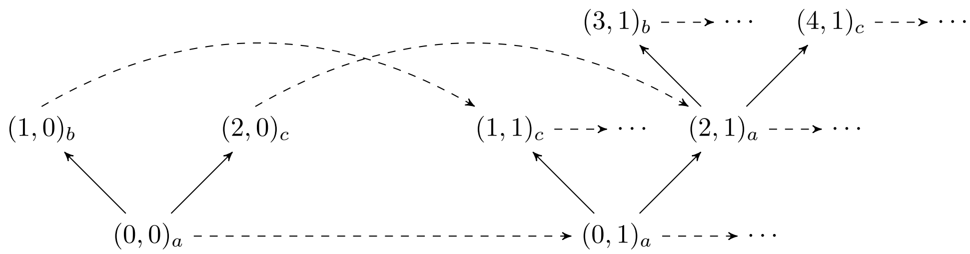

Figure 2 The strata W_0^{e}, W_1^{e} obtained from the model $\mathcal{M}$ defined in Figure 1. The subindexes indicate the value of $h = \bigcup_{k\in\mathbb{N}} h_k$.

Proof. Suppose that $(x,y) \preccurlyeq^{\mathrm{e}} (x',y')$. There is a sequence $(x_0,y_0)\dots(x_n,y_n)$ such that $(x_0,y_0) = (x,y)$, $(x_n,y_n) = (x',y')$ and for all $i < n$, $(x_i,y_i) \uparrow^{\mathrm{e}} (x_{i+1},y_{i+1})$. By construction, for all $i < n$, $(x_i,y_i) \uparrow_{x_{i+1}} (x_{i+1},y_{i+1})$ and $y_i = y_{i+1}$. Let $(x''_{i+1}, y''_{i+1}, H''_{i+1})$ be the $(x_{i+1}-1)^{\text{th}}$ element of $\mathcal{D}$ and v''_{i+1} the element of W choosen at the $(x_{i+1}-1)^{\text{th}}$ step. By construction, $x''_{i+1} = x_i$, $y''_{i+1} \le y_i$ and by the ordering of $\mathcal{D}$, $x_i \le x_{i+1}-1$. Moreover, $h_{x_i}(x_i, y''_{i+1}) \preccurlyeq v''_{i+1}$. Since $h_{x_i}(x_i,y_i) = S^{y_i - y''_{i+1}}\left(h_{x_i}\left(x_i, y''_{i+1}\right)\right)$ and $h_{x_{i+1}}(x_{i+1},y_{i+1}) = S^{y_{i+1}-y''_{i+1}}\left(v''_{i+1}\right)$, by the confluence condition for $\mathcal{M}$, $h_{x_i}(x_i,y_i) \preccurlyeq h_{x_{i+1}}(x_{i+1},y_{i+1})$. ◀

▶ **Lemma 6.** *$\mathcal{M}^{\mathrm{e}}$ is an expanding model.*

Proof. First we check that $\mathcal{M}^{\mathrm{e}}$ is a model. By Lemma 5, $\preccurlyeq^{\mathrm{e}}$ is antisymetric, hence a partial order. For the monotonicity condition, suppose that $(x,y) \preccurlyeq^{\mathrm{e}} (x',y')$. By Lemma 5, $h_x(x,y) \preccurlyeq h_{x'}(x',y')$ and by the monotonicity condition for $\mathcal{M}$, $V\left(h_x(x,y)\right) \subseteq V\left(h_{x'}(x',y')\right)$. For the confluence condition, it suffices to observe that by construction, if $(x,y) \uparrow^{\mathrm{e}} (x',y')$ then $(x,y+1) \uparrow^{\mathrm{e}} (x',y'+1)$. Therefore, $\mathcal{M}^{\mathrm{e}}$ is a model. To prove that $\mathcal{M}^{\mathrm{e}}$ is stratified, define $W_n^{\mathrm{e}} = \{(x,y) \in W^{\mathrm{e}} \mid y = n\}$ for all $n \in \mathbb{N}$. Conditions 3 of Def. 4 trivially holds and condition 1 comes directly from Lemma 5. To prove condition 2, it suffices to observe that by construction, for all $(x,y) \in W^{\mathrm{e}}$, either $x = 0$ or there is exactly one state $(x',y') \in W^{\mathrm{e}}$ such that $(x',y') \uparrow^{\mathrm{e}} (x,y)$. Therefore, by Lemma 5, for all $(x,y) \in W^{\mathrm{e}}$, there is a unique path from $(0,y)$ to (x,y). Finally, to prove that $\mathcal{M}^{\mathrm{e}}$ is expanding, suppose that $(c,b) \in W^{\mathrm{e}}$ and $(a,b+1) \uparrow^{\mathrm{e}} (c,b+1)$. Then the $(c-1)^{\text{th}}$ element of $\mathcal{D}$ is (a,y,H) for some y,H. Moreover, since $(c,b) \in W^{\mathrm{e}}$, $b \ge y$ and since $(a,y) \in W^{\mathrm{e}}$, $(a,b) \in W^{\mathrm{e}}$ and $(a,b) \uparrow^{\mathrm{e}} (c,b)$. Therefore it can easily be proved by induction on the length of the path from $S^{\mathrm{e}}(w)$ to $S^{\mathrm{e}}(v)$ that $S^{\mathrm{e}}(w) \preccurlyeq^{\mathrm{e}} S^{\mathrm{e}}(v)$ implies $w \preccurlyeq^{\mathrm{e}} v$. ◀

▶ **Lemma 7.** *For any state $(x,y) \in W^{\mathrm{e}}$ and any $\psi \in \Sigma$, $\mathcal{M}^{\mathrm{e}},(x,y) \models \psi$ if and only if $\mathcal{M}, h_x(x,y) \models \psi$.*

Proof. The proof is by induction on the size $|\psi|$ of the formula. The cases for propositional variables, falsum, conjunctions and disjunctions are straightforward. For the temporal modalities, it suffices to observe that for all $(x,y) \in W^{\mathrm{e}}$ and all $n \in \mathbb{N}$, $(x,y+n) \in W^{\mathrm{e}}$ and $h_x(x,y+n) = S^n\left(h_x(x,y)\right)$. Finally, for implication, suppose first that $\mathcal{M}^{\mathrm{e}},(x,y) \not\models \psi_1 \to \psi_2$. Then there is (x',y') such that $(x,y) \preccurlyeq^{\mathrm{e}} (x',y')$, $\mathcal{M}^{\mathrm{e}},(x',y') \models \psi_1$ and $\mathcal{M}^{\mathrm{e}},(x',y') \not\models \psi_2$. By Lemma 5, $h_x(x,y) \preccurlyeq h_{x'}(x',y')$ and by induction hypothesis, $\mathcal{M}, h_{x'}(x',y') \models \psi_1$ and $\mathcal{M}, h_{x'}(x',y') \not\models \psi_2$. Therefore, $\mathcal{M}, h_x(x,y) \not\models \psi_1 \to \psi_2$. For the other direction suppose that $\mathcal{M}, h_x(x,y) \not\models \psi_1 \to \psi_2$. There is $v' \in W$ such that $h_x(x,y) \preccurlyeq v'$, $\mathcal{M}, v' \models \psi_1$ and $\mathcal{M}, v' \not\models \psi_2$. Let k be such that $(x,y,\Sigma(v'))$ is the k^{th} element of $\mathcal{D}$. Condition (D3) trivially holds and since $x \le k$, condition (D1) holds too. Hence, there is $(x',y') \in W^{\mathrm{e}}$ such that

$\Sigma(h_{x'}(x', y')) = \Sigma(v')$ and either $(x', y') = (x, y)$ or $(x, y) \uparrow^{\mathrm{e}} (x', y')$. By induction hypothesis, $\mathcal{M}^{\mathrm{e}}, (x', y') \vDash \psi_1$ and $\mathcal{M}^{\mathrm{e}}, (x', y') \nvDash \psi_2$, hence $\mathcal{M}^{\mathrm{e}}, (x, y) \nvDash \psi_1 \to \psi_2$. ◀

In conclusion, we obtain the following:

▶ **Theorem 8.** *A formula φ is satisfiable (resp. falsifiable) on an intuitionistic dynamic model if and only if it is satisfiable (resp. falsifiable) on an expanding model.*

3.2 Persistent frames

Expanding models were introduced as a weakening of product models. They often lead to logics with a less complex validity problem. Thus it is natural to also consider a variant of $\mathsf{ITL}^{\mathsf{e}}$ interpreted over product models, or over the somewhat wider class of persistent models.

▶ **Definition 9.** Let $(W, \preccurlyeq)$ be a poset. If $S \colon W \to W$ is such that, whenever $v \succcurlyeq S(w)$, there is $u \succcurlyeq w$ such that $v = S(u)$, we say that S is *backward confluent*. If S is both forward and backward confluent, we say that it is *persistent*. A tuple $(W, \preccurlyeq, S)$ where S is persistent is a *persistent intuitionistic temporal frame*, and the set of valid formulas over the class of persistent intuitionistic temporal frames is denoted $\mathsf{ITL}^{\mathsf{p}}$, or *persistent domain* LTL.

The name 'persistent' comes from the fact that Theorem 8 can be modified to obtain a stratified model $\mathcal{M}'$ where $S' \colon W'_k \to W'_{k+1}$ is an isomorphism, i.e. whose domains are persistent with respect to S'. As we will see, the finite model property fails over the class of persistent models.

▶ **Lemma 10.** *The formula $\varphi = \neg\neg\Diamond\Box p \to \Diamond\neg\neg\Box p$ is not valid over the class of persistent models.*

Proof. Consider the model $\mathcal{M} = (W, \preccurlyeq, S, V)$, where $W = \mathbb{Z} \cup \{r\}$ with r a fresh world not in $\mathbb{Z}$, $w \preccurlyeq v$ if and only if $w = r$ or $w = v$, $S(r) = r$ and $S(n) = n + 1$ for $n \in \mathbb{Z}$, and $[\![p]\!] = [0, \infty)$. It is readily seen that $\mathcal{M}$ is a persistent model, that $\mathcal{M}, r \vDash \neg\neg\Diamond\Box p$ (since every maximal world above r satisfies $\Diamond\Box p$), yet $\mathcal{M}, r \nvDash \Diamond\neg\neg\Box p$, since there is no n such that $\mathcal{M}, S^n(r) \vDash \neg\neg\Box p$. It follows that $\mathcal{M}, r \nvDash \varphi$, and hence φ is not valid, as claimed. ◀

▶ **Lemma 11.** *The formula φ (from Lemma 10) is valid over the class of finite, persistent models.*

Proof. Let $\mathcal{M} = (W, \preccurlyeq, S, V)$ be a finite, persistent model, and assume that $\mathcal{M}, w \vDash \neg\neg\Diamond\Box p$. Let $v_1, \ldots, v_n$ enumerate the maximal elements of $\{v \in W \mid w \preccurlyeq v\}$. For each $i \leq n$, let k_i be large enough so that $\mathcal{M}, S^{k_i}(v_i) \vDash \Box p$, and let $k = \max k_i$. We claim that $\mathcal{M}, S^k(w) \vDash \neg\neg\Box p$, which concludes the proof. Let $u \succcurlyeq S^k(w)$ be any leaf. Then, there is $v_i \succcurlyeq w$ such that $u = S^k(v_i)$ (since compositions of persistent functions are persistent). But, since $k \geq k_i$, we obtain $\mathcal{M}, u \vDash \Box p$, as desired. ◀

The following is then immediate from Lemmas 10 and 11:

▶ **Theorem 12.** $\mathsf{ITL}^{\mathsf{p}}$ *does not have the finite model property.*

Thus our decidability proof for $\mathsf{ITL}^{\mathsf{e}}$, which proceeds by first establishing a strong finite model property, does not carry over to $\mathsf{ITL}^{\mathsf{p}}$. Whether $\mathsf{ITL}^{\mathsf{p}}$ is decidable remains open.

4 Combinatorics of intuitionistic models

In this section we introduce some combinatorial tools we will need in order to prove that $\mathsf{ITL}^{\mathsf{e}}$ has the strong finite model property, and hence is decidable. We begin by discussing labeled structures, which allow for a graph-theoretic approach to intuitionistic models.

4.1 Labeled structures and quasimodels

▶ **Definition 13.** Given a set Λ whose elements we call 'labels' and a set W, a Λ-*labeling function* on W is any function $\lambda\colon W \to \Lambda$. A structure $\mathcal{S} = (W, R, \lambda)$ where W is a set, $R \subseteq W \times W$ and λ is a labeling function on W is a Λ-*labeled structure,* where 'structure' may be replaced with 'poset', 'directed graph', etc.

A useful measure of the complexity of a labeled poset or graph is given by its level:

▶ **Definition 14.** Given a labeled poset $\mathcal{A} = (W, \preccurlyeq, \lambda)$ and an element $w \in W$, an *increasing chain* from w of length n is a sequence $v_1 \ldots v_n$ of elements of W such that $v_1 = w$ and $\forall i < n,\ v_i \prec v_{i+1}$, where $u \prec u'$ is shorthand for $u \preccurlyeq u'$ and $u' \not\preccurlyeq u$. The chain $v_1 \ldots v_n$ is *proper* if it moreover satisfies $\forall i < n,\ \lambda(v_i) \neq \lambda(v_{i+1})$. The *depth* $\mathrm{dpt}(w) \in \mathbb{N} \cup \{\omega\}$ of w is defined such that $\mathrm{dpt}(w) = m$ if m is the maximal length of all the increasing chains from w and $\mathrm{dpt}(w) = \omega$ if there is no such maximum. Similarly, the *level* $\mathrm{lvl}(w) \in \mathbb{N} \cup \{\omega\}$ of w is defined such that $\mathrm{lvl}(w) = m$ if m is the maximal length of all the proper increasing chains from w and $\mathrm{lvl}(w) = \omega$ if there is no such maximum. The level $\mathrm{lvl}(\mathcal{A})$ of $\mathcal{A}$ is the maximal level of all its elements.

The notions of *depth* and *level* are extended to any acyclic directed graph $(W, \uparrow, \lambda)$ by taking the respective values on $(W, \uparrow^*, \lambda)$.

An important class of labeled posets comes from intuitionistic models.

▶ **Definition 15.** Given an intuitionistic Kripke model $\mathcal{M} = (W, \preccurlyeq, V)$ and a set $\Sigma \subseteq \mathcal{L}$ closed under subformulas, it can easily be checked that for all $w, v \in W$, if $w \preccurlyeq v$ then $\Sigma(w) \subseteq \Sigma(v)$. We denote the labeled poset $(W, \preccurlyeq, \Sigma(\cdot))$ by $\mathcal{M}^\Sigma$. Conversely, given a labeled poset $\mathcal{A} = (W, \preccurlyeq, \lambda)$ over 2^Σ such that if $w \preccurlyeq v$ then $\lambda(w) \subseteq \lambda(v)$, the valuation V_λ is defined such that $V_\lambda(w) = \{p \in \mathbb{P} \mid p \in \lambda(w)\}$ for all $w \in W$, and denote the resulting model by $\mathcal{A}^{\mathrm{mod}}$.

▶ **Definition 16.** Let Σ be a finite set of formulas closed under subformulas and $\mathcal{A} = (W, \preccurlyeq, \lambda)$ be a 2^Σ-labeled poset. We say that $\mathcal{A}$ is a Σ-*quasimodel* if λ is monotone in the sense that $w \preccurlyeq v$ implies that $\lambda(w) \subseteq \lambda(v)$, and whenever $\varphi \to \psi \in \Sigma$ and $w \in W$, we have that $\varphi \to \psi \in \lambda(w)$ if and only if, for all v such that $w \preccurlyeq v$, if $\varphi \in \lambda(v)$ then $\psi \in \lambda(v)$. If $(W, \preccurlyeq)$ is a tree, we say that $\mathcal{A}$ is *tree-like.*

4.2 Simulations, immersions and condensations

As is well-known, truth in intuitionistic models is preserved by bisimulation, and thus this is usually the appropriate notion of equivalence between different models. However, for our purposes, it is more convenient to consider a weaker notion, which we call *bimersion.*

▶ **Definition 17.** Given two labeled posets $\mathcal{A} = (W_\mathcal{A}, \preccurlyeq_\mathcal{A}, \lambda_\mathcal{A})$ and $\mathcal{B} = (W_\mathcal{B}, \preccurlyeq_\mathcal{B}, \lambda_\mathcal{B})$ and a relation $R \subseteq W_\mathcal{A} \times W_\mathcal{B}$, we write $\mathrm{dom}(R)$ for $\{w \in W_\mathcal{A} \mid \exists v \in W_\mathcal{B},\ (w, v) \in R\}$ and $\mathrm{rng}(R)$ for $\{v \in W_\mathcal{B} \mid \exists w \in W_\mathcal{A},\ (w, v) \in R\}$. A relation $\sigma \subseteq W_\mathcal{A} \times W_\mathcal{B}$ is a *simulation* from $\mathcal{A}$ to $\mathcal{B}$

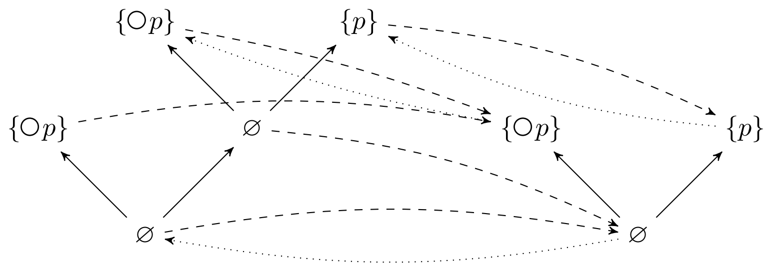

Figure 3 A condensation from the labeled frame on the left to the labeled frame on the right. Dashed arrows indicate ρ, dotted arrows ι.

if $\mathrm{dom}(\sigma) = W_{\mathcal{A}}$ and whenever $w\ \sigma\ v$, it follows that $\lambda_{\mathcal{A}}(w) = \lambda_{\mathcal{B}}(v)$, and if $w \preccurlyeq_{\mathcal{A}} w'$ then there is v' so that $v \preccurlyeq_{\mathcal{B}} v'$ and $w'\ \sigma\ v'$.

A simulation is called an *immersion* if it is a function. If an immersion $\sigma\colon W_{\mathcal{A}} \to W_{\mathcal{B}}$ exists, we write $\mathcal{A} \trianglelefteq \mathcal{B}$. If, moreover, there is an immersion $\tau\colon W_{\mathcal{B}} \to W_{\mathcal{A}}$, we say that they are *bimersive*, write $\mathcal{A} \triangleq \mathcal{B}$, and call the pair (σ, τ) a *bimersion*. A *condensation from $\mathcal{A}$ to $\mathcal{B}$* is a bimersion (ρ, ι) so that $\rho\colon W_{\mathcal{A}} \to W_{\mathcal{B}}$, $\iota\colon W_{\mathcal{B}} \to W_{\mathcal{A}}$, ρ is surjective, and $\rho\iota$ is the identity on $W_{\mathcal{B}}$. If such a condensation exists we write $\mathcal{B} \ll \mathcal{A}$. Observe that $\mathcal{B} \ll \mathcal{A}$ implies that $\mathcal{B} \triangleq \mathcal{A}$.

If $\mathcal{M}, \mathcal{N}$ are models and Σ a set of formulas closed under subformulas, we write $\mathcal{M} \trianglelefteq_{\Sigma} \mathcal{N}$ if $\mathcal{M}^{\Sigma} \trianglelefteq \mathcal{N}^{\Sigma}$, and define $\triangleq_{\Sigma}, \ll_{\Sigma}$ similarly. We may also write e.g. $\mathcal{A} \ll \mathcal{M}$ if $\mathcal{A}$ is 2^{Σ}-labeled and $\mathcal{A} \ll \mathcal{M}^{\Sigma}$.

In this text, simulations will *always* be between posets, and if instead $\mathcal{A}$ or $\mathcal{B}$ is an acyclic directed graph, a simulation between $\mathcal{A}$ and $\mathcal{B}$ will be one between their respective transitive closures. It will typically be convenient to work with immersions rather than simulations: however, as the next lemma shows, not much generality is lost by this restriction.

▶ **Lemma 18.** *Let $\mathcal{A} = (W_{\mathcal{A}}, \preccurlyeq_{\mathcal{A}}, \lambda_{\mathcal{A}})$ and $\mathcal{B} = (W_{\mathcal{B}}, \preccurlyeq_{\mathcal{B}}, \lambda_{\mathcal{B}})$ be labeled posets. If a simulation $\sigma \subseteq W_{\mathcal{A}} \times W_{\mathcal{B}}$ exists, $W_{\mathcal{A}}$ is a finite tree, and $w\ \sigma\ w'$, then there is an immersion $\sigma' \subseteq W_{\mathcal{A}} \times W_{\mathcal{B}}$ such that $w' = \sigma'(w)$.*

Proof. By a straightforward induction on the depth of w we show that there is a partial immersion σ_w with $w \in \mathrm{dom}(\sigma_w)$, whose domain is the subtree generated by w, and such that $w\ \sigma\ w'$. Let D be the set of childrenset of daughters of w, and for each $v \in D$, choose v' so that $v\ \sigma\ v'$ and $w' \preccurlyeq_{\mathcal{B}} v'$. By the induction hypothesis, there is a partial immersion σ'_v with $v \in \mathrm{dom}(\sigma'_v)$. Then, one readily checks that $\{(w, w')\} \cup \bigcup_{v \in D} \sigma'_v$ is also an immersion, as needed. ◀

Condensations are useful for producing (small) quasimodels out of models.

▶ **Proposition 19.** *Given an intuitionistic model $\mathcal{M} = (W_{\mathcal{M}}, \preccurlyeq_{\mathcal{M}}, V_{\mathcal{M}})$, a set Σ of intuitionistic formulas that is closed for subformulas, and a 2^{Σ}-labeled poset $\mathcal{A} = (W_{\mathcal{A}}, \preccurlyeq_{\mathcal{A}}, \lambda_{\mathcal{A}})$ over Σ, if $\mathcal{A} \ll \mathcal{M}$, then $\mathcal{A}$ is a quasimodel.*

Proof. See Appendix. ◀

4.3 Normalized labeled trees

In order to count the number of different labeled trees up to bimersion, we construct, for any set Λ of labels and any $k \geq 1$, the labeled directed acyclic graph $\mathcal{G}_k^{\Lambda} = \left(W_k^{\Lambda}, \uparrow_k^{\Lambda}, \lambda_k^{\Lambda}\right)$ by induction on k as follows.

Base case. For $k = 1$, let $\mathcal{G}_1^\Lambda = \left(W_1^\Lambda, \uparrow_1^\Lambda, \lambda_1^\Lambda\right)$ with $W_1^\Lambda = \Lambda W_1^\Lambda = L$, $\uparrow_1^\Lambda = \varnothing$, and $\lambda_1^\Lambda(w) = w$ for all $w \in W_1^\Lambda$.

Inductive case. Suppose that $\mathcal{G}_k^\Lambda = \left(W_k^\Lambda, \uparrow_k^\Lambda, \lambda_k^\Lambda\right)$ has already been defined. The graph $\mathcal{G}_{k+1}^\Lambda = \left(W_{k+1}^\Lambda, \uparrow_{k+1}^\Lambda, \lambda_{k+1}^\Lambda\right)$ is constructed such that:

$$W_{k+1}^\Lambda = W_k^\Lambda \cup P$$
$$\uparrow_{k+1}^\Lambda = \uparrow_k^\Lambda \cup \left\{(x,y) \in W_{k+1}^\Lambda \times W_{k+1}^\Lambda \mid \exists(\ell, C) \in P,\ x = (\ell, C) \text{ and } y \in C\right\}$$
$$\lambda_{k+1}^\Lambda(w) = \begin{cases} \lambda_k^\Lambda(w) & \text{if } w \in W_k^\Lambda \\ \ell & \text{if } w = (\ell, C) \in P \end{cases}$$

where $P = \left\{(\ell, C) \in \Lambda \times \mathcal{P}\left(W_k^\Lambda\right) \mid \forall y \in C,\ \lambda_k^\Lambda(y) \neq \ell\right\}$.

Note that $\mathcal{G}_k^\Lambda = \left(W_k^\Lambda, \uparrow_k^\Lambda, \lambda_k^\Lambda\right)$ is typically not a tree, but we may unravel it to obtain one.

▶ **Definition 20.** Given a labeled directed acyclic graph $\mathcal{G} = (W, \uparrow, \lambda)$ and a node $w \in W$, the *unraveling* of $\mathcal{G}$ from w is the labeled tree $\mathcal{T}_w = (W_w, \uparrow_w, \lambda_w)$ such that W_w is the set of all the paths from w in $\mathcal{G}$, $\xi \uparrow_w \zeta$ if and only if there is $v \in W$ such that $\zeta = \xi v$, and $\lambda_w(v_0 \ldots v_n) = \lambda(v_n)$.

▶ **Proposition 21.** *For any rooted labeled tree $\mathcal{T} = (W, \uparrow, \lambda)$ over a set Λ of labels, if the level of $\mathcal{T}$ is finite then there is a condensation from $\mathcal{T}$ to an unraveling of $\mathcal{G}_{\mathrm{lvl}(\mathcal{T})}^\Lambda = \left(W_{\mathrm{lvl}(\mathcal{T})}^\Lambda, \uparrow_{\mathrm{lvl}(\mathcal{T})}^\Lambda, \lambda_{\mathrm{lvl}(\mathcal{T})}^\Lambda\right)$.*

Proof. Let $\mathcal{T} = (W, \uparrow, \lambda)$ be a labeled treedirected acyclic graph with root r. We write $\prec$ for the transitive closure of $\uparrow$ and $\preccurlyeq$ for the reflexive closure of $\prec$. The proof is by induction on the level $n = \mathrm{lvl}(\mathcal{T})$ of $\mathcal{T}$. For $n = 1$, observe that this means that $\lambda(w) = \lambda(r)$ for all $w \in W$. Let $\rho = W \times \{\lambda(r)\}$ and $\iota = \{(\lambda(r), r)\}$. It can easily be checked that (ρ, ι) is a condensation.[2] For $n > 1$, suppose the property holds for all rooted labeled trees $\mathcal{T}'$ such that $\mathrm{lvl}(\mathcal{T})' < n$. Define the following sets:

$$N = \{w \in W \mid \lambda(w) \neq \lambda(r) \text{ and for all } v \prec w,\ \lambda(v) = \lambda(r)\}$$
$$M = \{w \in W \mid \text{for all } v \preccurlyeq w,\ \lambda(v) = \lambda(r)\}$$

Clearly, for all $w \in N$, $\mathrm{lvl}(w) < n$. Therefore, by induction, there is a condensation (ρ_w, ι_w) from the subgraph of $\mathcal{T}$ generated by w to the unraveling of $\mathcal{G}_{n-1}^\Lambda$ from some $y_w \in W_{n-1}^\Lambda$. Let us define $r' = (\lambda(r), \{y_w \mid w \in N\})$ and consider the unraveling $\mathcal{G}$ of $\mathcal{G}_n^\Lambda$ from r'. It can easily be checked that $\rho = (M \times \{r'\}) \cup \bigcup_{w \in N} \rho_w \rho = (S \times \{r'\}) \cup \bigcup_{w \in W} \rho_w$ is an immersion from $\mathcal{T}$ to $\mathcal{G}$, $\iota' = \{(r', r)\} \cup \bigcup_{w \in N} \iota_w \iota' = \{(r', r)\} \cup \bigcup_{w \in W} \iota_w$ is a simulation from $\mathcal{G}$ to $\mathcal{T}$ and $\iota' \subseteq \rho^{-1}$. Using Lemma 18, we can then choose an immersion $\iota \subseteq \iota'$, so that (ρ, ι) is a condensation from $\mathcal{T}$ to $\mathcal{G}$. ◀

Finally, let us define recursively E_k^n and Q_k^n for all $n, k \in \mathbb{N}$ by:

$$E_k^n = \begin{cases} 0 & \text{if } k = 0 \\ E_{k-1}^n + n2^{E_{k-1}^n} & \text{otherwise} \end{cases} \qquad Q_k^n = \begin{cases} 0 & \text{if } k = 0 \\ 1 + E_{k-1}^n Q_{k-1}^n & \text{otherwise} \end{cases}$$

[2] Recall that as per our convention, this means that (ρ, ι) is a condensation between the respective transitive closures.

The following lemma can be proven by a straightforward induction, left to the reader.

▶ **Lemma 22.** *For any finite set Λ with cardinality n and all $k \in \mathbb{N}$,*

1. *the cardinality of $\mathcal{G}_k^\Lambda$ is bounded by E_k^n, and*
2. *the cardinality of any unraveling of $\mathcal{G}_k^\Lambda$ is bounded by Q_k^n.*

From these and Proposition 21, we obtain the following:

▶ **Theorem 23.**

1. *Given a set of labels Λ and a Λ-labeled tree $\mathcal{T}$ of level $k < \omega$, there is a Λ-labeled tree $\mathcal{T}'$ bounded by $Q_k^{|\Lambda|}$ such that $\mathcal{T}' \triangleq \mathcal{T}$. We call $\mathcal{T}'$ the* normalized Λ-labeled tree *for $\mathcal{T}$.*
2. *Given a sequence of Λ-labeled trees $\mathcal{T}_1, \ldots, \mathcal{T}_n$ of level $k < \omega$ with $n > E_k^{|\Lambda|}$, there are indexes $i < j \leq n$ such that $\mathcal{T}_i \triangleq \mathcal{T}_j$.*

The second item may be viewed as a finitary variant of Kruskal's theorem for labeled trees [19]. When applied to quasimodels, we obtain the following:

▶ **Proposition 24.** *Let Σ be a set of formulas closed under subformulas with $|\Sigma| = s < \omega$.*

1. *Given a tree-like Σ-quasimodel $\mathcal{T}$ and $s = |\Sigma|$, there is a tree-like Σ-quasimodel $\mathcal{T}' \triangleq_\Sigma \mathcal{T}$ bounded by $Q_{s+1}^{2^s}$. We call $\mathcal{T}'$ the* normalized Σ-quasimodel *for $\mathcal{T}$.*
2. *Given a sequence of tree-like Σ-quasimodels $\mathcal{T}_1, \ldots, \mathcal{T}_n$ with $n > E_{s+1}^{2^s}$, there are indexes $i < j \leq n$ such that $\mathcal{T}_i \triangleq \mathcal{T}_j$.*

Proof. Immediate from Proposition 19 and Theorem 23Lemma 22 using the fact that any Σ-quasimodel has level at most $s + 1$. ◀

Finally, we obtain an analogous result for pointed structures.

▶ **Definition 25.** A *pointed labeled poset* is a structure $(W, \preccurlyeq, \lambda, w)$ consisting of a labeled tree with a designated world $w \in W$. Given a labeled poset $\mathcal{A} = (W_\mathcal{A}, \preccurlyeq_\mathcal{A}, \lambda_\mathcal{A})$ and $w \in W_\mathcal{A}$, we denote by $\mathcal{A}^w$ the pointed labeled poset given by $\mathcal{A}^w = (W_\mathcal{A}, \preccurlyeq_\mathcal{A}, \lambda_\mathcal{A}, w)$. A *pointed simulation* between pointed labeled posets $\mathcal{A} = (W_\mathcal{A}, \preccurlyeq_\mathcal{A}, \lambda_\mathcal{A}, w_\mathcal{A})$ and $\mathcal{B} = (W_\mathcal{B}, \preccurlyeq_\mathcal{B}, \lambda_\mathcal{B}, w_\mathcal{B})$ is a simulation $\sigma \subseteq W_\mathcal{A} \times W_\mathcal{B}\sigma \subset W_\mathcal{A} \times W_\mathcal{B}$ such that if $w \ \sigma \ v$, then $w = w_\mathcal{A}$ if and only if $v = w_\mathcal{B}$. The notions of *pointed immersion, pointed condensation,* etc. are defined analogously to Definition 17.

▶ **Lemma 26.** *If Λ has n elements, any pointed Λ-labeled poset of level at most k condenses to a labeled pointed tree bounded by Q_{k+2}^{2n}, and there are at most E_{k+2}^{2n} bimersion classes.*

Proof. We may view a pointed labeled poset $\mathcal{A} = (W, \preccurlyeq, \lambda, w)$ as a (non-pointed) labeled poset as follows. Let $\Lambda' = \Lambda \times \{0, 1\}$. Then, set $\lambda'(v) = (\lambda(v), 0)$ if $v \neq w$, $\lambda'(w) = (\lambda(w), 1)$. Note that $\mathcal{A}$ may now have level $k+2$, since we may have that $u \preccurlyeq w \preccurlyeq v$, $\lambda(u) = \lambda(w) = \lambda(v)$, yet $\lambda'(u) \neq \lambda'(w)$ and $\lambda'(w) \neq \lambda'(v)$. By Proposition 21, $\mathcal{A}$ condenses to a generated tree $\mathcal{T}$ of $\mathcal{G}_{k+2}^{\Lambda'}$ by some condensation (ρ, ι). Let $w' = \rho(w)$, and consider $\mathcal{T}$ as a pointed structure with distinguished point w'. Given that ρ is a surjective, label-preserving function, w, w' are the only points whose label has second component 1, and therefore (ρ, ι) must be a pointed condensation, as claimed. ◀

▶ **Proposition 27.** *Let Σ be a set of formulas closed under subformulas with $|\Sigma| = s < \omega$.*

1. *Given a tree-like pointed Σ-quasimodel $\mathcal{T}$ and a formula φ, there is a tree-like pointed Σ-quasimodel $\mathcal{T}' \triangleq \mathcal{T}$ bounded by $Q_{s+3}^{2^{s+1}}$. We call $\mathcal{T}'$ the* normalized pointed Σ-quasimodel *for $\mathcal{T}$.*

2. *Given a sequence of tree-like pointed Σ-quasimodels $\mathcal{T}_1, \ldots, \mathcal{T}_n$ with $n > E_{s+3}^{2^{s+1}}$, there are indexes $i < j \leq n$ such that $\mathcal{T}_i \triangleq \mathcal{T}_j$.*

With these tools at hand, we are ready to prove that $\mathsf{ITL^e}$ has the strong finite model property, and hence is decidable.

5 Decidability

The following transformations are defined for any stratified model $\mathcal{M} = (W, \preccurlyeq, S, V)$ and any finite, non-empty set of formulas Σ closed under subformulas. In each case, given a stratified model $\mathcal{M} = (W, \preccurlyeq, S, V)$, we will produce another stratified model $\mathcal{M}' = (W', \preccurlyeq', S', V')$ and a map $\pi \colon W' \to W$ such that $\Sigma_{\mathcal{M}}(\pi(w)) = \Sigma_{\mathcal{M}'}(w)$ for all $w \in W'\Sigma_{\mathcal{M}}(w) = \Sigma_{\mathcal{M}'}(\pi(w))$ for all $w \in W'$.

Replace $\mathcal{M}_k$ with a copy of the normalized Σ-quasimodel of $\mathcal{M}_k$. Let $\mathcal{T} = (W_{\mathcal{T}}, \uparrow_{\mathcal{T}}, \lambda_{\mathcal{T}})$ be a copy of the normalized labeled tree of $\mathcal{M}_k^{\Sigma}$ such that $W_{\mathcal{T}} \cap W = \varnothing$, and (ρ, ι) the condensation from $\mathcal{M}_k^{\Sigma}$ to $\mathcal{T}$. The result of the transformation is the tuple $(W', \preccurlyeq', S', V')$ such that $W' = W \cup W_{\mathcal{T}} \setminus W_k$, $\preccurlyeq' = \preccurlyeq\restriction_{W \setminus W_k} \cup (\uparrow_{\mathcal{T}})^*$,

$$S'(w) = \begin{cases} \rho\left(S\left(w\right)\right) & \text{if } w \in W_{k-1} \\ S\left(\iota\left(w\right)\right) & \text{if } w \in W_{\mathcal{T}} \\ S(w) & \text{otherwise} \end{cases} \qquad V'(w) = \begin{cases} \{p \mid p \in \lambda_{\mathcal{T}}(w)\} & \text{if } w \in W_{\mathcal{T}} \\ V(w) & \text{otherwise} \end{cases}$$

The map π is the identity on $W'_i = W_i$ for $i \neq k$, and $\pi(w) = \iota(w)$ for $w \in W_{\mathcal{T}}$.

Replace $\mathcal{M}_k$ with a copy of the normalized, pointed Σ-quasimodel of $\mathcal{M}_k$ preserving w, where $w \in W_k$. The transformation is similar to the previous one except that $\mathcal{M}_k$ is regarded as a pointed structure with distinguished point w.

Replace $\mathcal{M}_\ell$ with $\mathcal{M}_k$, where $k < \ell$ and there is an immersion $\sigma \colon W_k \to W_\ell$ (seen as 2^{Σ}-labeled trees). The result of the transformation is the tuple $(W', \preccurlyeq', S', V')$ such that $W' = W \setminus \bigcup_{k<m\leq \ell} W_m$, $\preccurlyeq' = \preccurlyeq\restriction_{W'}$,

$$S'(w) = \begin{cases} S\left(\sigma\left(w\right)\right) & \text{if } S(w) \in W_k \\ S(w) & \text{otherwise} \end{cases}$$

and $V' = V\restriction_{W'}$.

The map π is the identity on $W'_i = W_i$ for $i < k$, on $W'_i = W_{i+\ell-k}$ for $i > k$, and $\pi(w) = \sigma(w)$ for all $w \in W'_k$.

Replace $\mathcal{M}_\ell$ with $\mathcal{M}_k$ connecting w_k to w_ℓ, where $k < \ell$, $w_k \in W_k$, $w_\ell \in W_\ell$ and there is an immersion $\sigma \colon W_k \to W_\ell$ such that $\sigma(w_k) = w_\ell$. The transformation is defined as the previous one.

▶ **Lemma 28.** *The result of any previous transformation is a stratified model such that $\Sigma_{\mathcal{M}}(\pi(w)) = \Sigma_{\mathcal{M}'}(w)\Sigma_{\mathcal{M}}(w) = \Sigma_{\mathcal{M}'}(\pi(w))$ for any $w \in W'$.*

Proof. The proof that $\mathcal{M}' = (W', \preccurlyeq', S', V')$ is a stratified modela model is straightforward and left to the reader. We prove by structural induction on φ that for all transformations, all $w \in W'$ and all $\varphi \in \Sigma$, $\mathcal{M}', w \vDash \varphi$ iff $\mathcal{M}, \pi(w) \vDash \varphi$. We only detail the case for the next modality when $\mathcal{M}_k$ is replaced with a copy of the normalized Σ-quasimodel $\mathcal{T}$ of $\mathcal{M}_k$ and $w \in W_{k-1}w \in W'_{k-1}$. The cases for the other temporal modalities are similar (see also the proof of Lemma 29). The cases for the implication are similar as in the proof of Proposition 19. The remaining cases are straighforward. Suppose that $w \in W_{k-1}w \in W'_{k-1}$ and $\mathcal{M}, \pi(w) \vDash \circ\psi$. Then $\psi \in \Sigma(S(w))$. Since $S'(w) = \rho(S(w))$, $\pi(S'(w)) = \iota(S'(w))$ and (ρ, ι) is a condensation, $\Sigma(S(w)) = \lambda_{\mathcal{T}}(S'(w)) = \Sigma(\pi(S'(w)))$. Since $\mathcal{M}, \pi(S'(w)) \vDash \psi$, by induction hypothesis, $\mathcal{M}', S'(w) \vDash \psi$. Hence $\mathcal{M}', w \vDash \circ\psi$.and $\mathcal{M}, \pi(S'(w)) \vDash \psi$. By induction hypothesis, $\mathcal{M}', S'(w) \vDash \psi$, hence $\mathcal{M}', w \vDash \circ\psi$. The other direction is similar. ◀

Now, let us consider a stratified model $\mathcal{M} = (W, \preccurlyeq, S, V)$ with w_0 the root of W_0. The finite model $\mathcal{M}^{\text{fin}} = \left(W^{\text{fin}}, \preccurlyeq^{\text{fin}}, S^{\text{fin}}, V^{\text{fin}}\right)$ with a state w_0^{fin} such that $\Sigma(w_0^{\text{fin}}) = \Sigma(w_0)$ is constructed by the following procedure. This procedure is in three phases plus a final step. At each step, the model $\mathcal{M}$ is modified in place. Moreover, three index variables are maintained by the procedure:

- The variable i, initialized to 0, indicates the current labeled tree W_i which is considered.
- The variable j, initially undefined, indicates the index of the first labeled tree occuring infinitely often up to bimersion.
- The variable ℓ, initially undefined, holds the index of the last labeled tree that must not be modified.

As an invariant, $\mathcal{M}$ is stratified until the final step and for all $k < i$, $\mathcal{M}_k$ is a copy of a normalized labeled tree.

First phase

- If there is $k < i$ such that $\mathcal{M}_k \trianglelefteq \mathcal{M}_i$, replace $\mathcal{M}_i$ with $\mathcal{M}_k$, set i to $k+1$ and redo the same phase.
- If not, and for all $x > i$ there is $y > x$ such that $\mathcal{M}_y \trianglelefteq \mathcal{M}_i$, then replace $\mathcal{M}_i$ with a copy of its normalized Σ-quasimodel, increase i by one, set j and ℓ to i and start the next phase.
- Otherwise, replace $\mathcal{M}_i$ with its normalized Σ-quasimodel, increase i by one and redo the same phase.

Second phase. In this phase, we need to care about eventualities. To this end, a current eventuality (w, ψ), initially undefined, is maintained across the executions of the phase. Let w_x denote the element of the fulfillment of (w, ψ) belonging to W_x (if it exists), and $\mathcal{M}_x^+$ be the pointed structure $\mathcal{M}_x^{w_x}$. The phase proceeds through the following steps:

- If (w, ψ) is defined and the last element of the fulfillment of (w, ψ) belongs to some W_k with $k \leq i$ then undefine (w, ψ), set ℓ to i and repeat the same phase.
- If (w, ψ) is undefined then choose an eventuality (w, ψ) such that $w \in W_j$ and the last element of its fulfillment belongs to some W_k with $k > i$ and we and repeat the same phaseMissing text. If there is no such eventuality then start the next phase.
- If (w, ψ) is defined and there is k such that $\ell < k < i$ and $\mathcal{M}_k^+ \trianglelefteq \mathcal{M}_i^+$, then replace $\mathcal{M}_i$ with $\mathcal{M}_k$ connecting w_k to w_i, set i to $k+1$ and redo the same phase.
- Otherwise, replace $\mathcal{M}_i$ with a copy of the normalized labeled tree of $\mathcal{M}_k$ preserving w_i, increase i and redo the same phase.

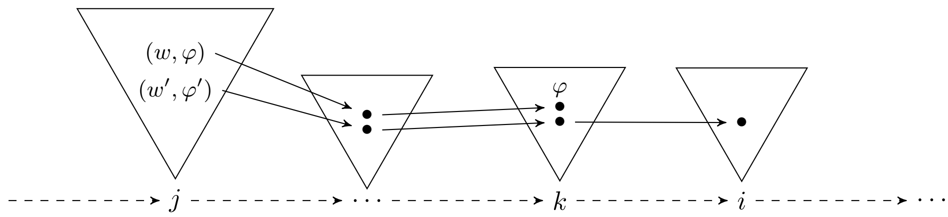

Figure 4 The stratum $\mathcal{M}_j$ and two of its eventualities. The fulfillment of (w,φ) is displayed, as well as the initial portion of the fulfillment of (w',φ').

Third phase

- If $\mathcal{M}_i \trianglelefteq \mathcal{M}_j$, then start the final step.
- If there is k such that $\ell < k < i$ and $\mathcal{M}_k \trianglelefteq \mathcal{M}_i$, then replace $\mathcal{M}_i$ with $\mathcal{M}_k$, set i to $k+1$ and redo the same phase.
- Otherwise, replace $\mathcal{M}_i$ with a copy of its normalized Σ-quasimodel, increase i by one and redo the same phase.

Final step. There is an immersion $\sigma\colon W_i \to W_j$. Construct the final tuple $(W^{\text{fin}}, \preccurlyeq^{\text{fin}}, S^{\text{fin}}, V^{\text{fin}})$ such that $W^{\text{fin}} = \bigcup_{0\le m<i} W_m$, $\preccurlyeq^{\text{fin}} = \preccurlyeq\downharpoonright_{W^{\text{fin}}}$,

$$S^{\text{fin}}(w) = \begin{cases} \sigma\left(S\left(w\right)\right) & \text{if } w \in W_{i-1} \\ S(w) & \text{otherwise} \end{cases}$$

$V^{\text{fin}} = V\downharpoonright_{W^{\text{fin}}}$, and w_0^{fin} is the root of W_0 (note that $w_0^{\text{fin}} \in W^{\text{fin}}$).

▶ **Lemma 29.** *The final tuple is a model and* $\Sigma(w_0^{\text{fin}}) = \Sigma(w_0)$.

Proof. The proof that $\mathcal{M}^{\text{fin}} = \left(W^{\text{fin}}, \preccurlyeq^{\text{fin}}, S^{\text{fin}}, V^{\text{fin}}\right)$ is a model is straightforward and left to the reader. We prove by structural induction on φ that for all $w \in W^{\text{fin}}$ and all $\varphi \in \Sigma$, $\mathcal{M}^{\text{fin}}, w \vDash \varphi$ iff $\mathcal{M}, w \vDash \varphi$. The cases for propositional variables and the boolean connectives are straightforward. The case for the next temporal modality is similar as in the proof of Lemma 28. For the eventually and henceforth temporal modalities, suppose first that (w,φ) is an eventuality in $\mathcal{M}$ and $w \in W^{\text{fin}}$. Let $w_0 \ldots w_n$ be the fulfillment of (w,φ) in $\mathcal{M}$. If $w_n \in W^{\text{fin}}$ then by induction hypothesis, (w,φ) is an eventuality in $\mathcal{M}^{\text{fin}}$. Otherwise, there is $k \le n$ such that $w_k \in W_i$. Therefore, (w_k,φ) is an eventuality in $\mathcal{M}$ and so is $(\sigma(w_k),\varphi)$. Since by construction, after the second phase, the length of the fulfillment of any eventuality (v,φ) such that $v \in W_j$ is bounded by $1+i-j$, (w,φ) is an eventuality in $\mathcal{M}^{\text{fin}}$. Conversely, suppose now that (w,φ) is an eventuality in $\mathcal{M}^{\text{fin}}$ and let $w_0 \ldots w_n$ be its fulfillment. For each $k \le n$ let m_k be such that $w_k \in W_{m_k}$. The proof is by a subinduction on the number r of $k \in \{1..n\} k \in 1..n$ such that $m_k = j$. If $r = 0$ then by induction hypothesis, (w,φ) is an eventuality in $\mathcal{M}$. If $r > 0$, let $k > 0$ be the least index such that $m_k = j$. If $k = n$ then suppose that $\varphi = \Diamond\psi$, the other case being symmetric. We have the other case beeing symmetric When have$\mathcal{M}^{\text{fin}}, w_k \vDash \psi$ and by induction $\mathcal{M}, w_k \vDash \psi$. Since $k > 0$, $w_k = S^{\text{fin}}(w_{k-1}) = \sigma(S(w_{k-1}))$ and since σ is an immersion, $\mathcal{M}, S(w_{k-1}) \vDash \psi$. Therefore (w,φ) is an eventuality in $\mathcal{M}$. Finally, if $r > 0$ and $k < n$ then (w_k,φ) is an eventuality in $\mathcal{M}^{\text{fin}}$ and by the subinduction hypothesis (w_k,φ) is an eventuality in $\mathcal{M}$. Since $k > 0$, $w_k = S^{\text{fin}}(w_{k-1}) = \sigma(S(w_{k-1}))$. MoreoverMorevoer, since σ is an immersion, $(S(w_{k-1}),\varphi)$ is an eventuality in $\mathcal{M}$. Hence (w,φ) is an eventuality in $\mathcal{M}$. ◀

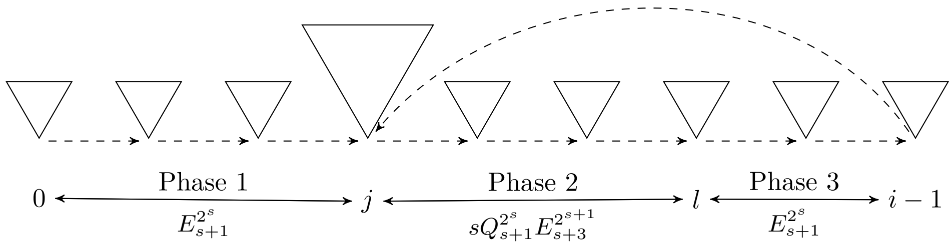

Figure 5 An illustration of the three phases of $\mathcal{M}^{\text{fin}}$. Below each phase we indicate the number of strata, used for the computations in the proof of Lemma 30.

▶ **Lemma 30.** *The cardinality of* W^{fin} *is bounded by*

$$B(s) \stackrel{def}{=} Q_{s+3}^{2^{s+1}} \left(2E_{s+1}^{2^s} + sQ_{s+1}^{2^s} E_{s+3}^{2^{s+1}}\right)$$

where $s = |\Sigma|$.

Proof. See Appendix. ◀

We have proved the following strong finite model property.

▶ **Theorem 31.** *There exists a computable function* B *such that for any formula* $\varphi \in \mathcal{L}$*, if* φ *is satisfiable (resp. unsatisfiable) then* φ *is satisfiable (resp. falsifiable) in a model* $\mathcal{M} = (W, \preccurlyeq, S, V)$ *such that* $|W| \leq B(|\varphi|)$.

Proof. In view of Theorem 8, a formula φ is satisfiable (resp. falsifiable) in a model $\mathcal{M}$ if and only if it is satisfied (resp. falsified) at the root of a stratified model $\mathcal{M}^{\text{e}}$. Then, by Lemma 29, φ is satisfied (resp. falsified) in $\mathcal{M}^{\text{e}}$ if and only if it is satisfied (res. falsified) on $(\mathcal{M}^{\text{e}})^{\text{fin}}$, which is effectively bounded by $B(|\varphi|)$ by Lemma 30. ◀

As a corollary, we get the decidability of $\mathsf{ITL}^{\mathsf{e}}$.

▶ **Corollary 32.** *The satisfiability and validity problems for* $\mathsf{ITL}^{\mathsf{e}}$ *are decidable.*

6 Conclusion

We have introduced $\mathsf{ITL}^{\mathsf{e}}$, an intuitionistic analogue of LTL based on expanding domain models from modal logic. In the literature, intuitionistic modal logic is typically interpreted over persistent models, but as we have shown this interpretation has the technical disadvantage of not enjoying the finite model property. Of course, this fact alone does not imply that $\mathsf{ITL}^{\mathsf{p}}$ is undecidable, and whether the latter is true remains an open problem. Meanwhile, our semantics are natural in the sense that we impose the minimal conditions on S so that all truth values are upwards closed under $\preccurlyeq$, and a wider class of models is convenient as they can more easily be tailored for specific applications.

This is an exploratory work, being the first to consider the logic $\mathsf{ITL}^{\mathsf{e}}$. As can be gathered from the tools we have developed, understanding this logic poses many technical challenges, and many interesting questions remain open. Perhaps the most pressing is the complexity of validity and satisfiability: the decision procedure we have given is non-elementary, but there seems to be little reason to assume that this is optimal. It may be possible to further 'trim' the model $\mathcal{M}^{\text{fin}}$ to obtain one that is elementarily bounded. However, we should not expect polynomially bounded models, as $\mathsf{ITL}^{\mathsf{e}}$ is conservative over intuitionistic propositional logic, which is already PSPACE-complete. Finally, we leave open the problem of finding a sound and complete axiomatization for $\mathsf{ITL}^{\mathsf{e}}$.

Acknowledgements. We would like to thank the anonymous referees for their useful suggestions which helped improve the presentation.

References

1 J.-M. Alliot, M. Diéguez, and L. Fariñas del Cerro. Metabolic pathways as temporal logic programs. In *Logics in Artificial Intelligence – 15th European Conference, JELIA 2016, Larnaca, Cyprus, November 9–11, 2016, Proceedings*, pages 3–17, 2016.
2 P. Balbiani and M. Diéguez. Temporal here and there. In M. Loizos and A. Kakas, editors, *Logics in Artificial Intelligence*, pages 81–96. Springer, 2016.
3 G. Boenn, M. Brain, M. De vos, and J. Ffitch. Automatic music composition using answer set programming. *Theory Pract. Log. Program.*, 11(2-3):397–427, 2011.
4 P. Cabalar and M. Diéguez. Strong Equivalence of Non-Monotonic Temporal Theories. In *Principles of Knowledge Representation and Reasoning: Proceedings of the 14th International Conference, KR'14, Vienna, Austria*, July 20–24, 2014.
5 P. Cabalar and G. Pérez. Temporal Equilibrium Logic: A First Approach. In *EUROCAST'07*, page 241–248, Las Palmas de Gran Canaria, Spain, 2007.
6 D. Van Dalen. Intuitionistic logic. In *Handbook of Philosophical Logic*, volume 166, pages 225–339. Springer Netherlands, 1986.
7 R. Davies. A temporal-logic approach to binding-time analysis. In *Proceedings, 11th Annual IEEE Symposium on Logic in Computer Science, New Brunswick, New Jersey, USA, July 27–30, 1996*, pages 184–195, 1996.
8 J. M. Davoren. On intuitionistic modal and tense logics and their classical companion logics: Topological semantics and bisimulations. *Annals of Pure and Applied Logic*, 161(3):349–367, 2009.
9 W. B. Ewald. Intuitionistic tense and modal logic. *The Journal of Symbolic Logic*, 51(1):166–179, 1986.
10 L. Fariñas del Cerro, A. Herzig, and E. Iraz Su. Epistemic equilibrium logic. In *IJCAI'15*, pages 2964–2970, Buenos Aires, Argentina, 2015. AAAI Press.
11 D. Fernández-Duque. The intuitionistic temporal logic of dynamical systems. *arXiv*, 1611.06929 [math.LO], 2016.
12 D. M. Gabbay, A. Kurucz, F. Wolter, and M. Zakharyaschev. *Many-Dimensional Modal Logics: Theory and Applications, Volume 148 (Studies in Logic and the Foundations of Mathematics)*. North Holland, 1 edition, 2003.
13 D. Gabelaia, A. Kurucz, F. Wolter, and M. Zakharyaschev. Non-primitive recursive decidability of products of modal logics with expanding domains. *Annals of Pure and Applied Logic*, 142(1-3):245–268, 2006.
14 M. Gebser, C. Guziolowski, M. Ivanchev, T. Schaub, A. Siegel, S. Thiele, and P. Veber. Repair and prediction (under inconsistency) in large biological networks with answer set programming. In *KR'10*, 2010.
15 A. Heyting. *Die formalen Regeln der intuitionistischen Logik.* Sitzungsberichte der Preussischen Akademie der Wissenschaften. Physikalisch-mathematische Klasse. Deütsche Akademie der Wissenschaften zu Berlin, Mathematisch-Naturwissenschaftliche Klasse, 1930.
16 D. Inclezan. An application of ASP to the field of second language acquisition. In *LPNMR'13*, pages 395–400, 2013.

17 N. Kamide and H. Wansing. Combining linear-time temporal logic with constructiveness and paraconsistency. *J. Applied Logic*, 8(1):33–61, 2010.

18 K. Kojima and A. Igarashi. Constructive linear-time temporal logic: Proof systems and Kripke semantics. *Information and Computation*, 209(12):1491 – 1503, 2011.

19 J. B. Kruskal. Well-quasi-ordering, the tree theorem, and vazsonyi's conjecture. *Transactions of the American Mathematical Society*, 95(2):210–225, 1960.

20 V. Lifschitz, D. Pearce, and A. Valverde. Strongly Equivalent Logic Programs. *ACM Transactions on Computational Logic*, 2(4):526–541, 2001.

21 V. Marek and M. Truszczyński. *Stable models and an alternative logic programming paradigm*, pages 169–181. Springer-Verlag, 1999.

22 G. Mints. *A Short Introduction to Intuitionistic Logic*. 2000.

23 I. Niemelä. Logic Programs with Stable Model Semantics as a Constraint Programming Paradigm. *Annals of Mathematics and Artificial Intelligence*, 25(3-4):241–273, 1999.

24 H. Nishimura. Semantical analysis of constructive PDL. *Publications of the Research Institute for Mathematical Sciences, Kyoto University*, 18:427–438, 1982.

25 M. Nogueira, M. Balduccini, M. Gelfond, R. Watson, and M. Barry. An A-Prolog decision support system for the space shuttle. In *AAAI Spring Symposium*, 2001.

26 D. Pearce. A New Logical Characterisation of Stable Models and Answer Sets. In *Proc. of Non-Monotonic Extensions of Logic Programming (NMELP'96)*, pages 57–70, Bad Honnef, Germany, 1996.

27 G. Plotkin and C. Stirling. A framework for intuitionistic modal logics: Extended abstract. In *Proceedings of the 1986 Conference on Theoretical Aspects of Reasoning About Knowledge*, TARK'86, pages 399–406, San Francisco, CA, USA, 1986. Morgan Kaufmann Publishers Inc.

28 A. K. Simpson. *The proof theory and semantics of intuitionistic modal logic*. PhD thesis, University of Edinburgh, UK, 1994.

A Proof of Lemma 30

Proof. Let us consider the stratified model $\mathcal{M} = (W, \preccurlyeq, S, V)$ obtained after the third phase. For all $k < i$ (where i has the value assigned at the end of this phase), W_k is a copy either of a normalized Σ-quasimodel or of a pointed normalized Σ-quasimodel. By Propositions 24 and 27, for all $k < i$, $|W_k| \leq Q_{s+3}^{2^{s+1}}$. We prove now that

$$i \leq 2E_{s+1}^{2^s} + sQ_{s+1}^{2^s}E_{s+3}^{2^{s+1}}.$$

After the first phase, by Proposition 24, we have $j \leq E_{s+1}^{2^s}$ and $|W_j| \leq Q_{s+1}^{2^s}$. Therefore, during the second phase, the current eventuality is defined at most $sQ_{s+1}^{2^s}$ times. Moreover, each time the current eventuality is undefined, by Proposition 27 we have that $i - \ell \leq E_{s+3}^{2^{s+1}}$ we have that $i - \ell \leq E_{n+3}^{2^{n+1}}$. Therefore, when the second phase terminates,

$$\ell - j \leq sQ_{s+1}^{2^s}E_{s+3}^{2^{s+1}}.$$

Finally, after the third phase, by Proposition 24, $i - \ell \leq E_{s+1}^{2^s}$. ◀