On the Hybrid Extension of CTL and CTL$^+$

Ahmet Kara[1], Volker Weber[1], Martin Lange[2], and Thomas Schwentick[1]

[1] Technische Universität Dortmund
[2] Ludwig-Maximilians-Universität München

Abstract. The paper studies the expressivity, relative succinctness and complexity of satisfiability for hybrid extensions of the branching-time logics CTL and CTL$^+$ by variables. Previous complexity results show that only fragments with *one variable* do have elementary complexity. It is shown that H^{1}CTL$^+$ and H^{1}CTL, the hybrid extensions with one variable of CTL$^+$ and CTL, respectively, are expressively equivalent but H^{1}CTL$^+$ is exponentially more succinct than H^{1}CTL. On the other hand, HCTL$^+$, the hybrid extension of CTL with arbitrarily many variables does not capture CTL*, as it even cannot express the simple CTL* property EGFp. The satisfiability problem for H^{1}CTL$^+$ is complete for triply exponential time, this remains true for quite weak fragments and quite strong extensions of the logic.

1 Introduction

Reasoning about trees is at the heart of many fields in computer science , such as verification and semistructured data. A wealth of sometimes quite different frameworks has been proposed for this purpose, according to the needs of the respective application. For reasoning about computation trees as they occur in verification, branching-time logics like CTL and tree automata are two such frameworks. In fact, they are closely related [25].
In some settings, the ability to mark a node in a tree and to refer to this node turned out to be useful. As neither classical branching-time logics nor tree automata provide this feature, many different variations have been considered, including tree automata with pebbles [9, 24, 29], memoryful CTL* [16], branching-time logics with forgettable past [18, 19], and logics with the "freeze" operator [14], the latter ones in the context of data trees [23]. As classical logic naturally provides means to refer to a node, namely constants and variables, it is an obvious question how these means can be incorporated into branching-time logics without losing their desirable properties which made them prevailing in verification [26].
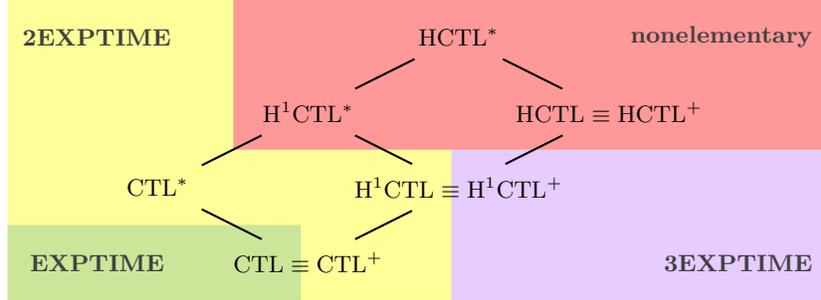
This question leads into the field of hybrid logics, where such extensions of temporal logics are studied [3]. In particular, a hybrid extension of CTL has been introduced in [29].
As usual for branching-time logics, formulas of their hybrid extensions are evaluated at nodes of a computation tree, but it is possible to bind a variable to the current node, to evaluate formulas relative to the root and to check whether the current node is bound to a variable. As an example, the HCTL-formula $\downarrow x @_{\mathrm{root}} \mathrm{EF}(p \wedge \mathrm{EF}x)$ intuitively says "I can place x at the current node, jump back to the root, go to a node where p holds and follow some (downward) path to reach x again. Or, equivalently: "there was a node fulfilling p in the past of the current node".

In this paper we continue the investigation of hybrid extensions of classical branching-time logics started in [29]. The main questions considered are (1) expressivity, (2) complexity of the satisfiability problem, and (3) succinctness. Figure 1 shows our results in their context. The complexity of the model checking problems will be studied in future work.

Classical branching-time logics are CTL (with polynomial time model checking and exponential time satisfiability) and CTL* (with polynomial space model checking and doubly exponential time satisfiability test). As CTL is sometimes not expressive enough[3] and CTL* is considered too expensive for some applications, there has been an intense investigation of

---

[3] Some things cannot be expressed at all, some only in a very verbose way.

2EXPTIME $\quad$ HCTL* $\quad$ nonelementary

H^{1}CTL* $\qquad$ HCTL $\equiv$ HCTL$^+$

CTL* $\qquad$ H^{1}CTL $\equiv$ H^{1}CTL$^+$

EXPTIME $\quad$ CTL $\equiv$ CTL$^+$ $\qquad$ 3EXPTIME

Fig. 1. Expressivity and complexity of satisfiability for hybrid branching-time logics. The lines indicate strict inclusion, unrelated logics are incomparable.

intermediate logics. We take up two of them here: CTL$^+$, where a path formula is a Boolean combination of basic path formulas[4] and ECTL, where fairness properties can be stated explicitly.

Whereas (even simpler) hybrid logics are undecidable over arbitrary transition systems [1], their restriction to trees is decidable via a simple translation to Monadic Second Order logic. However, the complexity of the satisfiability problem is high even for simple hybrid temporal logics over the frame of natural numbers: nonelementary [10] , even if only two variables are allowed [22, 29]. The one variable extension of CTL, H^{1}CTL, behaves considerably better, its satisfiability problem can be solved in **2EXPTIME** [29]. This is the reason why this paper concentrates on natural extensions of this complexity-wise relatively modest logic.
Even H^{1}CTL can express properties that are not bisimulation-invariant (e.g., that a certain configuration can be reached along two distinct computation paths) and is thus not captured by CTL*. In fact, [29] shows that H^{1}CTL captures and is strictly stronger than CTL with past, another extension of CTL studied in previous work [15]. One of our main results is that H^{1}CTL (and actually even HCTL$^+$) does not capture ECTL (and therefore not CTL*) as it cannot express simple fairness properties like EGFp. To this end, we introduce a simple Ehrenfeucht-style game (in the spirit of [2]). We show that existence of a winning strategy for the second player in the game for a property P implies that P cannot be expressed in HCTL$^+$.

In [29] it is also shown that the satisfiability problem for H^{1}CTL* has nonelementary complexity. We show here that the huge complexity gap between H^{1}CTL and H^{1}CTL* does not yet occur between H^{1}CTL and H^{1}CTL$^+$: we prove that there is only an exponential complexity gap between H^{1}CTL and H^{1}CTL$^+$, even when H^{1}CTL$^+$ is extended by past modalities and fairness operators. We pinpoint the exact complexity by proving the problem complete for **3EXPTIME**.

The exponential gap between the complexities for satisfiability of H^{1}CTL and H^{1}CTL$^+$ already suggests that H^{1}CTL$^+$ might be exponentially more succinct than H^{1}CTL. In fact, we show an exponential succinctness gap between the two logics by a proof based on the height of finite models. This refines the method of [17] based on model size. It should be noted that an $\mathcal{O}(n)!$-succinctness gap between CTL and H^{1}CTL was established in [29]. We mention that there are other papers on hybrid logics and hybrid tree logics that do not study expressiveness or complexity issue, e.g., [11, 21].

The paper is organized as follows. Definitions of the logics we use are in Section 2. Expressivity results are presented in Section 3. The complexity results can be found in Section 4, the succinctness results in Section 5.

---

[4] Precise definitions can be found in Section 2.

2 Definitions

Tree logics. In this section, we define syntax and semantics of the logics we use. We assume the reader is familiar with the tree logics CTL and CTL* [5]. However, we review the definition of the syntax and semantics of them next. Formulas of CTL* are composed from *state formulas* φ and *path formulas* ψ. They have the following abstract syntax.

$$\varphi ::= p \mid \neg\varphi \mid \varphi \vee \varphi \mid \varphi \wedge \varphi \mid \mathrm{E}\psi \mid \mathrm{A}\psi$$
$$\psi ::= \varphi \mid \neg\psi \mid \psi \vee \psi \mid \psi \wedge \psi \mid \mathrm{X}\psi \mid \psi\mathrm{U}\psi$$

We use the customary abbreviations $\mathrm{F}\psi$ for $\top\mathrm{U}\psi$ and $\mathrm{G}\psi$ for $\neg\mathrm{F}\neg\psi$.

The semantics of formulas is defined inductively. The semantics of path formulas is defined relative to a tree[5] $\mathcal{T}$, a path π of $\mathcal{T}$ and a position $i \geq 0$ of this path. E.g., $\mathcal{T}, \pi, i \models \psi_1\mathrm{U}\psi_2$ if there is some $j \geq i$ such that $\mathcal{T}, \pi, j \models \psi_2$ and, for each $l, i \leq l < j$, $\mathcal{T}, \pi, l \models \psi_1$.

The semantics of state formulas is defined relative to a tree $\mathcal{T}$ and a node v of $\mathcal{T}$. E.g., $\mathcal{T}, v \models \mathrm{E}\psi$ if there is a path π in $\mathcal{T}$, starting from v such that $\mathcal{T}, \pi, 0 \models \psi$. A state formula φ holds in a tree $\mathcal{T}$ if it holds in its root. Thus, sets of trees can be defined by CTL* state formulas.

CTL is a strict sub-logic of CTL*. It allows only path formulas of the forms $\mathrm{X}\varphi$ and $\varphi_1\mathrm{U}\varphi_2$ where $\varphi, \varphi_1, \varphi_2$ are state formulas. CTL$^+$ is the sub-logic of CTL* where path formulas are Boolean combinations of formulas of the forms $\mathrm{X}\varphi$ and $\varphi_1\mathrm{U}\varphi_2$ and $\varphi, \varphi_1, \varphi_2$ are state formulas.

Hybrid logics. In hybrid logics, a limited use of variables is allowed. For a general introduction to hybrid logics we refer to [3]. As mentioned in the introduction, we concentrate in this paper on hybrid logic formulas with *one* variable x. However, as we also discuss logics with more variables, we define hybrid logics H^{k}CTL* with k variables. For each $k \geq 1$, the syntax of H^{k}CTL* is defined by extending CTL* with the following rules for state formulas.

$$\varphi ::= \downarrow x_i\, \varphi \mid x_i \mid @_{x_i}\, \varphi \mid \mathrm{root} \mid @_{\mathrm{root}}\, \varphi$$

where $i \in \{1, \ldots, k\}$. The semantics is now relative to a vector $\boldsymbol{u} = (u_1, \ldots, u_k)$ of nodes of $\mathcal{T}$ representing an assignment $x_i \mapsto u_i$. For a node v and $i \leq k$ we write $\boldsymbol{u}[i/v]$ to denote $(u_1, \ldots, u_{i-1}, v, u_{i+1}, \ldots, u_k)$. For a tree $\mathcal{T}$ a node v and a vector $\boldsymbol{u}$, the semantics of the new state formulas is defined as follows.

$$
\begin{aligned}
\mathcal{T}, v, \boldsymbol{u} &\models\, \downarrow x_i\, \varphi & \text{if} \quad & \mathcal{T}, v, \boldsymbol{u}[i/v] \models \varphi \\
\mathcal{T}, v, \boldsymbol{u} &\models\, x_i & \text{if} \quad & v = u_i \\
\mathcal{T}, v, \boldsymbol{u} &\models\, @_{x_i}\, \varphi & \text{if} \quad & \mathcal{T}, u_i, \boldsymbol{u} \models \varphi \\
\mathcal{T}, v, \boldsymbol{u} &\models\, \mathrm{root} & \text{if} \quad & v \text{ is the root of } \mathcal{T} \\
\mathcal{T}, v, \boldsymbol{u} &\models\, @_{\mathrm{root}}\, \varphi & \text{if} \quad & \mathcal{T}, r, \boldsymbol{u} \models \varphi, \text{ where } r \text{ is the root of } \mathcal{T}
\end{aligned}
$$

Similarly, the semantics of path formulas is defined relative to a tree $\mathcal{T}$, a path π of $\mathcal{T}$, a position $i \geq 0$ of π and a vector $\boldsymbol{u}$. E.g., $\mathcal{T}, \pi, i, \boldsymbol{u} \models \mathrm{X}\psi$ if $\mathcal{T}, \pi, i+1, \boldsymbol{u} \models \psi$.

Intuitively, to evaluate a formula $\downarrow x_i\, \varphi$ one puts a pebble x_i on the current node v and evaluates φ. During the evaluation, x_i refers to v (unless it is bound again by another $\downarrow x_i$-quantifier).

The hybrid logics H^{k}CTL$^+$ and H^{k}CTL are obtained by restricting H^{k}CTL* in the same fashion as for CTL$^+$ and CTL, respectively. The logic HCTL is the union of all logics H^{k}CTL, likewise HCTL$^+$ and HCTL*.

A state formula φ of a hybrid logic is *satisfiable* if there exists a tree $\mathcal{T}$ with $\mathcal{T}, r, \boldsymbol{u} \models \varphi$, where r is the root of $\mathcal{T}$ and $\boldsymbol{u} = (r, \ldots, r)$ is a vector of adequate length. In this case we also

---

[5] In general, we consider finite and infinite trees and, correspondingly, finite and infinite paths in trees. It should always be clear from the context whether we restrict attention to finite or infinite trees.

say that $\mathcal{T}$ is a *model* of φ (denoted as $\mathcal{T} \models \varphi$). A state formula φ is *finitely satisfiable* if it has a *finite* model.

Two path formulas ψ and ψ' are *equivalent* (denoted as $\psi \equiv \psi'$) if for all trees $\mathcal{T}$, all paths π of $\mathcal{T}$ and all vectors $\boldsymbol{u}$ of adequate length it holds: $\mathcal{T}, \pi, 0, \boldsymbol{u} \models \psi$ iff $\mathcal{T}, \pi, 0, \boldsymbol{u} \models \psi'$. Similarly, two state formulas φ and φ' are equivalent (denoted as $\varphi \equiv \varphi'$) if for all trees $\mathcal{T}$, all nodes v and all vectors $\boldsymbol{u}$ of adequate length it holds: $\mathcal{T}, v, \boldsymbol{u} \models \varphi$ iff $\mathcal{T}, v, \boldsymbol{u} \models \varphi'$. We say that a logic $\mathcal{L}'$ is at least as expressive as $\mathcal{L}$ (denoted as $\mathcal{L} \leq \mathcal{L}'$) if for every $\varphi \in \mathcal{L}$ there is a $\varphi' \in \mathcal{L}'$ such that $\varphi \equiv \varphi'$. $\mathcal{L}$ and $\mathcal{L}'$ have the *same expressive power* if $\mathcal{L} \leq \mathcal{L}'$ and $\mathcal{L}' \leq \mathcal{L}$. $\mathcal{L}'$ is *strict more expressive* than $\mathcal{L}$ if $\mathcal{L} \leq \mathcal{L}'$ but not $\mathcal{L}' \leq \mathcal{L}$.

Size, depth and succinctness. For each formula φ, we define its *size* $|\varphi|$ as usual and its *depth* $d(\varphi)$ as the nesting depth with respect to path quantifiers.

It should be remarked that the definition of $d(\varphi)$ is tailored for the proof of inexpressibility with respect to H^kCTL. For general H^kCTL^* formulas one would count also the nesting of temporal operators.

The formal notion of *succinctness* is a bit delicate. We follow the approach of [12] and refer to the discussion there. We say that a logic $\mathcal{L}$ is *h-succinct in* a logic $\mathcal{L}'$, for a function $h : \mathbb{N} \to \mathbb{R}$, if for every formula φ in $\mathcal{L}$ there is an equivalent formula φ' in $\mathcal{L}'$ such that $|\varphi'| \leq h(|\varphi|)$. $\mathcal{L}$ is *$\mathcal{F}$-succinct in* $\mathcal{L}'$ if $\mathcal{L}$ is h-succinct in $\mathcal{L}'$, for some h in function class $\mathcal{F}$. We say that $\mathcal{L}$ is *exponentially more succinct* than $\mathcal{L}'$ if $\mathcal{L}$ is *not* h-succinct in $\mathcal{L}'$, for any function $h \in 2^{o(n)}$.

Normal forms. It will sometimes be convenient to restrict the set of operators that have to be considered. To this end, we say that a H^kCTL formula is in E-*normal form*, if it does not use the path quantifier A at all. A formula is in U-*normal form* if it only uses the combinations EX, EU and AU (but not, e.g., EG and AX).

Proposition 1. *Let $k \geq 1$.*

(a) For each H^kCTL formula φ there is an equivalent H^kCTL-formula in U-normal form and the size of ψ is linear in the size of φ.

(b) For each H^kCTL formula φ there is an equivalent H^kCTL-formula in E-normal form.

Proof. (a) This can be easily shown just as for CTL. Actually, the original definition of CTL by Emerson and Clarke [5] used only EX, EU and AU.

(b) This is straightforward as $\text{A}(\psi\text{U}\chi)$ can equivalently expressed as $(\neg\text{E}(\neg\chi)\,\text{U}(\neg\chi \wedge \neg\psi)) \wedge (\neg\text{EG}\neg\chi)$. However, it should be noted that the recursive application of this replacement rule may result in a formula of exponential size.

$\square$

3 Expressivity of HCTL and HCTL$^+$

3.1 The expressive power of HCTL$^+$ compared to HCTL

Syntactically CTL$^+$ extends CTL by allowing Boolean combinations of path formulas in the scope of a path quantifier A or E. Semantically this gives CTL$^+$ the ability to fix a path and test its properties by *several* path formulas. However in [6] it is shown that every CTL$^+$-formula can be translated to an equivalent CTL-formula. The techniques used there are applicable to the hybrid versions of these logics.

Theorem 2. *For every $k \geq 1$, H^kCTL has the same expressive power as H^kCTL^+.*

Proof. The main difficulty in the translation from CTL^+ to CTL can be described as follows: In a formula like $\text{E}[\text{F}\varphi_1 \wedge \ldots \wedge \text{F}\varphi_n]$ it is not determined in which order the formulas $\varphi_1, \ldots, \varphi_n$ hold on the path fixed by the quantifier E. In [6] this problem is solved by listing *all* possible orders. For instance the formula $\varphi = \text{E}[\text{F}\varphi_1 \wedge \text{F}\varphi_2]$ is equivalent to $\varphi' = \text{EF}(\varphi_1 \wedge \text{EF}\varphi_2) \vee \text{EF}(\varphi_2 \wedge \text{EF}\varphi_1)$. The transformation algorithm for CTL^+ to CTL in [6] is based on the following equivalences of CTL^+-formulas[6]:

(1) $\neg\text{X}\varphi \equiv \text{X}\neg\varphi$

(2) $\neg(\varphi\text{U}\varphi') \equiv [(\varphi \wedge \neg\varphi')\text{U}(\neg\varphi \wedge \neg\varphi')] \vee \text{G}\neg\varphi'$

(3) $\text{E}(\psi \vee \psi') \equiv \text{E}\psi \vee \text{E}\psi'$

(4) $\text{X}\varphi \wedge \text{X}\varphi' \equiv \text{X}(\varphi \wedge \varphi')$

(5) $\text{G}\varphi \wedge \text{G}\varphi' \equiv \text{G}(\varphi \wedge \varphi')$

(6) $\text{E}[\bigwedge\limits_{i=1}^{n}(\varphi_i\text{U}\varphi_i') \wedge \text{X}\chi \wedge \text{G}\xi] \equiv \bigvee\limits_{I \subseteq \{1,\ldots,n\}} [\bigwedge\limits_{i \in I}\varphi_i' \wedge \xi \wedge \bigwedge\limits_{i \notin I}\varphi_i \wedge \text{EX}(\chi \wedge \text{E}(\bigwedge\limits_{i \notin I}(\varphi_i\text{U}\varphi_i') \wedge \text{G}\xi))]$

(7) $\text{E}[\bigwedge\limits_{i=1}^{n}(\varphi_i\text{U}\varphi_i') \wedge \text{G}\chi] \equiv \bigvee\limits_{\pi \in Permutation(\{1,\ldots,n\})} [\text{E}[(\bigwedge\limits_{i=1}^{n}\varphi_i \wedge \chi)\text{U}(\varphi_{\pi(1)}' \wedge$
$\text{E}[(\bigwedge\limits_{i \neq \pi(1)}\varphi_i \wedge \chi)\text{U}(\varphi_{\pi(2)}' \wedge \text{E}[(\bigwedge\limits_{i \neq \pi(1),\pi(2)}\varphi_i \wedge \chi)\text{U}(\varphi_{\pi(3)}' \ldots \text{U}(\varphi_{\pi(n)}' \wedge \text{EG}\chi)\ldots)])])]]$

As explained above in equivalence (7) a disjunction of all possible orders of the formulas φ_i' is formulated. These equivalences also hold for H^kCTL^+. It can easily be shown that the occurence of root, $@_{\text{root}}$, $\downarrow x$ or $@_x$ for a variable x does not destroy any of the equivalences. Furthermore, as already indicated, if a formula on the right or the left side of one of the equivalences is in the scope of $\downarrow x$ then the node to which x is assigned is (up to a new $\downarrow x$) unique which means that x can be treated like a usual proposition. Altogether the translation algorithm for CTL^+ to CTL presented in [6] also gives a translation algorithm from H^kCTL^+ to H^kCTL for every $k \geq 1$. In [6] it is noticed that the factorial blowup introduced by equivalence (7) is the worst blowup in the whole transformation process and since $n! = 2^{\mathcal{O}(n \ log \ n)}$ the transformation of a formula φ results in a formula of length $2^{\mathcal{O}(n \ log \ n)}$. □

The transformation algorithm in Theorem 2 also yields an upper bound for the succinctness between H^1CTL^+ and H^1CTL.

Corollary 3. H^1CTL^+ *is* $2^{\mathcal{O}(n \log n)}$-*succinct in* H^1CTL.

3.2 Fairness is not expressible in HCTL^+.

In this subsection, we show the following result.

Theorem 4. *There is no formula in* HCTL^+ *which is logically equivalent to* $\text{E}\overset{\infty}{\text{F}}p$.

Here, $\mathcal{T}, v, \boldsymbol{u} \models \text{E}\overset{\infty}{\text{F}}\varphi$ if there is a path π starting from v that has infinitely many nodes v' with $\mathcal{T}, v', \boldsymbol{u} \models \varphi$. As an immediate consequence of this theorem, HCTL^+ does not capture CTL^*.

In order to prove Theorem 4, we define an Ehrenfeucht-style game that corresponds to the expressive power of HCTL. A game for a different hybrid logic was studied in [2]. We show that if a set L of trees can be characterized by a HCTL-formula, the spoiler has a winning strategy in the game for L. We expect the converse to be true as well but do not attempt to prove it as it is not needed for our purposes here.

Let L be a set of (finite or infinite) trees. The HCTL-*game* for L is played by two players, the *spoiler* and the *duplicator*. First, the spoiler picks a number k which will be the number

---

[6] In [6] and its journal version [7] there are some slight inaccuracies in the equivalences. Here we list the corrected ones.

of rounds in the core game. Afterwards, the duplicator chooses two trees, $\mathcal{T} \in L$ and $\mathcal{T}' \notin L$. The goal of the spoiler is to make use of the difference between $\mathcal{T}$ and $\mathcal{T}'$ in the core game.

The *core game* consists of k rounds of moves, where in each round i a node from $\mathcal{T}$ and a node from $\mathcal{T}'$ are selected according to the following rules. The spoiler can choose whether she starts her move in $\mathcal{T}$ or in $\mathcal{T}'$ and whether she plays a node move or a path move.

In a *node move* she simply picks a node from $\mathcal{T}$ (or $\mathcal{T}'$) and the duplicator picks a node in the other tree. We refer to these two nodes by a_i (in $\mathcal{T}$) and a_i' (in $\mathcal{T}'$), respectively, where i is the number of the round.

In a *path move*, the spoiler first chooses one of the trees. Let us assume she chooses $\mathcal{T}$, the case of $\mathcal{T}'$ is completely analogous. She picks an already selected node a_j of $\mathcal{T}$, for some $j < i$ and a path π starting in a_j. However, a node a_j can only be selected if there is no other node a_l, $l < i$ below a_j. The duplicator answers by selecting a path π' from a_j'. Then, the spoiler selects some node a_i' from π' and the duplicator selects a node a_i from π.

The duplicator wins the game if at the end the following conditions hold, for every $i, j \leq k$:

- a_i is the root iff a_i' is the root;
- $a_i = a_j$ iff $a_i' = a_j'$;
- for every proposition p, p holds in a_i iff it holds in a_i';
- there is a (downward) path from a_i to a_j iff there is a path from a_i' to a_j';
- a_j is a child of a_i iff a_j' is a child of a_i'.

Theorem 5. *If a set L of (finite and infinite) trees can be characterized by a HCTL-formula, the spoiler has a winning strategy on the HCTL-game for L.*

Proof. Let L be a set of trees and $\varphi \in$ HCTL such that, for every tree $\mathcal{T}$, $\mathcal{T}$ is in L if and only if $\mathcal{T} \models \varphi$. We show that the spoiler has a winning strategy with k_φ rounds in the game for L, where k_φ only depends on φ.

The proof is by induction on the structure of φ. As usual, we have to prove a slightly stronger statement for the induction step. We show that, for every HCTL-formula φ with variables from $X_l := \{x_1, \ldots, x_l\}$, there is k_φ such that, for trees $\mathcal{T}, \mathcal{T}'$, nodes v from $\mathcal{T}$ and v' from $\mathcal{T}'$ and node vectors $\boldsymbol{u}$ and $\boldsymbol{u}'$, the spoiler has a winning strategy in the k_φ-round core game on $(\mathcal{T}, v, \boldsymbol{u})$ and $(\mathcal{T}', v', \boldsymbol{u}')$ if $\mathcal{T}, v, \boldsymbol{u} \models \varphi$ and $\mathcal{T}', v', \boldsymbol{u}' \not\models \varphi$.

Thus, the proof uses a slightly extended game, in which the duplicator does not only choose $\mathcal{T}$ and $\mathcal{T}'$ but also nodes v, v' and node vectors $\boldsymbol{u}, \boldsymbol{u}'$. The game starts in a situation where nodes $a_0 := v$ and $a_i := u_i$, for $1 \leq i \leq l$ are already selected in $\mathcal{T}$ and correspondingly in $\mathcal{T}'$. In the remaining k rounds $a_{l+1}, \ldots, a_{l+k}$ and $a_{l+1}', \ldots, a_{l+k}'$ are selected and the winning condition applies to $a_0, \ldots, a_{l+k}$ and $a_0', \ldots, a_{l+k}'$.

It is easy to see that the theorem follows from this extended statement.

If φ is atomic, it can only test propositional properties of v and hence the spoiler wins the game by choosing $k_\varphi = 0$.

The rest of the proof is by case distinction on the outermost operator or quantifier of φ.

- If $\varphi = \neg\psi$, the spoiler has a winning strategy for ψ by the hypothesis. She can simply follow that winning strategy whilst switching the roles of $\mathcal{T}$ and $\mathcal{T}'$. In particular, $k_\varphi = k_\psi$.
- If $\varphi = \psi \lor \chi$ the spoiler chooses $k_\varphi = \max(k_\psi, k_\chi)$. In the core game she either follows the winning strategy for ψ or for χ depending on whether $\mathcal{T}, v, \boldsymbol{u} \models \psi$ or $\mathcal{T}, v, \boldsymbol{u} \models \chi$.
- The case that $\varphi = \psi \land \chi$ is analogous to the previous one.
- If $\varphi = \mathrm{EX}\psi$ the spoiler chooses $k_\varphi = k_\psi + 1$. Let $\mathcal{T}, v, \boldsymbol{u}, \mathcal{T}', v', \boldsymbol{u}'$ be selected by the duplicator. As $\mathcal{T}, v, \boldsymbol{u} \models \varphi$ there is a child w of v such that $\mathcal{T}, w, \boldsymbol{u} \models \psi$. On the other hand, there is no child w' of v' with $\mathcal{T}', w', \boldsymbol{u}' \models \psi$. Thus, the spoiler can select $a_{l+1} := w$ and win the remaining k_ψ rounds no matter which child of v' is chosen by the duplicator. She simply mimics the strategy of the game for ψ on $(\mathcal{T}, a_{l+1}, \boldsymbol{u})$ and $(\mathcal{T}', a_{l+1}', \boldsymbol{u}')$. If the duplicator does not choose a child of v' the spoiler wins instantly.
- As $\mathrm{AX}\psi \equiv \neg\mathrm{EX}\neg\psi$, the case of $\varphi = \mathrm{AX}\psi$ is already covered by the previous cases.

– If $\varphi = \mathrm{E}(\psi\mathrm{U}\chi)$ the spoiler chooses $k_\varphi = \max(k_\psi + 2, k_\chi + 1)$. Let $\mathcal{T}, v, \boldsymbol{u}, \mathcal{T}', v', \boldsymbol{u}'$ be selected by the duplicator. As $\mathcal{T}, v, \boldsymbol{u} \models \mathrm{E}(\psi\mathrm{U}\chi)$, there is some node w below v such that $\mathcal{T}, w, \boldsymbol{u} \models \chi$ and, for each node z on the path from v to w it holds $\mathcal{T}, z, \boldsymbol{u} \models \psi$. The spoiler does a node move in $\mathcal{T}$ and selects $a_{l+1} = w$.

Let w' be the node selected by the duplicator. As $\mathcal{T}', v', \boldsymbol{u}' \not\models \varphi$, we can conclude that $\mathcal{T}', w', \boldsymbol{u}' \not\models \chi$ or, for some z' on the path from v' to w', $\mathcal{T}', z', \boldsymbol{u}' \not\models \psi$. In the former case, the game continues, in the latter case she selects $v'_{l+2} = z'$. In either case, she has a winning strategy for the remaining $\max(k_\psi, k_\chi)$ rounds by induction.

– If $\varphi = \mathrm{A}(\psi\mathrm{U}\chi)$, φ is equivalent to $\varphi_1 \wedge \varphi_2$ where $\varphi_1 = \neg\mathrm{EG}\neg\chi$ and $\varphi_2 = \neg\mathrm{E}((\neg\chi)\mathrm{U}(\neg\psi \wedge \neg\chi))$. The spoiler chooses $k_\varphi = \max(k_\psi, k_\chi) + 2$.

Let $\mathcal{T}, v, \boldsymbol{u}, \mathcal{T}', v', \boldsymbol{u}'$ be selected by the duplicator. If $\mathcal{T}', v', \boldsymbol{u}' \not\models \varphi_2$ the winning strategy of the spoiler is already given by the previous cases. Otherwise, $\mathcal{T}', v', \boldsymbol{u}' \models \mathrm{EG}\neg\chi$. Let ρ' be a path starting from v' such that $\mathcal{T}', \rho', 0 \models \mathrm{G}\neg\chi$. Let w' be the first node on this path for which none of the nodes $a'_1, \ldots, a'_l$ is below w'. The spoiler selects w' in a node move. Let w be the node selected by the duplicator. If there is a node z on the path from v to w such that $\mathcal{T}, z, \boldsymbol{u} \models \chi$, the spoiler chooses z in a subsequent node move and wins by induction as there is no corresponding node between v' and w'. Otherwise she makes a path move in which she first selects the sub-path π' of ρ' starting in w'. Let π be a path in $\mathcal{T}$ starting from w as selected by the duplicator. As $\mathcal{T}, w, \boldsymbol{u} \models \varphi_1$, there is a node z on π such that $\mathcal{T}, z, \boldsymbol{u} \models \chi$. The spoiler selects this node as a_{i+1} and, as the duplicator cannot find such a node on π' wins by induction.

– If $\varphi = \downarrow x_i \psi$, for some i, the spoiler simply chooses $k_\varphi = k_\psi$. After the selection of $\mathcal{T}, v, \boldsymbol{u}, \mathcal{T}', v', \boldsymbol{u}'$ by the duplicator she mimics the game for ψ on the structures $(\mathcal{T}, v, \boldsymbol{u}[i/v])$ and $(\mathcal{T}', v', \boldsymbol{u}'[i/v'])$.

– If $\varphi = @_{x_i}\psi$, for some i, the spoiler mimics the game on $(\mathcal{T}, u_i, \boldsymbol{u})$ and $(\mathcal{T}', u'_i, \boldsymbol{u}')$.

□

Now we turn to the proof of Thm. 4. It makes use of the following lemma which is easy to prove using standard techniques (see, e.g., [20]). The lemma will be used to show that the duplicator has certain move options on paths starting from the root. The parameter S_k given by the lemma will be used below for the construction of the structures $\mathcal{B}_k$.

For a string $s \in \Sigma^*$ and a symbol $a \in \Sigma$ let $|s|$ denote the length of s and $|s|_a$ the number of occurrences of a in s.

Lemma 6. *For each $k \geq 0$ there is a number $S_k \geq 0$ such that, for each $s \in \{0,1\}^*$ there is an $s' \in \{0,1\}^*$ such that $|s'| \leq S_k$ and $s \equiv_k s'$.*

Here, $\equiv_k$ is equivalence with respect to the k-round Ehrenfeucht game on strings (or equivalently with respect to first-order sentences of quantifier depth k). It should be noted that, if $k \geq 3$ and $s \equiv_k s'$, then the following conditions hold.
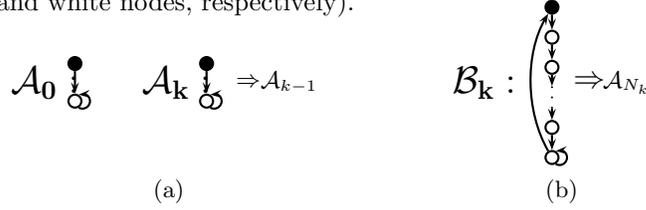
– $s \in \{0\}^*$ implies $s' \in \{0\}^*$.
– If the first symbol of s is 1 the same holds for s'.
– If s does not have consecutive 1's, s' does not either.

We fix some S_k, for each k.

The proof of Thm. 4 uses the HCTL-game defined above. Remember that the spoiler opens the game with the choice of a $k \in \mathbb{N}$ and the duplicator responds with two trees $\mathcal{T} \in L$ and $\mathcal{T}' \notin L$. We want to show that the duplicator has a winning strategy so we need to construct such trees, and then need to show that the duplicator has a winning strategy for the k-round core game on $\mathcal{T}$ and $\mathcal{T}'$.

We will use transition systems in order to finitely represent infinite trees. A transition system is a $\mathcal{K} = (V, E, v_0, \ell)$ where (V, E) is a directed graph, $v_0 \in V$, and ℓ labels each state $v \in V$ with a finite set of propositions. The *unraveling* $T(\mathcal{K})$ is a tree with node set V^+ and root v_0. A node $v_0 \ldots v_{n-1} v_n$ is a child of $v_0 \ldots v_{n-1}$ iff $(v_{n-1}, v_n) \in E$. Finally, the label of a node $v_0 \ldots v_n$ is $\ell(v_n)$.

Inspired by [8] we define transition systems $\mathcal{A}_i$, for each $i \geq 0$, as depicted in Fig. 2 (a). Nodes in which p holds are depicted black, the others are white (and we subsequently refer to them as black and white nodes, respectively).

$$\mathcal{A_0} \quad \mathcal{A_k} \Rightarrow \mathcal{A}_{k-1} \qquad \mathcal{B_k} : \Rightarrow \mathcal{A}_{N_k}$$

(a) (b)

Fig. 2. Illustration of the definition of (a) $\mathcal{A}_k$ and (b) $\mathcal{B}_k$. The path of white nodes in $\mathcal{B}_k$ consists of S_k nodes. The double arrow $\Rightarrow$ indicates that every white node on the left is connected to every black node on the right.

Thus, $\mathcal{A}_0$ has a black (root) node and a white node with a cycle. $\mathcal{A}_i$ has a black (root) node, a white node with a cycle and a copy of $\mathcal{A}_{i-1}$. Furthermore, there is an edge from the white node below the root of $\mathcal{A}_i$ to each black node in the copy of $\mathcal{A}_{i-1}$ (as indicated by $\Rightarrow$). Let $\mathcal{T}_i := T(\mathcal{A}_i)$. We first introduce some notation and state some simple observations concerning the tree $\mathcal{T}_i$.

(1) For a node v in $\mathcal{T}_i$ we denote the maximum number of black nodes on a path starting in v (and not counting v itself) the *height* $h(v)$ of v. Then the root of $\mathcal{T}_i$ has height i.
(2) If u and v are black nodes of some $\mathcal{T}_i$ with $h(u) = h(v)$ then the subtrees $T(u)$ and $T(v)$ induced by u and v are isomorphic.
(3) The height of a tree is defined as the height of its root.
(4) A white node v of height i has one white child (of height i) and i black children of heights $0, \ldots, i-1$. A black node has exactly one white son.
(5) Each finite path π of $\mathcal{T}_i$ induces a string $s(\pi) \in \{0,1\}^*$ in a natural way: $s(\pi)$ has one position, for each node of π, carrying a 1 iff the corresponding node is black.
(6) The root of $\mathcal{T}_i$ has only one child. We call the subtree induced by this (white!) child $\mathcal{U}_i$. If v is a white node of height i then $T(v)$ is isomorphic to $\mathcal{U}_i$.

Next we define numbers N_k inductively as follows: $N_0 := 0$ and $N_k := N_{k-1} + \max(S_3, S_k) + 1$.

The following lemma shows that the duplicator has a winning strategy in two structures of the same kind, provided they both have sufficient depth.

Lemma 7. *Let i, j, k be numbers such that $i, j \geq N_k$. Then the duplicator has a winning strategy in the k-round core game on*

(a) $\mathcal{T}_i$ and $\mathcal{T}_j$, and
(b) $\mathcal{U}_i$ and $\mathcal{U}_j$.

Proof of Lemma 7. The proof is by induction on k, the case $k = 0$ being trivial. Thus, let $k > 0$. We first show (a). Let us assume first that the spoiler makes a **node move** in $\mathcal{T}_i$ on v (node moves in $\mathcal{T}_j$ are symmetric). We distinguish two cases depending on the height of v. In both cases let r, r' be the roots of $\mathcal{T}_i$ and $\mathcal{T}_j$ respectively.

Case $h(v) > N_{k-1}$. Let π denote the path from r to v. By Lemma 6 there is a string s' with $|s'| \leq S_l$ such that $s(\pi) \equiv_l s'$, where $l = \max(k, 3)$. As $j \geq N_k = N_{k-1} + S_l + 1$, there is a node v' of height $\geq N_{k-1}$ in $\mathcal{T}_j$ such that the path[7] π' from r' to v' satifies $s(\pi') = s'$. The duplicator chooses v' as her answer in this round. We have to show that she has a winning strategy for the remaining $k - 1$ rounds. Her strategy is a composition of the following three strategies for different parts of the trees.[8]

---

[7] It should be noted that $l \geq 3$ guarantees in particular that s' does not have consecutive 1's.

[8] By a standard argument the different strategies can be combined into a strategy on $\mathcal{T}$ and $\mathcal{T}'$ (see again [20]). It is helpful here that path moves only involve paths that are below all previously selected nodes. Hence, a path move always only affects one of the sub-games.

8

(i) If the spoiler does a (node or path) move in $T(v)$ or $T(v')$ the duplicator can reply according to her winning strategy in the $(k-1)$-round game in these two trees which is guaranteed by induction as both have height $\geq N_{k-1}$.

(ii) If the spoiler chooses a node on π or π' then the duplicator answers following her strategy in the k-round Ehrenfeucht game on the strings $s(\pi)$ and $s(\pi')$.

(iii) The remaining (and most complicated) sub-strategy concerns moves elsewhere in $\mathcal{T}_i$ (the case of a move elsewhere in $\mathcal{T}'_j$ is again symmetric). Let w be a node chosen by the spoiler (in a node move or as the starting node of a path). Let y be the last node of π on the path ρ from r to w. Node y has two different successors – namely one on π and one on ρ, hence it must be white by fact (4) above. Let z be its successor on ρ. Let y' be the node corresponding to y on π' as induced by the winning strategy of the duplicator on $s(\pi)$ and $s(\pi')$.
   - If z has height $< N_{k-1}$ let z' be the unique (black!)[9] child of y' of the same height as z. By fact (2) above, $T(z)$ and $T(z')$ are isomorphic and the duplicator has a winning strategy in these two subtrees induced by an isomorphism.
   - If z has height $\geq N_{k-1}$ let z' be some child of y' (of the same color as z, not on π', and with height $\geq N_{k-1}$. (Note that z and z' can be both black or both white). By induction the duplicator has a winning strategy on $T(z)$ and $T(z')$.

Case $h(v) \leq N_{k-1}$. Let π be the path from r to v, and u_1 be the highest black node on π with $h(u_1) \leq N_{k-1}$. Then we must have $h(u_1) = N_{k-1}$ because π contains black nodes of height up to $i \geq N_k$. Hence, u_1 has a white parent u_2 s.t. $h(u_2) > N_{k-1}$. We determine a node u'_2 in $\mathcal{T}'$ in the same way we picked v' for v in the first case. In particular, $h(u'_2) \geq N_{k-1}$ and for the paths ρ leading from r to u_2 and ρ' leading from r' to u'_2 we have $s(\rho) \equiv_k s(\rho')$.

Let u'_1 be the black child of u'_2 of height $h(u_1)$. As $h(u_1) = h(u'_1)$ there is an isomorphism σ between $T(u_1)$ and $T(u_2)$ and we choose $v' := \sigma(v)$. An illustration is given in Figure 3.

The winning strategy of the duplicator for the remaining $k-1$ rounds follows σ on $T(u_1)$ and $T(u_2)$ and is analogous to the first case in the rest of the trees.

Next, we assume that the first move of the spoiler is a **path move** in $\mathcal{T}$ (path moves in $\mathcal{T}'$ are again symmetric).

Let π be the path chosen by the spoiler. Let v be the lowest black node of π. Let v' be the node chosen by the duplicator had the spoiler selected v in a node move and let π' be the path from r' to v' extended by the unique infinite white path below v'. We can distinguish the same two cases as for node moves.

If the node a' selected by the spoiler on π' is from $T(v')$ (in case 1) or from $T(u'_1)$ (in case 2), the duplicator can choose a corresponding node a by the isomorphism. Otherwise, the duplicator replies by the node a of ρ induced by the k-round (!) winning strategy of the duplicator on $s(\rho)$ and $s(\rho')$.

This completes the inductive step for (a).

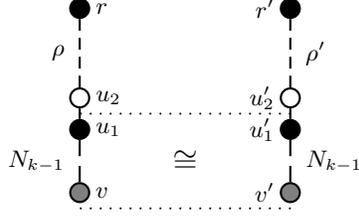For (b) the proof is completely analogous. This completes the proof of the lemma.

$\square$

We are now prepared to prove Thm. 4.

Proof of Theorem 4. By Theorem 2 it is sufficient to show that no formula equivalent to $\overset{\infty}{\mathrm{E}}\mathrm{F}p$ exists in HCTL. To this end, we prove that the duplicator has a winning strategy in the HCTL-game for the set of trees fulfilling $\overset{\infty}{\mathrm{E}}\mathrm{F}p$.

We next define transition systems $\mathcal{B}_k$, for $k \geq 0$. As illustrated in Figure 2 (b), $\mathcal{B}_k$ has a black root from which a path of length S_k of white nodes starts. The last of these white nodes has a self-loop and an edge back to the root. Furthermore, $\mathcal{B}_k$ has a copy of $\mathcal{A}_{N_k}$ and there is an edge from each white node of the initial path to each black node of the copy of $\mathcal{A}_{N_k}$.

---

[9] z must be black as all nodes on π have height $\geq N_{k-1}$ and only black nodes may have a smaller height than their parent.

Fig. 3. Illustration of the case where $h(v) \leq N_{k-1}$. The colors of v and v' are not known a priori.

Clearly, for each k, $\mathcal{B}_k \models \mathrm{E}\overset{\infty}{\mathrm{F}}p$ and $\mathcal{A}_k \not\models \mathrm{E}\overset{\infty}{\mathrm{F}}p$.

We show in the remainder of the proof that, for each k, the duplicator has a winning strategy in the k-round core game on $\mathcal{T} = T(\mathcal{B}_k)$ and $\mathcal{T}' = T(\mathcal{A}_{N_k})$.

We call the nodes of $\mathcal{T}$ induced by the extra nodes that $\mathcal{B}_k$ has over $\mathcal{A}_{N_k}$ *extra nodes*.

The proof is again by induction on k and it is very similar to the proof of Lemma 7. The case $k = 0$ is again trivial.

A node move v in $\mathcal{T}$ is answered just as it was in the proof of the lemma. The duplicator then wins by induction. Likewise node moves v' in $\mathcal{T}'$ are answered as in the proof of the lemma and (besides the root) no special black node of $\mathcal{T}$ is involved.

It remains to deal with path moves. Path moves starting in $\mathcal{T}'$ can be handled as in the proof of the lemma. Likewise path moves starting in $\mathcal{T}$ with a path π of finitely many black nodes can be handled along the same lines.

The only real new case is when the spoiler chooses a path π in $\mathcal{T}$ that has infinitely many black nodes. The duplicator answers by choosing the unique path π' of $\mathcal{T}'$ starting from r' that only consists of white nodes. Let v' be a node of π' that is chosen by the spoiler. By (the remark after) Lemma 6, there is a string s of $\leq S_k$ zeros such that $s \equiv_k s(\rho')$, where ρ' is the path from r' to v'. The duplicator thus picks the node v on π such that $s(\rho) = s$, where ρ denotes the path from r to v. Note that by construction, the initial white path of $\mathcal{T}$ is long enough to guarantee the existence of such a node v. That the duplicator has a winning strategy for the remaining $k - 1$ rounds can be shown along the same lines as before. Subsequent choices of paths with infinitely many black nodes are answered by paths with one black node and infinitely many white ones. □

4   Satisfiability of H¹CTL⁺

Theorem 8. *Satisfiability of* $\mathrm{H}^1\mathrm{CTL}^+$ *is hard for* **3EXPTIME**.

Proof. The proof is by reduction from a tiling game (with **3EXPTIME** complexity) to the satisfiability problem of $\mathrm{H}^1\mathrm{CTL}^+$. Actually we show that the lower bound even holds for the fragment of $\mathrm{H}^1\mathrm{CTL}^+$ without the U-operator (but with the F-operator instead).

An instance $I = (T, H, V, F, L, n)$ of the *2EXP-corridor tiling game* consists of a finite set T of *tile types*, two relations $H, V \subseteq T \times T$ which constitute the *horizontal and vertical constraints*, respectively, two sets $F, L \subseteq T$ which describe the starting and end conditions, respectively, and a number n given in unary. The game is played by two players, E and A, on a board consisting of 2^{2^n} columns and (potentially) infinitely many rows. Starting with player E and following the constraints H, V and F the players put tiles to the board consecutively from left to right and row by row. The constraints prescribe the following conditions:

- A tile t' can only be placed immediately to the right of a tile t if $(t, t') \in H$.
- A tile t' can only be placed immediately above a tile t if $(t, t') \in V$.
- The types of all tiles in the first row belong to the set F.

Player E wins the game if a row is completed containing only tiles from L or if A makes a move that violates the constraints. On the other hand, player A wins if E makes a forbidden move or the game goes on ad infinitum.

A winning strategy for E has to yield a countermove for all possible moves of A in all possible reachable situations. Furthermore, the starting condition and the horizontal and vertical constraints have to be respected. Finally, the winning strategy must guarantee that either player A comes into a situation where he can no longer make an allowed move or a row with tiles from L is completed.

The problem to decide for an instance I whether player E has a winning strategy on I is complete for **3EXPTIME**. This follows by a straightforward extension of [4].

Now we show in full detail how to build a formula φ_I of length $\mathcal{O}(|I| \cdot |T|)$ from an instance I with tile set T of the 2EXP-corridor tiling game such that φ_I is satisfiable if and only if player E has a winning strategy in the game for I. In fact, a tree will satisfy φ_I if and only if it encodes a winning strategy of player E. Here, the encoding tree represents all possible plays (for the various moves of player A) for a fixed (and winning) strategy of player E.

Encoding of the winning strategy for player E. We encode strategies for player E as T-labeled trees in which each move is represented by a sequence of nodes (see Figure 5). The first move of player E is represented by a sequence starting at the root. Each sequence corresponding to a move of E is followed by several branches, one for every possible next move of A. Each sequence corresponding to a move of A is followed by one sequence of nodes corresponding to the move of player E following the strategy. It is clear that such a tree represents a *winning* strategy if every root path corresponds to a sequence of moves resulting in a win for player E.

In the encoding we use the propositions $\{pos_e, pos_o, b_0, ..., b_{n-1}, row_e, row_o, b, o, c, q_\sharp\} \cup \{p_t \mid t \in T\}$. In order to be able to describe the constraints via a $\mathrm{H}^1\mathrm{CTL}^+$-formula of polynomial length we serially number all positions of a row of the board in the style of [27]. While in [27] a row[10] consists of 2^n positions we have to deal with 2^{2^n} positions in the current proof. We encode each position by a sequence of 2^n nodes, each of which we call *position bits*. For each of these nodes the propositions $b_0, ..., b_{n-1}$ encode a binary number. Each position bit in turn represents one bit of a binary number of length 2^n via proposition b[11]. Each such sequence is preceded by a *position node* which holds some additional information that will be described later. We call a sequence of length $2^n + 1$ representing one position of the tiling a *position sequence*. In each position bit of a position sequence proposition p_t holds for the tiling type t of its tile. A row is represented by 2^{2^n} position sequences preceded by a *row node*. Altogether, a row is represented by a *row sequence* consisting of $(2^n + 1)2^{2^n} + 1$ nodes.

For technical reasons, each position node of an even (odd) position is marked using the proposition pos_e, (pos_o, respectively). Likewise, the row nodes of even (odd) rows are marked using row_e (row_o). It is worth noting that the tree branches only[12] after position nodes in which pos_o holds as the odd positions are tiled by A.

To compare two nodes of a path that are far apart we use a technique that was originally invented in [27] and was also applied in [13]. To this end, we use two kinds of nodes: *original nodes* which are labelled with the proposition o and *copy nodes* labelled with the proposition c. For the encoding of the winning strategy only the original nodes are relevant. Each original node has a copy node with identical propositions (except for the proposition o) as a child. Likewise, each copy node has only copy nodes with identical propositions as children (see Figure 4). Copy nodes will enable us to "mark" an original node v by fixing a path π through v and its copy node child. Since copy nodes carry the propositions of their parent original nodes, assertions about an original node can be tested in any of their subsequent copy nodes.

Proposition $q_\sharp$ is used to label the part of the tree which does not belong to the encoding of the winning strategy.

---

[10] Actually, the proof in [27] uses alternating Turing machines and thus encodes configurations rather than rows.

[11] It should be noted that the *lowest* bit of this binary number is represented by the position bit with the *highest* number.

[12] After the modification in the next paragraph, this statement only holds with respect to *original nodes*.

Testing vertical constraints. The most difficult condition to test is that a tree respects the vertical constraints. Thus, we have to check the following condition: *For every row, except the first one, the tile of every position is consistent with the tile of the corresponding position in the previous row.*

To this end, we have to compare two position sequences representing corresponding positions in consecutive rows. Two sequences represent corresponding positions if, whenever two position bits are equivalent with respect to $b_0, \ldots, b_{n-1}$, they are also equivalent with respect to b. The technical challenge is to do this comparison with a formula of linear (as opposed to exponential) size. Finally, it has to be checked that the position bits of the two position sequences are consistent with respect to V.

Let r and r' be row sequences representing two consecutive rows. Let s' be some position sequence of r' for which consistency with the corresponding position of the previous row r shall be checked.

We first give an informal description of the formula that checks the vertical constraints. Assume that x is associated with the last position bit of the position sequence s' representing the j-th position in r'. We first construct a formula ξ that becomes true exactly at the first position bit v of the position sequence s representing position j in row r on the path to x.

To this end, ξ checks that there is a path starting in v, continuing at least to the last node of s (from where it might follow copy nodes) and from each non-copy node u on this path, there is a path leading to some copy node of a node u' in s' with exactly the same propositions $b_0, \ldots, b_{n-1}, b$.

This can be expressed as

$$\xi = o \wedge \bigwedge_{i=0}^{n-1} \neg b_i \wedge \mathrm{E}\big[\mathrm{F} \bigwedge_{i=0}^{n-1} b_i \wedge \mathrm{G}\big(o \rightarrow \mathrm{E}[(\mathrm{F}row_o \wedge \neg \mathrm{F}row_e \vee \mathrm{F}row_e \wedge \neg \mathrm{F}row_o) \wedge$$

$$\mathrm{G}(\neg c \rightarrow \mathrm{EF}x) \wedge \mathrm{FE}(\mathrm{F}x \wedge \mathrm{G}\neg\varphi_{pos}) \wedge \bigwedge_{i=0}^{n-1} (b_i \leftrightarrow \mathrm{F}(c \wedge b_i)) \wedge b \leftrightarrow \mathrm{F}(c \wedge b)]\big)\big]$$

Here, φ_{pos} is an abbreviation for $pos_o \vee pos_e$ indicating that a node is a position node. The path formula $\mathrm{F} \bigwedge_{i=0}^{n-1} b_i$ ensures that the current path continues at least to the last position bit of the current position from where it might follow copy nodes. The path formula $\mathrm{F}row_o \wedge \neg \mathrm{F}row_e \vee \mathrm{F}row_e \wedge \neg \mathrm{F}row_o$ makes sure that the current path meets exactly two rows. The path formula $\mathrm{G}(\neg c \rightarrow \mathrm{EF}x) \wedge \mathrm{FE}(\mathrm{F}x \wedge \mathrm{G}\neg\varphi_{pos})$ tests that the path reaches the same position sequence as the node carrying x but no subsequent position sequence.

The vertical constraints now hold if, whenever x is put to the last position node of some position sequence s' with tile type t', if at some node v the formula ξ holds then the tile type t at v has to be such that $(t, t') \in V$. This can now be expressed by $\mathrm{AG}\big[\bigwedge_{t' \in T}[(o \wedge \bigwedge_{i=0}^{n-1} b_i \wedge p_{t'}) \rightarrow \downarrow x.@_{\mathrm{root}}\mathrm{AG}(\xi \rightarrow \bigvee_{(t,t') \in V} p_t)]\big]$.

For an instance I of 2EXP-corridor tiling game we present the whole formula φ_I of length in $\mathcal{O}(|I||T|)$ such that φ_I is satisfiable if and only if player E has a winning strategy in the game for I. The formula φ_I is composed of the conjunction of $\chi = \bigwedge_{i=1}^{10} \chi_i$ and $\psi = \bigwedge_{i=1}^{7} \psi_i$ where χ describes the basic tree structure that is needed to formulate a strategy and the ψ guarantees that the model of φ_I corresponds to a winning strategy for player E. Each of the subformulas χ_i, ψ_i is of length $\mathcal{O}(|I||T|)$.

We first introduce some abbreviations:

- $\varphi_{pos} = pos_e \vee pos_o$ (*position* node)
- $\varphi_{row} = row_e \vee row_o$ (*row* node)
- $\varphi_{first} = o \wedge \bigwedge_{i=0}^{n-1} \neg b_i$ (*first* (and original) node in a position sequence)
- $\varphi_{last} = \bigwedge_{i=0}^{n-1} b_i$ (*last* node in a position sequence, not necessarily original)
- $\psi_{full} = \mathrm{G}\neg\varphi_{pos} \wedge \mathrm{F}(c \wedge \varphi_{last})$ (the path extends exactly until the end of the position sequence and continues with copy nodes; in this sense it is a *full* path)

- $\psi_{2pos} = (\mathrm{F}pos_e \wedge \neg \mathrm{F}pos_o) \vee (\mathrm{F}pos_o \wedge \neg \mathrm{F}pos_e)$ (the path meets *two* (consecutive) *position* sequences)
- $\psi_{2row} = (\mathrm{F}row_e \wedge \neg \mathrm{F}row_o) \vee (\mathrm{F}row_o \wedge \neg \mathrm{F}row_e)$ (the path meets *two* (consecutive) *row* sequences)

The first formula helps to describe some properties in a simple way.

T1: *Each node of the tree is labelled with exactly one of the propositions* row_e, row_o, pos_e, pos_o, $q_\sharp$, o *and* c.

$$
\begin{aligned}
\chi_1 = \mathrm{AG}[&(\varphi_{row} \vee \varphi_{pos} \vee q_\sharp \vee o \vee c)\wedge \\
&(pos_e \rightarrow \neg pos_o \wedge \neg row_e \wedge \neg row_o \wedge \neg q_\sharp \wedge \neg o \wedge \neg c)\wedge \\
&(pos_o \rightarrow \neg pos_e \wedge \neg row_e \wedge \neg row_o \wedge \neg q_\sharp \wedge \neg o \wedge \neg c)\wedge \\
&(row_e \rightarrow \neg pos_e \wedge \neg pos_o \wedge \neg row_o \wedge \neg q_\sharp \wedge \neg o \wedge \neg c)\wedge \\
&(row_o \rightarrow \neg pos_e \wedge \neg pos_o \wedge \neg row_e \wedge \neg q_\sharp \wedge \neg o \wedge \neg c)\wedge \\
&(q_\sharp \rightarrow \neg pos_e \wedge \neg pos_o \wedge \neg row_e \wedge \neg row_o \wedge \neg o \wedge \neg c)\wedge \\
&(o \rightarrow \neg pos_e \wedge \neg pos_o \wedge \neg row_e \wedge \neg row_o \wedge \neg q_\sharp \wedge \neg c)\wedge \\
&(c \rightarrow \neg pos_e \wedge \neg pos_o \wedge \neg row_e \wedge \neg row_o \wedge \neg q_\sharp \wedge \neg o)]
\end{aligned}
$$

T2: *The root induces with the proposition* row_e *the encoding of the first row and every node labelled with* row_e *or* row_o *has exactly one child labled with* pos_e *signalising the encoding of a new position.*

$$
\chi_2 = row_e \wedge \mathrm{AG}[\varphi_{row} \rightarrow \mathrm{EX}(pos_e \wedge {\downarrow}x.@_{\mathrm{root}}\mathrm{EF}(\mathrm{EX}x \wedge \mathrm{AX}x))]
$$

T3: *Every child of a* pos_e-*or* pos_o-*node is labelled with the initial position bit number encoded by* $b_0 \ldots b_{n-1}$.

$$
\chi_3 = \mathrm{AG}[\varphi_{pos} \rightarrow \mathrm{AX}\varphi_{first}]
$$

T4: *As long as the last position bit of a position is not reached, the next node is labelled with the next position bit number.*

In order to keep the length of χ_4 within the bound $\mathcal{O}(|I||T|)$, we make use of additional propositions $d_0, \ldots, d_{n-1}$ and $e_0, \ldots, e_{n-1}$. The idea is that $d_i = 1$ iff $b_j = 1$, for all $j < i$ and that $e_i = 1$ iff $b_j = 0$, for all $j < i$.

$$
\begin{aligned}
\chi_4 &= \chi_{4a} \wedge \chi_{4b} \\
\chi_{4a} &= \mathrm{AG}\Big(d_0 \wedge \bigwedge_{i=1}^{n-1}[d_i \leftrightarrow (d_{i-1} \wedge b_{i-1})] \wedge e_0 \wedge \bigwedge_{i=1}^{n-1}[e_i \leftrightarrow (e_{i-1} \wedge \neg b_{i-1})]\Big) \\
\chi_{4b} &= \mathrm{AG}\Big[(o \wedge \neg\varphi_{last}) \rightarrow \\
&\qquad \Big({\downarrow}x.\mathrm{EX}(o \wedge [(e_{n-1} \wedge b_{n-1} \wedge @_x\neg b_{n-1}) \vee (\neg e_{n-1} \wedge (b_{n-1} \leftrightarrow @_x b_{n-1}))]\wedge \\
&\qquad\qquad \bigwedge_{i=0}^{n-2}[(e_{i+1} \wedge @_x d_{i+1}) \vee (e_i \wedge b_i \wedge @_x\neg b_i) \vee (\neg e_i \wedge (b_i \leftrightarrow @_x b_i))])\Big)\Big]
\end{aligned}
$$

T5: *Each position bit has a child, which represents a copy of it. The nodes of a subtree rooted at a copy node are labelled exactly with the same propositions.*

$$
\begin{aligned}
\chi_5 = \mathrm{AG}[&(o \rightarrow {\downarrow}x.\mathrm{EX}(c \wedge \bigwedge_{i=0}^{n-1}(b_i \leftrightarrow @_x b_i) \wedge b \leftrightarrow @_x b))\wedge \\
&(c \rightarrow {\downarrow}x.\mathrm{AG}(c \wedge \bigwedge_{i=0}^{n-1}(b_i \leftrightarrow @_x b_i) \wedge b \leftrightarrow @_x b))]
\end{aligned}
$$

T6: *Each position bit has exactly two children. One of them is its copy node and the other one is:*

(a) the next position bit if the current node is not already the last position bit or
(b) a node containing one of the propositions pos_e, pos_o, row_e, row_o and $q_\sharp$, otherwise.

$$\chi_6 = \chi_{6a} \wedge \chi_{6b}$$

$$\chi_{6a} = \mathrm{AG}[(o \wedge \neg\varphi_{last}) \rightarrow \mathrm{EX}(o \wedge {\downarrow}x.@_{root}\mathrm{EF}(\mathrm{EX}x \wedge \mathrm{AX}(o \rightarrow x))) \wedge$$
$$\mathrm{EX}(c \wedge {\downarrow}x.@_{root}\mathrm{EF}(\mathrm{EX}x \wedge \mathrm{AX}(c \rightarrow x))) \wedge$$
$$\mathrm{AX}(o \vee c)]$$

$$\chi_{6b} = \mathrm{AG}[(o \wedge \varphi_{last}) \rightarrow \mathrm{EX}(c \wedge {\downarrow}x.@_{root}\mathrm{EF}(\mathrm{EX}x \wedge \mathrm{AX}(c \rightarrow x))) \wedge$$
$$\mathrm{EX}(\neg c \wedge {\downarrow}x.@_{root}\mathrm{EF}(\mathrm{EX}x \wedge \mathrm{AX}(\neg c \rightarrow x))) \wedge$$
$$\mathrm{AX}(c \vee \varphi_{pos} \vee \varphi_{row} \vee q_\sharp)]$$

T7: *With the propositions row_e and row_o the counting of the positions of the current row starts. This means that the next position gets the initial position number. Therefore the proposition b is set to false in every position bit of this position.*

$$\chi_7 = \mathrm{AG}[\varphi_{row} \rightarrow \mathrm{EXAXE}[\psi_{full} \wedge \mathrm{G}\neg b]]$$

Compared to the increasing of a position bit number the increasing of a position number is a little bit complicated because in the latter case the bits are distributed over several nodes. In addition, we have to account for the case that player A cannot make a move without violating the horizontal or vertical constraints. In this case the game is over and $q_\sharp$-nodes follow only. We describe the increasing of a position number in three parts.

$$\chi_8 = \chi_{8a} \wedge \chi_{8b} \wedge \chi_{8c}$$

T8a: *If the last position is reached then a new row is started or the game is over.*

$$\chi_{8a} = \mathrm{AG}[\varphi_{first} \wedge \mathrm{E}[\psi_{full} \wedge \mathrm{G}b] \rightarrow \mathrm{E}[\mathrm{G}\neg\varphi_{pos} \wedge \mathrm{F}(o \wedge \varphi_{last} \wedge \mathrm{EX}(\varphi_{row} \vee q_\sharp))]]$$

T8b: *If the last position is not reached and it is the turn of player E then the next position sequence definitely has to be encoded and it gets the next position number.*

$$\chi_{8b} = \mathrm{AG}[\varphi_{first} \wedge \mathrm{E}[\mathrm{G}\neg\varphi_{pos} \wedge \mathrm{F}(c \wedge \varphi_{last} \wedge b)] \wedge \mathrm{E}[\psi_{full} \wedge \mathrm{F}\neg b] \rightarrow \theta]$$

Find the highest bit which has to be flipped.

$$\theta = \mathrm{E}[\psi_{full} \wedge \mathrm{F}(o \wedge \neg b \wedge \mathrm{EX}(\neg c \wedge (o \rightarrow \mathrm{E}[\psi_{full} \wedge \mathrm{G}b])) \wedge \theta')]$$

Flip the same bit in the next position sequence.

$$\theta' = {\downarrow}x.[\theta'' \wedge \mathrm{E}[\psi_{full} \wedge \mathrm{F}(o \wedge \varphi_{last} \wedge$$
$$\mathrm{EX}(\varphi_{pos} \wedge \mathrm{AXE}[\psi_{full} \wedge \mathrm{F}(o \wedge \bigwedge_{i=0}^{n-1}(b_i \leftrightarrow @_x b_i) \wedge (b \leftrightarrow @_x b) \wedge$$
$$\mathrm{EX}(\neg c \wedge (o \rightarrow \mathrm{E}[\psi_{full} \wedge \mathrm{G}\neg b])))])))]]$$

14

If the proposition b is satisfied in a position bit preceding the flipped bits then it is also satisfied in the same position bit in the next position sequence.

$$\theta'' = @_{\mathrm{root}} \mathrm{EF}[\varphi_{first} \wedge \mathrm{E}[\mathrm{G} \neg \varphi_{pos} \wedge \mathrm{F}x \wedge \mathrm{F}(c \wedge \bigwedge_{i=0}^{n-1}(b_i \leftrightarrow @_x b_i)) \wedge$$

$$\mathrm{G}(o \wedge \neg x \to \mathrm{A}[\psi_{2pos} \wedge \bigwedge_{i=0}^{n-1}(b_i \leftrightarrow \mathrm{F}(c \wedge b_i)) \to b \leftrightarrow \mathrm{F}(c \wedge b)])]]$$

T8c: *If the last position is not reached and it is the turn of player A then the next position sequence has not to be encoded but the game could be over.*

$$\chi_{8c} = \mathrm{AG}[\varphi_{first} \wedge \mathrm{E}[\mathrm{G} \neg \varphi_{pos} \wedge \mathrm{F}(c \wedge \varphi_{last} \wedge \neg b)] \to (\mathrm{E}[\psi_{full} \wedge \mathrm{F}(o \wedge \varphi_{last} \wedge \mathrm{EX}q_\sharp)] \vee \theta)]$$

T9: *Determining whether a position/row has an even or uneven number.*

$$\chi_9 = \chi_{9a} \wedge \chi_{9b}$$

$$\chi_{9a} = \mathrm{AG}[\varphi_{pos} \to \downarrow x. @_{\mathrm{root}} \mathrm{AG}[\varphi_{pos} \wedge \mathrm{EXE}[\psi_{full} \wedge \mathrm{F}(o \wedge \varphi_{last} \wedge \mathrm{EX}x)] \to$$
$$(pos_e \leftrightarrow @_x pos_o)]]$$

$$\chi_{9b} = \mathrm{AG}[\varphi_{row} \to \downarrow x. @_{\mathrm{root}} \mathrm{AG}[\varphi_{row} \wedge \mathrm{EXE}[\mathrm{G} \neg \varphi_{row} \wedge \mathrm{F}(c \wedge \varphi_{last}) \wedge$$
$$\mathrm{F}(o \wedge \varphi_{last} \wedge \mathrm{EX}x)] \to (row_e \leftrightarrow @_x row_o)]]$$

T10: *Each move of player A is followed by exactly one counter move of player E. Because the even positions correspond to the moves of player E, every position node labeled with pos_e must have exactly one child.*

$$\chi_{10} = \mathrm{AG}[pos_e \to \mathrm{EX} \downarrow x. @_{\mathrm{root}} \mathrm{EF}(\mathrm{EX}x \wedge \mathrm{AX}x)]$$

Now we use the tree structure described above to encode a winning strategy for player E.
W1: *The game ends after a finite sequence of moves.*
We express this by postulating that on every rootpath which does not contain any copy nodes a $q_\sharp$-node is eventually reached. A $q_\sharp$-node is followed only by $q_\sharp$-nodes.

$$\psi_1 = \mathrm{A}(\mathrm{G} \neg c \to \mathrm{F}q_\sharp) \wedge \mathrm{AG}(q_\sharp \to \mathrm{AG}q_\sharp)$$
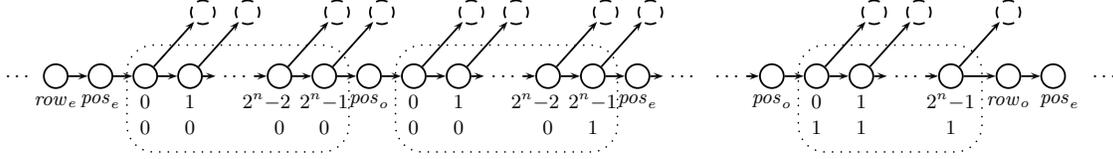
W2: *To each position belongs exactly one tile type.*
We remember that a tile type t is represented by the proposition p_t in all position bits of a position sequence.

$$\psi_2 = \mathrm{AG}[[o \to \bigvee_{t \in T}(p_t \wedge \bigwedge_{t \neq t' \in T} \neg p_{t'})] \wedge [o \wedge \neg \varphi_{last} \to \bigwedge_{t \in T}(t \leftrightarrow EX(o \wedge t))]]$$

W3: *According to the horizontal constraints, the tile type of every position, except the first one on each row, is consistent with the tile type of the precedent position.*

$$\psi_3 = \mathrm{AG}[\varphi_{first} \to$$
$$\bigwedge_{t' \in T}[p_{t'} \to \downarrow x. @_{\mathrm{root}} \mathrm{AG}[\varphi_{first} \wedge \neg x \wedge \mathrm{E}[\mathrm{F}x \wedge \psi_{2pos} \wedge \mathrm{G} \neg \varphi_{row}] \to \bigvee_{(t,t') \in H} p_t]]]$$

15

Fig. 4. A row sequence in the encoding of a winning strategy for player E. The upper nodes are copy nodes.

W4: *According to the vertical constraints, for every row, except the first one, it holds that the tile type of every position on this row is consistent with the tile type of the same position on the precedent row.*

$$\psi_4 = \mathrm{AG}[\varphi_{first} \to \bigwedge_{t' \in T} [p_{t'} \to \mathrm{E}[\psi_{full} \wedge \mathrm{F}(o \wedge \varphi_{last} \wedge \downarrow x.@_{\mathrm{root}} \mathrm{AG}[\xi \to \bigvee_{(t,t') \in V} p_t])]]]]$$

$$\xi = \varphi_{first} \wedge \mathrm{E}[\mathrm{F}x \wedge \mathrm{G}(\neg c \to \mathrm{EF}x) \wedge \psi_{2row}] \wedge$$
$$\mathrm{E}[\psi_{full} \wedge \mathrm{G}(o \to \mathrm{E}[\mathrm{G}(\neg c \to \mathrm{EF}x) \wedge \mathrm{FE}[\mathrm{G}\neg\varphi_{pos} \wedge \mathrm{F}x] \wedge$$
$$\bigwedge_{i=0}^{n-1} (b_i \leftrightarrow \mathrm{F}(c \wedge b_i)) \wedge b \leftrightarrow \mathrm{F}(c \wedge b)])]$$

W5: *All possible moves of player A are represented in the encoding.*

$$\psi_5 = \mathrm{AG}[\varphi_{first} \wedge \mathrm{E}[\psi_{full} \wedge \mathrm{F}(o \wedge \varphi_{last} \wedge \neg b)] \to$$
$$\bigwedge_{t \in T} (p_t \to \bigwedge_{(t,t') \in H} (\mathrm{E}[\psi_{full} \wedge \mathrm{F}(o \wedge \varphi_{last} \wedge \mathrm{EX}(\varphi_{pos} \wedge \mathrm{EX}p_{t'}))]) \vee$$
$$\mathrm{E}[\psi_{full} \wedge \mathrm{F}(o \wedge \varphi_{last} \wedge$$
$$\downarrow x.@_{\mathrm{root}} \mathrm{EF}(\xi \wedge$$
$$\bigvee_{(t'',t') \notin V} \mathrm{E}[\psi_{full} \wedge$$
$$\mathrm{F}(o \wedge \varphi_{last} \wedge \mathrm{EX}(p_{t''} \wedge \mathrm{EF}x))])])])))]$$

W6: *All tile types in the first row are from the set F.*

$$\psi_6 = \mathrm{AG}[\varphi_{first} \wedge \downarrow x.@_{\mathrm{root}} \mathrm{EXE}[\mathrm{G}\neg\varphi_{row} \wedge \mathrm{F}x] \to \bigvee_{t \in F} p_t]$$

W7: *Unless the game terminates prematurely (because player A is not able to make a further move) all tile types in the last row are from the set L.*

$$\psi_7 = \mathrm{AG}[o \wedge \varphi_{last} \wedge b \wedge \mathrm{EX}q_\sharp \to$$
$$\downarrow x.@_{\mathrm{root}} \mathrm{EF}[\varphi_{row} \wedge \mathrm{EXE}[\mathrm{G}\neg\varphi_{row} \wedge \mathrm{F}x \wedge \mathrm{G}(\varphi_{first} \to \bigvee_{t \in L} p_t)]]]]$$

Finally we obtain the formula φ_I as a conjunction of the formulas encoding the properties T1-T10 and W1-W7. It should be noticed that a model for φ_I can contain multiple possible moves for player E on each position. In this case it suffices to choose one of the suggested moves because every rootpath without copy nodes represents a win for E. □

We can obtain, by simple instantiation, a consequence of this lower complexity bound which will be useful later on in proving the exponential succinctness of $\mathrm{H}^1\mathrm{CTL}^+$ in $\mathrm{H}^1\mathrm{CTL}$.

Corollary 9. *There are finitely satisfiable* H^1CTL^+ *formulas* φ_n, $n \in \mathbb{N}$, *of size* $\mathcal{O}(n)$ *s.t. every tree model* $\mathcal{T}_n$ *of* φ_n *has height at least* $2^{2^{2^n}}$.

Proof. Consider the following instances of the 2EXP-tiling game: $I_n := (T, H, V, F, L, n)$ where $T = \{0, 1\} \times \{\mathtt{l}, \mathtt{f}, \mathtt{s}\}$. Tiles are supposed to model bit values in their first component. The second component describes whether or not a bit has to be *flipped* (value $\mathtt{f}$) or remains the *same* (value $\mathtt{s}$) in the increase of a value encoded in binary through these bits. The value $\mathtt{l}$ is used to mark the *lowest* bit in a number. Remember that in binary increase the lowest bit always gets flipped whereas the flipping of any other bit is determined by the value of itself, the value of the next lower bit and the question whether or not that bit gets flipped. We denote a tile as $0^\mathtt{l}$ for instance rather than $(0, \mathtt{l})$.

I_n is constructed in a way that forces both players to put down the number 0 in binary coding into the first row of the 2^{2^n}-arena and, whenever a row encodes a number i, then the players need to place the binary encoding of $i + 1$ into the next row. Thus, there will be (almost) no choices for the players. Player E should win when the highest possible number $2^{2^{2^n}}$ is placed in a row.

The starting constraints are $F := \{0^\mathtt{l}, 0^\mathtt{s}\}$. Hence, the first row must encode 0. The horizontal constraints are as follows.

$$H \;\; := \;\; \{(0^\mathtt{l}, b^\mathtt{s}), (1^\mathtt{l}, b^\mathtt{f}), (0^\mathtt{f}, b^\mathtt{s}), (1^\mathtt{f}, b^\mathtt{f}), (0^\mathtt{s}, b^\mathtt{s}), (1^\mathtt{s}, b^\mathtt{s})\}$$

where b is, in any case, an arbitrary value in $\{0, 1\}$.

Now note that with this H and F, there are only two possible first rows that the players can lay down: $0^\mathtt{l}0^\mathtt{s} \ldots 0^\mathtt{s}$ or $0^\mathtt{s} \ldots 0^\mathtt{s}$, and it is player E who determines entirely through his first choice which of these it is going to be.

Next we will translate the informal description of binary increase given above into the vertical constraints.

$$V \;\; := \;\; \{(0^\mathtt{l}, 1^\mathtt{l}), (1^\mathtt{l}, 0^\mathtt{l}), (0^\mathtt{f}, 1^x), (1^\mathtt{f}, 0^x), (0^\mathtt{s}, 0^x), (1^\mathtt{s}, 1^x)\}$$

where x is, in any case this time, an arbitrary value in $\{\mathtt{f}, \mathtt{s}\}$.

Now note that, if a row i is layed down entirely, then the vertical constraints determine uniquely the bit value of each tile in the next row. Furthermore, if the first tile in row i is of the form $b^\mathtt{l}$ then the first tile in row $i + 1$ is uniquely determined to be $(1 - b)^\mathtt{l}$, and H as well as the bit values in row $i + 1$ uniquely determine the flip-values of all the bits in row $i + 1$. Furthermore, all lowest bits that are all 1 in row i are 0 in row $i + 1$, the lowest bit that is 0 in row i is 1 in row $i + 1$, and all other bits in row $i + 1$ retain their value from row i. Hence, if row i encodes the number i in binary (starting with 0), then row $i + 1$ encodes the number $i + 1$ in binary.

Thus, if player E chooses tile $0^\mathtt{l}$ as the first one, then both players have no choice but to lay down the binary encodings of $0, 1, 2, \ldots$. On the other hand, if player E chooses tile $0^\mathtt{s}$ as the first one then all rows will encode the number 0 because the entire arena must be tiled with $0^\mathtt{s}$ only.

Finally, remember that the goal is to construct I_n in a way that enforces a play filling $2^{2^{2^n}}$ many rows. This can now easily be achieved by constuction the end constraints in a way that player E only wins when the number $2^{2^{2^n}} - 1$ has been placed down in a row. We therefore set $L := \{1^\mathtt{l}, 1^\mathtt{f}\}$. Note that in the successive increase as constructed above, the last row cannot contain the tile $1^\mathtt{s}$. □

Using the ideas of the transformation mentioned in Theorem 2 we can show that the lower bound for H^1CTL^+ is optimal. Even for strictly more expressive logics than H^1CTL^+ the satisfiability problem remains in **3EXPTIME**.

Theorem 10. *The satisfiability problem for* H^1CTL^+ *is* **3EXPTIME**-*complete.*

Proof. Let H^1PECTL^+ be the logic H^1CTL^+ augmented with the operators Y, S, $\overset{\infty}{F}$ and $\overset{\infty}{G}$ and similarly H^1PCTL be H^1CTL augmented with Y and S. We describe how to transform a H^1PECTL^+-formula φ into an satisfiability equivalent H^1PCTL-formula φ'. The transformation algorithm we use yields an exponential blowup in the formula length. As satisfiability of H^1PCTL is in **2EXPTIME** [28] we get the desired **3EXPTIME** upper bound for satisfiability of H^1PECTL^+.

The transformation algorithm uses the equivalences already used in the proof of Theorem 2 plus additional equivalences to deal with the extra operators. For convenience of the reader, we state all equivalences in the following.

(1) $\neg X\varphi \equiv X\neg\varphi$

(2) $\neg Y\varphi \equiv Y\neg\varphi$

(3) $\neg(\varphi U\varphi') \equiv [(\varphi \wedge \neg\varphi')U(\neg\varphi \wedge \neg\varphi')] \vee G\neg\varphi'$

(4) $\neg(\varphi S\varphi') \equiv (\varphi \wedge \neg\varphi')S(\neg\varphi \wedge \neg\varphi')$

(5) $\neg\overset{\infty}{G}\varphi \equiv \overset{\infty}{F}\neg\varphi$

(6) $E(\psi \vee \psi') \equiv E\psi \vee E\psi'$

(7) $X\varphi \wedge X\varphi' \equiv X(\varphi \wedge \varphi')$

(8) $Y\varphi \wedge Y\varphi' \equiv Y(\varphi \wedge \varphi')$

(9) $G\varphi \wedge G\varphi' \equiv G(\varphi \wedge \varphi')$

(10) $\overset{\infty}{G}\varphi \wedge \overset{\infty}{G}\varphi' \equiv \overset{\infty}{G}(\varphi \wedge \varphi')$

(11) Extraction of past operators

$$E[\bigwedge_{i=1}^{k} Y\varphi_i \wedge \bigwedge_{i=1}^{l}(\psi_i S\psi'_i) \wedge X\chi \wedge G\xi \wedge \overset{\infty}{G}\rho \wedge \bigwedge_{i=1}^{m}(\eta_i U\eta'_i) \wedge \bigwedge_{i=1}^{n}\overset{\infty}{F}\kappa_i \wedge \bigwedge_{i=1}^{o}\neg\overset{\infty}{F}\lambda_i]$$

$$\equiv$$

$$\bigwedge_{i=1}^{k} Y\varphi_i \wedge \bigwedge_{i=1}^{l}(\psi_i S\psi'_i) \wedge E[X\chi \wedge G\xi \wedge \overset{\infty}{G}\rho \wedge \bigwedge_{i=1}^{m}(\eta_i U\eta'_i) \wedge \bigwedge_{i=1}^{n}\overset{\infty}{F}\kappa_i \wedge \bigwedge_{i=1}^{o}\neg\overset{\infty}{F}\lambda_i]$$

(12) Elimination of the X-operator

$$E[X\varphi \wedge G\psi \wedge \overset{\infty}{G}\chi \wedge \bigwedge_{i=1}^{l}(\xi_i U\xi'_i) \wedge \bigwedge_{i=1}^{m}\overset{\infty}{F}\kappa_i \wedge \bigwedge_{i=1}^{n}\neg\overset{\infty}{F}\lambda_i]$$

$$\equiv$$

$$\bigvee_{I\subseteq\{1,\dots,l\}} [\bigwedge_{i\in I}\xi'_i \wedge \psi \wedge \bigwedge_{i\notin I}\xi_i \wedge EX(\varphi \wedge E[G\psi \wedge \overset{\infty}{G}\chi \bigwedge_{i\notin I}(\xi_i U\xi'_i) \wedge \bigwedge_{i=1}^{m}\overset{\infty}{F}\kappa_i \wedge \bigwedge_{i=1}^{n}\neg\overset{\infty}{F}\lambda_i])]$$

(13) Disjunction over all possible sequences in which the formulas ξ'_i with $1 \le i \le l$ can occur

$$E[G\psi \wedge \overset{\infty}{G}\chi \wedge \bigwedge_{i=1}^{l}(\xi_i U\xi'_i) \wedge \bigwedge_{i=1}^{m}\overset{\infty}{F}\kappa_i \wedge \bigwedge_{i=1}^{n}\neg\overset{\infty}{F}\lambda_i]$$

$$\equiv$$

$$\bigvee_{\pi\in Perm(\{1,\dots,n\})} [E[(\bigwedge_{i=1}^{n}\xi_i \wedge \psi)U(\xi'_{\pi(1)} \wedge E[(\bigwedge_{i\neq\pi(1)}\xi_i \wedge \psi)U(\xi'_{\pi(2)}\wedge$$

$$E[(\bigwedge_{i\neq\pi(1),\pi(2)}\xi_i \wedge \psi)U(\xi'_{\pi(3)} \wedge \dots U(\xi'_{\pi(n)} \wedge E[G\psi \wedge \overset{\infty}{G}\chi \wedge \bigwedge_{i=1}^{m}\overset{\infty}{F}\kappa_i \wedge \bigwedge_{i=1}^{n}\neg\overset{\infty}{F}\lambda_i])\dots)])])]]$$

(14) Elimination of the $\overset{\infty}{G}$-operator

$$E[G\psi \wedge \overset{\infty}{G}\chi \wedge \bigwedge_{i=1}^{m} \overset{\infty}{F}\kappa_i \wedge \bigwedge_{i=1}^{n} \neg\overset{\infty}{F}\lambda_i] \equiv E[\psi U(E[G(\psi \wedge \chi) \wedge \bigwedge_{i=1}^{m} \overset{\infty}{F}\kappa_i \wedge \bigwedge_{i=1}^{n} \neg\overset{\infty}{F}\lambda_i])]$$

(15) Elimination of $\neg\overset{\infty}{F}$

$$E[G\psi \wedge \bigwedge_{i=1}^{m} \overset{\infty}{F}\kappa_i \wedge \bigwedge_{i=1}^{n} \neg\overset{\infty}{F}\lambda_i] \equiv E[\psi U(E[G(\psi \wedge \bigwedge_{i=1}^{n} \neg\lambda_i) \wedge \bigwedge_{i=1}^{m} \overset{\infty}{F}\kappa_i])]$$

(16) Elimination of the $\overset{\infty}{F}$-operator

$$E[G\varphi \wedge \bigwedge_{i=1}^{m} \overset{\infty}{F}\kappa_i]$$

is satisfiable if and only if

$$\bigwedge_{i=1}^{m} AG(\neg EG(p_i \wedge \neg\kappa_i)) \wedge EG(\varphi \wedge \bigwedge_{i=1}^{m} p_i)$$

is satisfiable

Note that in (16) the two formulas are equivalent only with respect to satisfiability. For every formula κ_i, the new formula uses an additional proposition p_i which is supposed to hold on all paths satisfying $\overset{\infty}{F}\kappa_i$.

Let φ be a H^1PECTL^+-formula. We can assume that φ does not contain the path quantifier A because of $A\psi \equiv \neg E\neg\psi$. In a bottom up fashion the algorithm replaces each subformula $E\psi$ of φ by a H^1PCTL-formula. Thus, it only remains to describe how to transform a formula $E\psi$ where ψ is a boolean combination of path formulas of the form $Y\chi$, $\chi S\chi'$, $X\chi$, $\chi U\chi'$, $\overset{\infty}{F}\chi$ and $\overset{\infty}{G}\chi$ with $\chi \in H^1PCTL$. The transformation involves the following steps:

- Using De Morgan's laws the $\neg$-operators are pushed to the leaves of the Boolean combination.
- By applying equivalences (1)-(5) negations can be eliminated from the outermost Boolean combination and a formula $E\psi$ is obtained in which ψ is a positive Boolean combination of path formulas of the form $Y\chi$, $\chi S\chi'$, $X\chi$, $G\chi$, $\chi U\chi'$, $\overset{\infty}{F}\chi$, $\neg\overset{\infty}{F}\chi$ and $\overset{\infty}{G}\chi$
- By applying equivalences (6)-(10) the formula can be transformed into a formula of the form
  $E[\bigwedge_{i=1}^{k} Y\varphi_i \wedge \bigwedge_{i=1}^{l}(\psi_i S\psi_i') \wedge X\chi \wedge G\xi \wedge \overset{\infty}{G}\rho \wedge \bigwedge_{i=1}^{m}(\eta_i U\eta_i') \wedge \bigwedge_{i=1}^{n} \overset{\infty}{F}\kappa_i \wedge \bigwedge_{i=1}^{o} \neg\overset{\infty}{F}\lambda_i]$.
- Eventually, applying equivalences (11) - (16) a H^1PCTL-formula φ' is obtained which is satisfiable if and only if φ is satisfiable.

It can be shown that the factorial blowup in equivalence (7) is the worst blowup in the transformation algorithm [6]. As $n! = 2^{\mathcal{O}(n \log n)}$ we can conclude that $|\varphi'|$ is at most exponential in $|\varphi|$. □

5   The Succinctness of H^1CTL^+ w.r.t. H^1CTL

In Corollary 3 an upper bound of $2^{\mathcal{O}(n\log n)}$ for the succinctness of H^1CTL^+ in H^1CTL is given. In this section we establish the lower bound for the succinctness between the two logics. Actually we show that H^1CTL^+ is exponentially more succinct than H^1CTL. The model-theoretic approach we use in the proof is inspired by [17]. We first establish a kind of small model property for H^1CTL.

Theorem 11. *Every finitely satisfiable* H^1*CTL-formula* φ *with* $|\varphi| = n$ *has a model of depth* $2^{2^{\mathcal{O}(n)}}$.

Proof. In [28] it was shown that for every H^1CTL-formula φ, an equivalent nondeterministic Büchi tree automaton A_φ with $2^{2^{\mathcal{O}(|\varphi|)}}$ states can be constructed. It is easy to see by a pumping argument that if A_φ accepts some finite tree at all, it accepts one of depth $2^{2^{\mathcal{O}(|\varphi|)}}$. It should be noted that the construction in [28] only constructs an automaton that is equivalent to φ with respect to satisfiability. However, the only non-equivalent transformation step is from φ to a formula φ' without nested occurrences of the $\downarrow$-operator (Lemma 4.3 in [28]). It is easy to see that this step only affects the propositions of models but not their shape let alone depth. $\square$

Corollary 9 and Theorem 11 together immediately yield the following.

Corollary 12. H^1CTL$^+$ *is exponentially more succinct than* H^1CTL.

6 Conclusion

The aim of this paper is to contribute to the understanding of one-variable hybrid logics on trees, one of the extensions of temporal logics with reasonable complexity. We showed that H^1CTL$^+$ has no additional power over H^1CTL but is exponentially more succinct, we settled the complexity of H^1CTL$^+$ and showed that hybrid variables do not help in expressing fairness (as HCTL$^+$ cannot express EGFp).

However, we leave a couple of issues for further study, including the following.

- We conjecture that the succinctness gap between H^1CTL$^+$ and H^1CTL is actually $\theta(n)!$.
- We expect the HCTL-game to capture exactly the expressive power of HCTL. Remember that here we needed and showed only one part of this equivalence.
- The complexity of Model Checking for HCTL has to be explored thoroughly, on trees and on arbitrary transition systems. In this context, two possible semantics should be explored: the one, where variables are bound to nodes of the computation tree and the one which binds nodes to states of the transition system (the latter semantics makes the satisfiability problem undecidable on arbitrary transition systems [2])

References

1. C. Areces, P. Blackburn, and M. Marx. The computational complexity of hybrid temporal logics. *Logic Journal of the IGPL*, 8(5):653–679, 2000.
2. C. Areces, P. Blackburn, and M. Marx. Hybrid logics: Characterization, interpolation and complexity. *J. of Symbolic Logic*, 66(3):977–1010, 2001.
3. C. Areces and B. ten Cate. Hybrid logics. In *Handbook of Modal Logic*, volume 3 of *Studies in Logic*, pages 821–868. Elsevier, 2007.
4. B. S. Chlebus. Domino-tiling games. *J. Comput. Syst. Sci.*, 32(3):374–392, 1986.
5. E. M. Clarke and E. A. Emerson. Design and synthesis of synchronization skeletons using branching-time temporal logic. In *Proc. Logic of Programs*, volume 131 of *LNCS*, pages 52–71. Springer, 1981.
6. E. A. Emerson and J. Y. Halpern. Decision procedures and expressiveness in the temporal logic of branching time. In *Proceedings of the Fourteenth Annual ACM Symposium on Theory of Computing (STOC '82)*, pages 169–180. ACM, 1982.
7. E. A. Emerson and J. Y. Halpern. Decision procedures and expressiveness in the temporal logic of branching time. *J. Comput. Syst. Sci.*, 30(1):1–24, 1985.
8. E. A. Emerson and J. Y. Halpern. "Sometimes" and "not never" revisited: on branching versus linear time temporal logic. *J. ACM*, 33(1):151–178, 1986.
9. J. Engelfriet and H. J. Hoogeboom. Tree-walking pebble automata. In *Jewels are Forever*, pages 72–83. Springer, 1999.

10. M. Franceschet, M. de Rijke, and B.-H. Schlingloff. Hybrid logics on linear structures: Expressivity and complexity. In *Proc. of the 10th International Symposium on Temporal Representation and Reasoning / 4th International Conference on Temporal Logic (TIME-ICTL 2003)*, pages 192–202. IEEE, 2003.

11. V. Goranko. Temporal logics with reference pointers and computation tree logics. *Journal of Applied Non-Classical Logics*, 10(3-4), 2000.

12. M. Grohe and N. Schweikardt. The succinctness of first-order logic on linear orders. *Logical Methods in Computer Science*, 1(1), 2005.

13. J. Johannsen and M. Lange. CTL^+ is complete for double exponential time. In *Proc. of 30th ICALP 2003*, volume 2719 of *LNCS*, pages 767–775. Springer, 2003.

14. M. Jurdziński and R. Lazić. Alternation-free mu-calculus for data trees. In *Proc. of the 22th LICS 2007*. IEEE, 2007.

15. O. Kupferman and A. Pnueli. Once and for all. In *Proc. of the 10th LICS '95*, pages 25–35. IEEE, 1995.

16. O. Kupferman and M. Y. Vardi. Memoryful branching-time logic. In *Proc. of the 21st LICS 2006*, pages 265–274. IEEE, 2006.

17. M. Lange. A purely model-theoretic proof of the exponential succinctness gap between CTL^+ and CTL. *Information Processing Letters*, 108:308–312, 2008.

18. F. Laroussinie and P. Schnoebelen. A hierarchy of temporal logics with past. *Theor. Comput. Sci.*, 148(2):303–324, 1995.

19. F. Laroussinie and P. Schnoebelen. Specification in CTL+Past for verification in CTL. *Inf. Comput.*, 156(1-2):236–263, 2000.

20. L. Libkin. *Elements of Finite Model Theory*. Springer, 2004.

21. U. Sattler and M. Y. Vardi. The hybrid μ-calculus. In *Proc. of IJCAR2001*, volume 2083 of *LNCS*, pages 76–91. Springer, 2001.

22. T. Schwentick and V. Weber. Bounded-variable fragments of hybrid logics. In *Proc. of the 24th STACS 2007*, volume 4393 of *LNCS*, pages 561–572. Springer, 2007.

23. L. Segoufin. Automata and logics for words and trees over an infinite alphabet. In *Proc. of the 15th CSL 2006*, number 4207 in LNCS, pages 41–57. Springer, 2006.

24. B. ten Cate and L. Segoufin. XPath, transitive closure logic, and nested tree walking automata. In *Proc. of the 27th PODS 2008*, pages 251–260. ACM, 2008.

25. M. Y. Vardi. Alternating automata and program verification. In *Computer Science Today*, volume 1000 of *LNCS*, pages 471–485. Springer, 1995.

26. M. Y. Vardi. From church and prior to psl. In *25 Years of Model Checking*, volume 5000 of *LNCS*, pages 150–171. Springer, 2008.

27. M. Y. Vardi and L. J. Stockmeyer. Improved upper and lower bounds for modal logics of programs. In *Proc. of the 17th STOC '85*, pages 240–251. ACM, 1985.

28. V. Weber. Hybrid branching-time logics. *CoRR*, abs/0708.1723, 2007.

29. V. Weber. Branching-time logics repeatedly referring to states. Accepted to *JoLLI*, 2009. An extended abstract appeared in the procedings of HyLo 2007.
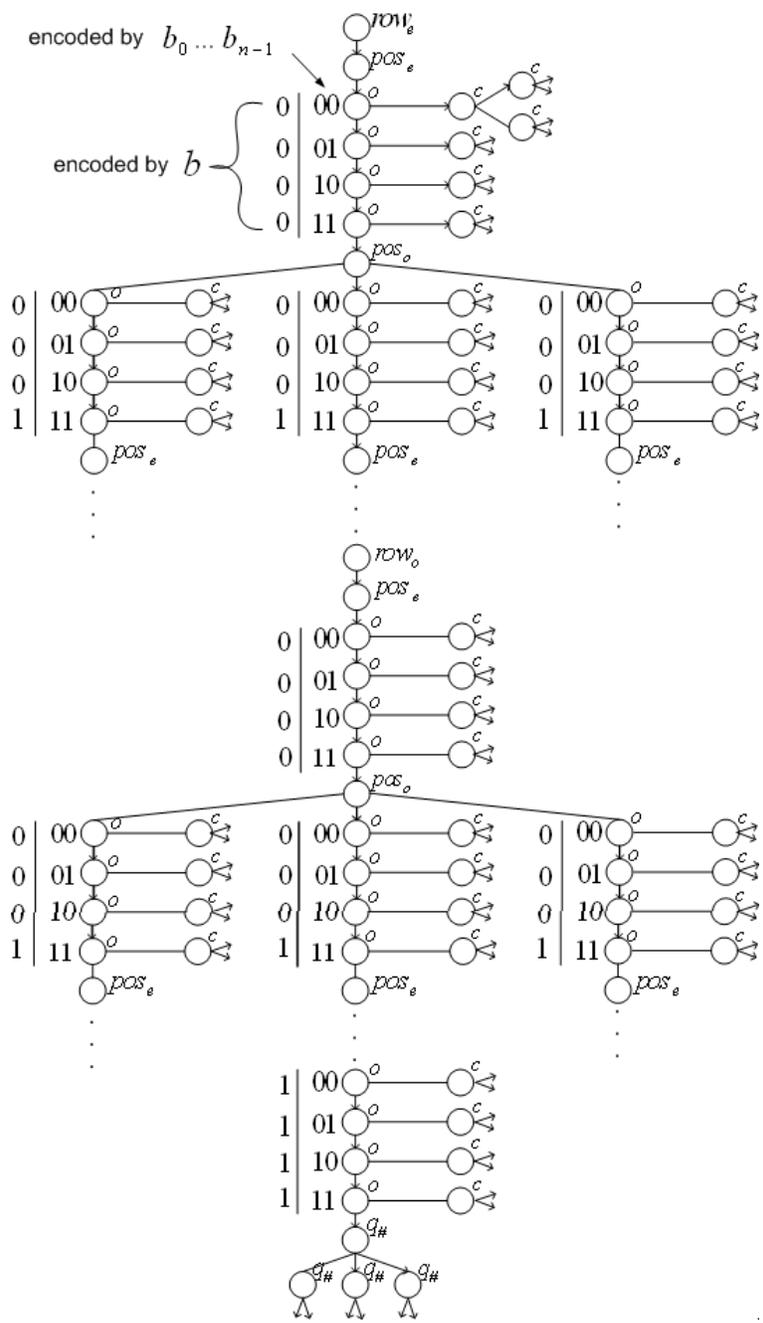
Fig. 5. An encoding of a winning strategy for player E in the case where $n = 2$.