Decidable and undecidable fragments of first-order branching temporal logics

Ian Hodkinson, Frank Wolter, and Michael Zakharyaschev

Department of Computing, Imperial College
180 Queen's Gate, London SW7 2BZ, U.K.;
Institut für Informatik, Universität Leipzig
Augustus-Platz 10-11, 04109 Leipzig, Germany;
Department of Computer Science, King's College
Strand, London WC2R 2LS, U.K.
(e-mails: imh@doc.ic.ac.uk, wolter@informatik.uni-leipzig.de, mz@dcs.kcl.ac.uk)

Contact author: Michael Zakharyaschev
Dept of Computer Science, King's College, Strand, London WC2R 2LS, UK
mz@dcs.kcl.ac.uk
Phone (+44) 20 7848 2725
Fax (+44) 20 7848 2851

Abstract

In this paper we analyze the decision problem for fragments of first-order extensions of branching time temporal logics such as computational tree logics CTL and CTL* or Prior's Ockhamist logic of historical necessity. On the one hand, we show that the one-variable fragments of logics like first-order $\mathcal{CTL}^*$—such as the product of propositional $\mathcal{CTL}^*$ with simple propositional modal logic S5, or even the one-variable bundled first-order temporal logic with sole temporal operator 'some time in the future'—are undecidable. On the other hand, it is proved that by restricting applications of first-order quantifiers to *state* (i.e., path-independent) formulas, and applications of temporal operators and path quantifiers to formulas with at most one free variable, we can obtain decidable fragments. The same arguments show decidability of 'non-local' propositional CTL*, in which truth values of propositional atoms depend on the history as well as the current time. The positive decidability results can serve as a unifying framework for devising expressive and effective time-dependent knowledge representation formalisms, e.g., temporal description or spatio-temporal logics.

1 Introduction

Temporal logics of branching time originate in philosophy and computer science. In philosophy, they formalize reasoning about indeterminate future; see e.g., [18, 19, 5, 20]. In computer science, they are used to reason about state transition systems (computations of reactive systems, agents in an unpredictable environment, etc.); see e.g., [6, 8, 16]. Families of different logics have been constructed in both disciplines, and in both cases most of the logics are *propositional* (the only exceptions we know of are [4, 22, 23]).

The main aim of this paper is to investigate the computational behavior—at least on the level of decidability—of first-order branching temporal logics (FOBTLs).

The starting point of our investigation lies in recent work [14] on first-order linear temporal logic (FOLTL), which showed that although very weak fragments (say, the monadic two-variable fragment) of standard FOLTLs can be highly undecidable, by restricting applications of temporal operators to formulas with at most one free variable, we obtain a fragment with much better computational behavior. This 'monodic' fragment can be axiomatized in a natural way [25] and becomes decidable if its 'purely first-order part' is restricted to a decidable fragment of first-order logic.

Of course, FOBTLs inherit bad computational properties of FOLTLs. For example, the guarded two-variable fragment of most FOBTLs is not even recursively enumerable; cf. [14]. Unfortunately (and to our surprise) it turned out that the situation is much worse. As will be shown in Section 3, the one-variable fragment of first-order $\mathcal{CTL}^*$ is undecidable. This is a monodic fragment with very simple decidable first-order part; it can be reformulated as the product of propositional $\mathcal{CTL}^*$ with propositional $\mathsf{S5}$. Even the one-variable fragment of the bundled FOBTL [5] with sole temporal operator 'now or some time in the future' is not decidable.

These 'negative' results become less surprising when we realize that first-order (bundled) $\mathcal{CTL}^*$ allows a form of quantification in three dimensions, something that is known to be associated with undecidability [17, 13]. A natural way to limit the interaction between the three dimensions is to restrict applications of first-order quantifiers to *state* (i.e., path-independent) formulas, and applications of temporal operators and path quantifiers to formulas with at most one free variable. The resulting fragment contains full propositional $\mathcal{CTL}^*$ and full first-order logic. In this paper, we show in a similar way to FOLTL that restricting its first-order part to a decidable first-order fragment yields a decidable monodic fragment of first-order $\mathcal{CTL}^*$. A number of 'positive' results of this sort will be proved in Section 4.

These results are interesting *per se*, and also because they identify a limit beyond which monodic fragments of temporal logics are no longer decidable. More practically, the positive results serve as a unifying framework for devising expressive and effective time-dependent knowledge representation formalisms, say, temporal description or spatio-temporal logics; for details see [15].

2 Syntax and semantics

There are a number of approaches to constructing temporal logics based on the branching time paradigm; see, e.g., [5, 7, 26, 10]. Many of the resulting languages are fragments of the language we call here $\mathcal{QPCTL}^*$, *quantified* $\mathcal{CTL}^*$ *with past operators*. It is obtained in the standard way by extending the language $\mathcal{PL}$ of classical predicate logic (without equality or function symbols) with the binary *temporal operators* S (Since) and U (Until) and the *path universal quantifier* A.

The intended models for $\mathcal{QPCTL}^*$ are based on ω-trees. Recall that a *tree* is a strict partial order $\mathfrak{F} = \langle W, < \rangle$ containing a unique $<$-minimal point (the *root* of $\mathfrak{F}$) and such that for all $w \in W$, the set

$\{v \in W : v < w\}$ is linearly ordered by $<$. A *full branch* of $\mathfrak{F}$ is a maximal linearly-ordered subset of W. An ω-*tree* is a tree whose full branches are all order-isomorphic to $\langle \mathbb{N}, < \rangle$.

$\mathcal{QPCTL}^*$ (as well as its sublanguages to be introduced below) is interpreted in structures of the form $\mathfrak{M} = \langle \mathfrak{F}, \mathcal{H}, D, I \rangle$, where $\mathfrak{F} = \langle W, < \rangle$ is an ω-tree, $\mathcal{H}$ is a set (*bundle*) of full branches of $\mathfrak{F}$ with $\bigcup \mathcal{H} = W$, D is a non-empty set called the *domain* of $\mathfrak{M}$, and I is a function associating with every moment of time $w \in W$ a first-order $\mathcal{PL}$-structure

$$I(w) = \left\langle D, P_0^{I(w)}, \dots, c_0^{I(w)}, \dots \right\rangle,$$

the *state* of $\mathfrak{M}$ at moment w. (Here, the $P_i^{I(w)}$ are predicates on D interpreting the predicate symbols P_i of $\mathcal{PL}$ and the $c_i^{I(w)}$ are elements of D interpreting its constants.) We require that $c_i^{I(w)} = c_i^{I(v)}$ for any $w, v \in W$—i.e., that constants are 'rigid'. The branches in the bundle $\mathcal{H}$ are called *histories*. If $\mathcal{H}$ contains all full branches of $\mathfrak{F}$, we say that $\mathfrak{M}$ is a *full tree model*, or simply a *tree model*.

An *assignment in D* is a function $\mathfrak{a}$ from the set of individual variables to D. (So assignments too are 'rigid'. We extend $\mathfrak{a}$ to constants via $\mathfrak{a}(c) = c^{I(w)}$ for any $w \in W$. The *truth relation* $(\mathfrak{M}, h, w) \models^{\mathfrak{a}} \varphi$, for $w \in h \in \mathcal{H}$ (or simply $(h, w) \models^{\mathfrak{a}} \varphi$ if $\mathfrak{M}$ is understood), is defined as follows:

- $(h,w) \models^{\mathfrak{a}} P_i(y_1, \dots, y_\ell)$ iff $I(w) \models P_i(\mathfrak{a}(y_1), \dots, \mathfrak{a}(y_\ell))$, where the y_i are variables or constants (this is the so-called 'local' approach in that there is no dependence on h),
- $(h,w) \models^{\mathfrak{a}} \forall x \psi$ iff $(h,w) \models^{\mathfrak{b}} \psi$ for every assignment $\mathfrak{b}$ in D that may differ from $\mathfrak{a}$ only on x,
- $(h,w) \models^{\mathfrak{a}} \chi \mathsf{S} \psi$ iff there is $v < w$ such that $(h,v) \models^{\mathfrak{a}} \psi$ and $(h,u) \models^{\mathfrak{a}} \chi$ for every $u \in (v,w)$, where $(v,w) = \{u \in W : v < u < w\}$,
- $(h,w) \models^{\mathfrak{a}} \chi \mathsf{U} \psi$ iff there is $v \in h$ such that $v > w$, $(h,v) \models^{\mathfrak{a}} \psi$ and $(h,u) \models^{\mathfrak{a}} \chi$ for every $u \in (w,v)$,
- $(h,w) \models^{\mathfrak{a}} \mathsf{A} \psi$ iff $(h',w) \models^{\mathfrak{a}} \psi$ for all $h' \in \mathcal{H}$ such that $w \in h'$,

plus the usual clauses for the Booleans. For a formula $\varphi(\bar{x})$ and a tuple $\bar{a}$ of elements of D, we write $(\mathfrak{M}, h, w) \models \varphi(\bar{a})$ if $(\mathfrak{M}, h, w) \models^{\mathfrak{a}} \varphi$ where $\mathfrak{a}(\bar{x}) = \bar{a}$. We will use the following standard abbreviations:

$$\Diamond_F \varphi = \top \mathsf{U} \varphi,\ \Box_F \varphi = \neg \Diamond_F \neg \varphi,\ \Box_F^+ \varphi = \varphi \wedge \Box_F \varphi,\ \Diamond_F^+ \varphi = \varphi \vee \Diamond_F \varphi, \bigcirc \varphi = \bot \mathsf{U} \varphi,\ \mathsf{E} \varphi = \neg \mathsf{A} \neg \varphi.$$

Thus, $\Diamond_F$ can be read as 'some time in the future', $\Box_F^+$ as 'from now on', $\bigcirc$ as 'at the next moment' or 'tomorrow', and E as 'there exists a history in $\mathcal{H}$'.

We will also be considering the following sublanguages of $\mathcal{QPCTL}^*$:

- $\mathcal{QCTL}^*$—that is, $\mathcal{QPCTL}^*$ without the past temporal operator S;
- $\mathcal{QCTL}_F^*$—that is, $\mathcal{QCTL}^*$ in which the binary operator U is replaced by $\Box_F^+$;
- $\mathcal{QPCTL}$—the fragment of $\mathcal{QPCTL}^*$ in which temporal operators occur only in the form $\mathsf{E}(\psi_1 \mathsf{U} \psi_2)$, $\mathsf{E}(\psi_1 \mathsf{S} \psi_2)$, $\mathsf{A}(\psi_1 \mathsf{U} \psi_2)$, or $\mathsf{A}(\psi_1 \mathsf{S} \psi_2)$.

To introduce one more fragment, $\mathcal{QPCTL}^s$, we need the definition of *state* and *path formulas*:

- all formulas without path quantifiers and temporal operators are state formulas;
- the set of state formulas is closed under the Booleans, path and first-order quantifications;
- every state formula is a path formula;

- the set of path formulas is closed under the Booleans and the temporal operators;
- if ψ is a path formula, then $\mathsf{A}\psi$ is a state formula.

So for a state formula ψ, whether $(h,w) \models^{\mathfrak{a}} \psi$ does not depend on h. Now, $\mathcal{QPCTL}^s$ consists of all state formulas in $\mathcal{QPCTL}^*$. The main difference between full $\mathcal{QPCTL}^*$ and $\mathcal{QPCTL}^s$ is that first-order quantifiers in $\mathcal{QPCTL}^s$ can be applied only to formulas which are history-independent. (That $\mathcal{QPCTL}^s$ only contains state formulas is not important for decidability, since a path formula φ is satisfiable iff the state formula $\mathsf{E}\varphi$ is satisfiable.) It should be clear that $\mathcal{QPCTL}^s$ contains propositional $\mathcal{PCTL}^*$ and $\mathcal{QPCTL}$. As examples, we give three formulas in the various fragments, trying to express that every ordered item is delivered in one day.

	$\mathcal{QPCTL}^*$	$\mathcal{QPCTL}^s$	$\mathcal{QPCTL}$
$\mathsf{E}\Box_F\forall x(\mathsf{ordered\_item}(x) \to \bigcirc\mathsf{delivered\_item}(x))$	✓	×	×
$\forall x\mathsf{E}\Box_F(\mathsf{ordered\_item}(x) \to \bigcirc\mathsf{delivered\_item}(x))$	✓	✓	×
$\mathsf{E}\Box_F\forall x(\mathsf{ordered\_item}(x) \to \mathsf{A}\bigcirc\mathsf{delivered\_item}(x))$	✓	✓	✓

For any of the languages $\mathcal{L}$ introduced above, denote by BL (respectively, L) the set of all $\mathcal{L}$-formulas that are true at all points in all histories under every assignment in every bundled (respectively, full) tree model. Such formulas are said to be *valid*, and a formula is *satisfiable* if its negation is not valid. Thus, BQPCTL^* is the set of $\mathcal{QPCTL}^*$-formulas valid in bundled tree models, while QPCTL^* is the set of $\mathcal{QPCTL}^*$-formulas valid in all tree models. The logic BQPCTL^* is the first-order version of the bundled Ockhamist logic of historical necessity; cf. e.g. [10].

$\mathcal{QPCTL}^*$-satisfiability in bundled tree models can be reduced to satisfiability in full tree models. Indeed, given a $\mathcal{QPCTL}^*$-formula φ, we take a propositional variable q not occurring in φ and denote by $\varphi^\uparrow$ the result of replacing each subformula of φ of the form $\mathsf{A}\psi$ by $\mathsf{A}(\Diamond_F\Box_F q \to \psi)$. Note that if φ is in $\mathcal{QPCTL}^s$ then so is $\varphi^\uparrow$ (however, $\varphi^\uparrow \notin \mathcal{QPCTL}$).

LEMMA 1. *φ is satisfiable in a bundled tree model iff $\mathsf{E}\Diamond_F\Box_F q \wedge \varphi^\uparrow$ is satisfiable in a full tree model.*

Proof. The implication $(\Leftarrow)$ is easily seen. We prove $(\Rightarrow)$. Using a Löwenheim–Skolem argument (cf. [5]), we may assume φ to be satisfied in a model $\mathfrak{M}$ with a countable bundle $\mathcal{H}$. We assume that $\mathcal{H}$ is infinite, leaving the (easy) other case to the reader. Let $h_0, h_1, \ldots$ be an enumeration of $\mathcal{H}$. We convert $\mathfrak{M}$ into a full tree model $\mathfrak{M}^\dagger$ and define a truth-relation for q in it inductively as follows: Put $(h_0,w) \models q$ for all $w \in h_0$. Suppose we have already defined truth of q in (h_i,w), for all $i \leq n$. Consider h_{n+1}. There must be a $w \in h_{n+1}$ such that the distance from w to each h_i, $i < n+1$, is ≥ 2 (the distance is the length of the shortest path from w to a point in h_i). Then we put $(h_{n+1},w') \models q$ if $w' \geq w$ and $w' \in h_{n+1}$.

Say that a full branch h of $\mathfrak{M}$ is *marked* if there is $w \in h$ such that $(h,w') \models q$ for all $w' \geq w$, $w' \in h$. One can easily see that h is marked iff $h \in \mathcal{H}$. In particular, if $h \notin \mathcal{H}$ and for each n, t_n is the least element of $h \setminus \bigcup_{m<n} h_m$, then $(h,t_n) \not\models q$ and $\{t_0,t_1,\ldots\}$ is infinite, so h is not marked. Now one can prove by induction that for every subformula ψ of φ and every (h,w), we have $(\mathfrak{M},h,w) \models \psi$ iff $(\mathfrak{M}^\dagger,h,w) \models \psi^\uparrow$. It follows that $\mathsf{E}\Diamond_F\Box_F q \wedge \varphi^\uparrow$ is satisfied in $\mathfrak{M}^\dagger$. ☐

3 Undecidable fragments

The following theorems indicate some limits beyond which one cannot hope to find decidable fragments of first-order temporal logics.

Given a first-order temporal language $\mathcal{L}$ and $\ell < \omega$, we denote by $\mathcal{L}^\ell$ the *ℓ-variable fragment* of $\mathcal{L}$ (i.e., every formula in $\mathcal{L}^\ell$ contains at most ℓ distinct individual variables). And by $\mathcal{L}^{mo}$ we denote the *monadic fragment* of $\mathcal{L}$ (i.e., the set of formulas which contain only unary predicates and propositional variables). Both the two-variable and the monadic fragments of classical (non-temporal) first-order logic are known to be decidable and have the finite model property; see [3] and references therein. The computational behavior of the corresponding fragments of first-order temporal logics turns out to be quite different. From linear time results (Theorem 2 of [14]) we easily obtain:

THEOREM 2. *For any of the FOBTLs* L *introduced above,* $\mathsf{L} \cap \mathcal{L}^2 \cap \mathcal{L}^{mo}$ *is not recursively enumerable.*

Another well-behaved fragment of classical predicate logic is the *guarded fragment* of [1]. The corresponding fragment $\mathcal{TGF}$ of FOBTL is obtained by restricting the $\mathcal{QPCTL}^*$ quantification formation rule to:

- if $\bar{x}, \bar{y}$ are tuples of variables, $G(\bar{x}, \bar{y})$ is an atomic formula, $\varphi(\bar{x}, \bar{y}) \in \mathcal{TGF}$, and every free variable occurring in $\varphi(\bar{x}, \bar{y})$ occurs in $G(\bar{x}, \bar{y})$ as well, then $\forall \bar{y}\, (G(\bar{x}, \bar{y}) \to \varphi(\bar{x}, \bar{y}))$ is in $\mathcal{TGF}$.

The set $\mathcal{TGF}$ is called the *guarded fragment* of FOBTL. Again in contrast to the case of classical predicate logic, we have the following consequence of Theorem 73 of [14]:

THEOREM 3. *For any of the FOBTLs* L *above,* $\mathsf{L} \cap \mathcal{L}^2 \cap \mathcal{TGF}$ *is not recursively enumerable.*

The main result of this section is the following:

THEOREM 4. *The one-variable fragments of the logics* BQCTL^*_F *and* QCTL^*_F *are undecidable. Hence, so are the one-variable fragments of* QPCTL^*, QCTL^*, BQPCTL^*, *and* BQCTL^*.

Proof. We only consider BQCTL^*_F; the other case follows from Lemma 1. The proof involves a rather indirect reduction of the following tiling problem. An instance is a finite set $\mathcal{T}$ of square tiles, the edges of each $\tau \in \mathcal{T}$ being colored $\mathit{Left}(\tau), \mathit{Right}(\tau), \mathit{Up}(\tau), \mathit{Down}(\tau)$ (see Fig. 2 in Appendix).[1] $\mathcal{T}$ is a yes-instance of the problem iff for each $\tau \in \mathcal{T}$ there is a $\mathcal{T}$-tiling of $\mathbb{Z} \times \mathbb{Z}$ that uses τ—a map $f : \mathbb{Z} \times \mathbb{Z} \to \mathcal{T}$ with $\mathit{Right}(f(i,j)) = \mathit{Left}(f(i+1,j))$, $\mathit{Up}(f(i,j)) = \mathit{Down}(f(i,j+1))$ for all $i, j \in \mathbb{Z}$, and (without loss of generality) $f(0,0) = \tau$. It is well known that this tiling problem is undecidable—see, e.g., [2].

To code this problem into BQCTL^*_F, we go via a construction originating in algebraic logic, which has already been used to prove undecidability of 3-dimensional modal logics [13]. It was shown in [12] that there is no algorithm to decide whether a finite Tarskian relation algebra is representable. We do not wish to assume any knowledge of algebraic logic here, so we will state some consequences of this result without using terms from that field. Given a tiling instance $\mathcal{T}$, we can effectively construct a certain finite signature $L(\mathcal{T})$ of binary relation symbols, and a set $\mathcal{C} \subseteq L(\mathcal{T})^3$ satisfying $(p,q,r) \in \mathcal{C} \Rightarrow (q,r,p) \in \mathcal{C}$ (and with further properties to be explained later). We call $\mathcal{C}$ the set of *consistent triples*. Results in [12] imply that if $\mathcal{T}$ is a yes-instance then there exists a countable $L(\mathcal{T})$-structure N satisfying:

- $\forall xy \big(\bigvee_{p \in L(\mathcal{T})} p(x,y) \wedge \neg \bigvee_{q \in L(\mathcal{T}) \setminus \{p\}} q(x,y) \big) \ \wedge \ \bigwedge_{p \in L(\mathcal{T})} \exists xy\, p(x,y)$,
- $\forall xy \big(\big(\bigvee_{(p,q,r) \in \mathcal{C}} p(x,y) \big) \leftrightarrow \exists z \big(q(y,z) \wedge r(z,x) \big) \big)$, for all $q, r \in L(\mathcal{T})$.

Thus, each pair of points in N is labelled by a unique $p \in L(\mathcal{T})$; only consistently-labelled triangles occur in N, but these must occur wherever possible. We will write a one-variable $\mathcal{QCTL}^*_F$-sentence $\widehat{\varphi}_{\mathcal{T}}$, essentially the Yankov–Fine formula of a frame derived from $\mathcal{C}$, that has a bundled model iff $\mathcal{T}$ is a yes-instance. This will establish undecidability of BQCTL^*_F.

[1] Strictly, we also require that there is a 'white' tile in $\mathcal{T}$ whose sides are all colored white; no other tile has any white edges. This is for technical reasons in [12] not relevant here.

In $\mathcal{QCTL}_F^*$, truth values of atomic formulas are history-independent (cf. Section 2). This is inconvenient here, so we temporarily drop this 'locality' restriction. We introduce three unary relation symbols p_0, p_1, p_2 for each $p \in L(\mathcal{T})$. Let $\varphi_{\mathcal{T}}$ be the conjunction of

$$\begin{array}{l}
\mathsf{A}\Box_F^+\forall x\big(\bigvee_p p_0(x) \wedge \neg \bigvee_{q\neq p} q_0(x)\big) \wedge \bigwedge_{p\in L(\mathcal{T})} \mathsf{E}\Diamond_F^+\exists x\, p_0(x), \quad \text{and similarly for } p_1, p_2 \\
\mathsf{A}\Box_F^+\forall x\big((\bigvee_{(p,q,r)\in\mathcal{C}} p_0(x)) \leftrightarrow \exists x(q_1(x)\wedge r_2(x))\big) \text{ for all } q,r \in L(\mathcal{T}), \\
\mathsf{A}\Box_F^+\forall x\big((\bigvee_{(p,q,r)\in\mathcal{C}} q_1(x)) \leftrightarrow \Diamond_F^+(r_2(x)\wedge p_0(x))\big) \text{ for all } p,r \in L(\mathcal{T}), \\
\mathsf{A}\Box_F^+\forall x\big((\bigvee_{(p,q,r)\in\mathcal{C}} r_2(x)) \leftrightarrow \mathsf{E}(p_0(x)\wedge q_1(x))\big) \text{ for all } p,q \in L(\mathcal{T}).
\end{array}$$

If $\mathcal{T}$ is a yes-instance, take N as above. We identify (notationally) N with its domain. Choose an ω-tree $\mathfrak{F} = \langle W, < \rangle$ with a countably infinite set $\mathcal{H}$ of full branches such that $\bigcup \mathcal{H} = W$. Let π_0 be the identity map on N. Choose $\mu : \omega \to N$ such that $\mu^{-1}(a)$ is infinite for all $a \in N$, and define $\pi_1 : W \to N$ by $\pi_1(w) = \mu(ht(w))$, where $ht(w) = |\{v \in W : v < w\}|$. Choose $\pi_2 : \mathcal{H} \to N$ such that $\bigcup(\pi_2^{-1}(a)) = W$ for all $a \in N$. Now we can define values for the p_{ij} in a model $\mathfrak{N}$ with domain N, underlying tree $\mathfrak{F}$, and bundle $\mathcal{H}$, by pulling back from N. For $a \in N$, $h \in \mathcal{H}$, $w \in h$, and $p,q,r \in L(\mathcal{T})$, we let

$$\begin{array}{lll}
(\mathfrak{N},h,w) \models p_0(a) & \text{iff} & N \models p(\pi_1(w), \pi_2(h)), \\
(\mathfrak{N},h,w) \models q_1(a) & \text{iff} & N \models q(\pi_2(h), \pi_0(a)), \\
(\mathfrak{N},h,w) \models r_2(a) & \text{iff} & N \models r(\pi_0(a), \pi_1(w)).
\end{array}$$

It can be checked that $\varphi_{\mathcal{T}}$ is true in $\mathfrak{N}$ at the root.

We now show how to code $\mathfrak{N}$ into an official $\mathcal{QCTL}_F^*$-model with 'local' history-independent atoms. We exploit the fact that the values of the atoms in $\mathfrak{N}$ depend on only two coordinates. r_2 depends on a, w, which is acceptable already. p_0 depends on w, h (actually, on $\pi_2(h)$) but not a; we can express this by coding h as $\pi_2(h)$, so reducing the dependence to w and a domain point. And q_1 depends on h, a but not w; by regarding q_1 as true at h, a iff it is in fact true at a, w for cofinitely many $w \in h$, we again reduce the dependence to w, a.

In detail, we translate $\varphi_{\mathcal{T}}$ into a $\mathcal{QCTL}_F^*$-sentence $\widehat{\varphi}_{\mathcal{T}}$ of the signature of $\varphi_{\mathcal{T}}$ plus a unary relation symbol Q. For $p,q,r \in L(\mathcal{T})$, we define formulas

$$\begin{array}{lcl}
\widehat{p_0}(x) & = & \exists x(p_0(x) \wedge \Diamond_F^+\Box_F^+ Q(x)) \\
\widehat{q_1}(x) & = & \Diamond_F^+\Box_F^+ q_1(x) \\
\widehat{r_2}(x) & = & r_2(x)
\end{array}$$

and let $\widehat{\varphi}_{\mathcal{T}}$ be the result of replacing each subformula $p_0(x)$ of $\varphi_{\mathcal{T}}$ by $\widehat{p_0}(x)$, and similarly for q_1, r_2. We form a $\mathcal{QCTL}_F^*$-structure $\mathfrak{M}$ with the same tree, bundle, and domain as $\mathfrak{N}$. Enumerate $\mathcal{H}$ as $\{h_0, h_1, \ldots\}$, and for $n < \omega$ let $h_n^- = h_n \setminus \bigcup_{m<n} h_m$, a cofinite subset of h_n. For $h \in \mathcal{H}$, $w \in h$, $a \in N$, and $p,q,r \in L(\mathcal{T})$, define

- $(\mathfrak{M},h,w) \models Q(a)$ iff there is $g \in \mathcal{H}$ with $w \in g^-$ and $a = \pi_2(g)$—so $(\mathfrak{M},h,w) \models \Diamond_F^+\Box_F^+ Q(a)$ iff $a = \pi_2(h)$,
- $(\mathfrak{M},h,w) \models p_0(a)$ iff there is $g \in \mathcal{H}$ with $w \in g$, $(\mathfrak{N},g,w) \models p_0(a)$, and $\pi_2(g) = a$,
- $(\mathfrak{M},h,w) \models q_1(a)$ iff there is $g \in \mathcal{H}$ with $w \in g^-$ and $(\mathfrak{N},g,w) \models q_1(a)$,
- $(\mathfrak{M},h,w) \models r_2(a)$ iff $(\mathfrak{N},h,w) \models r_2(a)$.

These values do not depend on h, so $\mathfrak{M}$ is a bona fide $\mathcal{QCTL}_F^*$-structure. It can now be checked that $(\mathfrak{M},h,w) \models \widehat{p_0}(a)$ iff $(\mathfrak{N},h,w) \models p_0(a)$, and similarly for $\widehat{q_1}, \widehat{r_2}$. Hence, $\widehat{\varphi}_{\mathcal{T}}$ is true in $\mathfrak{M}$.

For the converse, we will need more details about $L(\mathcal{T})$ and $\mathcal{C}$ from [12]. There is a 'converse' map $p \mapsto \breve{p}$ defined on $L(\mathcal{T})$, with $\breve{\breve{p}} = p$ for all p. $L(\mathcal{T})$ contains a relation symbol τ for each tile $\tau \in \mathcal{T}$ (we identify the two, so that $\mathcal{T} \subseteq L(\mathcal{T})$), the converse $\breve{\tau}$ of each tile τ, and also relation symbols $\mathsf{e}_0, \mathsf{e}_1, \mathsf{e}_2, +1_1, -1_1, +1_2, -1_2, \mathsf{g}_{01}, \mathsf{g}_{10}, \mathsf{g}_{02}, \mathsf{g}_{20}$ (and others that do not concern us here), all distinct. Each $p \in L(\mathcal{T})$ has a start index $st(p)$ and an end index $end(p)$. The start and end indices of g_{ij} are i, j, respectively; the start and end indices of $\mathsf{e}_i, +1_i$, and -1_i are i, and the tiles have start index 1 and end index 2. We have $\breve{g_{ij}} = g_{ji}$, $\breve{\mathsf{e}}_i = \mathsf{e}_i$, and $(+1_i)\breve{} = -1_i$. $\mathcal{C}$ satisfies:

- If $(p,q,r) \in \mathcal{C}$ then $end(p) = st(q)$, $end(q) = st(r)$, $end(r) = st(p)$, and $(q,r,p), (\breve{r}, \breve{q}, \breve{p}) \in \mathcal{C}$.
- For all $p, q \in L(\mathcal{T})$, $p = \breve{q}$ iff $(p,q,\mathsf{e}_i) \in \mathcal{C}$ for some $i < 3$.
- For all $p \in L(\mathcal{T})$, $(\mathsf{g}_{20}, \mathsf{g}_{01}, p) \in \mathcal{C}$ iff $p \in \mathcal{T}$.
- $(\mathsf{g}_{10}, \mathsf{g}_{01}, -1_1), (\mathsf{g}_{20}, \mathsf{g}_{02}, +1_2) \in \mathcal{C}$.
- for $\tau, \upsilon \in \mathcal{T}$, $(\breve{\upsilon}, +1_1, \tau) \in \mathcal{C}$ iff $Right(\upsilon) = Left(\tau)$, and $(\tau, -1_2, \breve{\upsilon}) \in \mathcal{C}$ iff $Down(\tau) = Up(\upsilon)$.

Now assume that $\widehat{\varphi}_{\mathcal{T}}$ is true (at the root, say) in some $\mathcal{QCTL}^*_F$-model $\mathfrak{M}$ with domain D, tree $\mathfrak{F} = \langle W, < \rangle$, and bundle $\mathcal{H}$. We will sketch a proof that $\mathcal{T}$ is a yes-instance. Let $\tau^{00} \in \mathcal{T}$ be given; we will read off from $\mathfrak{M}$ a tiling with τ^{00} at the origin. As notation, for $p, q, r \in L(\mathcal{T})$, $a \in D$, $h \in \mathcal{H}$, and $w \in h$, we write $(a,w,h) \models pqr$ iff $(\mathfrak{M}, h, w) \models \widehat{p}_0(a) \wedge \widehat{q}_1(a) \wedge \widehat{r}_2(a)$. It follows from the form of $\widehat{p}_0$ and of $\varphi_{\mathcal{T}}$ that

- for all a, w, h, there are unique $p, q, r \in L(\mathcal{T})$ with $(a,w,h) \models pqr$; moreover, $(p,q,r) \in \mathcal{C}$,
- for all $(p,q,r) \in \mathcal{C}$ there are a, h, w with $(a,w,h) \models pqr$,
- whenever $(a,w,h) \models pqr$ and $(a',w,h) \models p'q'r'$, we have $p = p'$,
- whenever $(a,w,h) \models pqr$ and $(p,q',r') \in \mathcal{C}$, there is $a' \in D$ with $(a',w,h) \models pq'r'$.

Analogous facts hold for q_1, r_2. We use these facts below without explicit mention.

As $\widehat{\varphi}_{\mathcal{T}}$ is true in $\mathfrak{M}$, and $(\mathsf{g}_{20}, \mathsf{g}_{01}, \tau^{00}) \in \mathcal{C}$, there are $x_0 \in D$, $h \in \mathcal{H}$, and $y_0 \in h$ with $(x_0, y_0, h) \models \mathsf{g}_{20}\mathsf{g}_{01}\tau^{00}$. From this, and since $(\mathsf{g}_{20}, \mathsf{g}_{02}, \mathsf{e}_2) \in \mathcal{C}$, there must be $u_0 \in D$ with $(u_0, y_0, h) \models \mathsf{g}_{20}\mathsf{g}_{02}\mathsf{e}_2$. Similarly, $(\mathsf{g}_{10}, \mathsf{g}_{01}, \mathsf{e}_1) \in \mathcal{C}$, so there is $v_0 \in h$ with $y_0 \leq v_0$ and $(x_0, v_0, h) \models \mathsf{g}_{10}\mathsf{g}_{01}\mathsf{e}_1$. In this way, by induction, for each $i \in \mathbb{Z}$ we can find $x_i, u_i \in D$ and $y_0 \leq y_i, v_i \in h$ with

- $(x_i, v_i, h) \models \mathsf{g}_{10}\mathsf{g}_{01}\mathsf{e}_1$ and $(x_{i+1}, v_i, h) \models \mathsf{g}_{10}\mathsf{g}_{01}{-1_1}$,
- $(u_i, y_i, h) \models \mathsf{g}_{20}\mathsf{g}_{02}\mathsf{e}_2$ and $(u_i, y_{i+1}, h) \models \mathsf{g}_{20}\mathsf{g}_{02}{+1_2}$.

See Fig. 3 in Appendix.

Consider now (x_i, y_j, h) $(i, j \in \mathbb{Z})$. We have $(x_i, y_j, h) \models pqr$ for some unique $(p,q,r) \in \mathcal{C}$. Since $(u_j, y_j, h) \models \mathsf{g}_{20}\mathsf{g}_{02}\mathsf{e}_2$, we have $p = \mathsf{g}_{20}$. Since $(x_i, v_i, h) \models \mathsf{g}_{10}\mathsf{g}_{01}\mathsf{e}_1$, we have $q = \mathsf{g}_{01}$. By the properties of $\mathcal{C}$, the only possibility for r is a tile, say τ^{ij}.

It remains to show that $(i,j) \mapsto \tau^{ij}$ is a tiling of $\mathbb{Z} \times \mathbb{Z}$ (with τ^{00} at the origin of course). Take $i, j \in \mathbb{Z}$; we check that $Right(\tau^{ij}) = Left(\tau^{i+1,j})$. We will use the points x_i, x_{i+1}, v_i, y_j, and one more point $t \in h$. Let m be whichever of v_i, y_j has greatest height in h. Since $(x_i, m, h) \models p\mathsf{g}_{01}r$ for some p, r, there must exist $t \geq m$ with $t \in h$ and $(x_i, t, h) \models \mathsf{g}_{10}\mathsf{g}_{01}\mathsf{e}_1$. Since $(\mathsf{e}_1, \mathsf{e}_1, \mathsf{e}_1) \in \mathcal{C}$, there is $g \in \mathcal{H}$ with $t \in g$ and $(x_i, t, g) \models \mathsf{e}_1\mathsf{e}_1\mathsf{e}_1$. Note that $v_i, y_j \in g$.

We consider some triples formed from the above with h and g (see Fig. 4 in Appendix). First, we show that $(x_i, v_i, g) \models \mathsf{e}_1\mathsf{e}_1\mathsf{e}_1$. Let $(x_i, v_i, g) \models pqr$, say, for unique $(p,q,r) \in \mathcal{C}$. We know that

$(x_i,t,g) \models \mathsf{e}_1\mathsf{e}_1\mathsf{e}_1$, so $q = \mathsf{e}_1$. Also, $(x_i,v_i,h) \models \mathsf{g}_{10}\mathsf{g}_{01}\mathsf{e}_1$, so $r = \mathsf{e}_1$ as well. So we must have $p = \mathsf{e}_1$ as required. In the same way we can successively show that $(x_{i+1},v_i,g) \models \mathsf{e}_1{+}1_1{-}1_1$, $(x_i,y_j,g) \models \breve{\tau}^{ij}\mathsf{e}_1\tau^{ij}$, and $(x_{i+1},y_j,g) \models \breve{\tau}^{ij}{+}1_1\tau^{i+1,j}$. Therefore, $(\breve{\tau}^{ij},+1_1,\tau^{i+1,j}) \in \mathcal{C}$, and the cited properties of $\mathcal{C}$ yield $\mathit{Left}(\tau^{i+1,j}) = \mathit{Right}(\tau^{ij})$, as required. That $\mathit{Up}(\tau^{ij}) = \mathit{Down}(\tau^{i,j+1})$ is shown similarly.

We have shown that $\mathcal{T}$ is a yes-instance of the cited tiling problem iff $\widehat{\varphi}_{\mathcal{T}}$ has a bundled $\mathcal{QCTL}^*_F$-model. The undecidability of the tiling problem now yields undecidability of satisfiability of BQCTL^*_F-sentences. $\square$

REMARK 5. A more general version of this result can be formulated in terms of products of propositional modal logics; see [9] and references therein. Namely, there is no decidable set L of formulas such that $\mathsf{K} \times \mathsf{CTL}^*_F \subseteq L \subseteq \mathsf{S5} \times \mathsf{CTL}^*_F$. The same holds for the propositional bundled case as well.

4 Decidable fragments

The ‘negative’ results of Section 3 can be ‘explained’ by the fact that all the undecidable fragments there are in a sense ‘three-dimensional’, which is often a cause of bad computational properties. The three-variable fragment of classical first-order logic is undecidable even without equality [17], and products of three propositional modal logics are usually undecidable [13]. In Theorem 4 we also have quantification in three dimensions: temporal operators, path quantifiers and the domain quantification.

A natural way to reduce the interaction between the dimensions is to restrict first-order quantification to state formulas—that is to work with the language $\mathcal{QPCTL}^s$—and to limit the scope of the path quantifiers (hence also of the temporal operators) to formulas with ≤ 1 free variable; cf. [14, 24].

DEFINITION 6 (MONODIC FORMULAS). Let $\mathcal{QPCTL}^s_{\boxed{1}}$ be the set of all $\mathcal{QPCTL}^s$-formulas φ such that any subformula of φ of the form $\mathsf{A}\psi$, $\psi_1 \mathsf{U} \psi_2$, or $\psi_1 \mathsf{S} \psi_2$ has at most one free variable. Such formulas φ will be called *monodic* and $\mathcal{QPCTL}^s_{\boxed{1}}$ the *monodic fragment* of $\mathcal{QPCTL}^s$.

It should be clear from the definition that $\mathcal{QPCTL}^s_{\boxed{1}}$ contains full propositional $\mathcal{CTL}^*$ and the full first-order (non-temporal) language. The latter means, in particular, that the monodic fragments of the logics under consideration are still undecidable.

The main aim of this section is to prove a satisfiability criterion for the monodic formulas (Theorem 8) and then apply it in order to obtain various decidable fragments of FOBTLs. As in [14], the idea is to encode models in structures called quasimodels and then express the statement ‘there exists a quasimodel satisfying a given monodic sentence’ as a monadic second-order sentence.

In what follows we assume that $\varphi \in \mathcal{QPCTL}^s_{\boxed{1}}$. Denote by $\mathit{sub}_n(\varphi)$ the closure under negation of the set of state subformulas of formulas in φ containing $\leq n$ free variables; $\mathit{con}(\varphi)$ is the set of constants in φ. Without loss of generality, we may identify ψ and $\neg\neg\psi$, so $\mathit{sub}_n(\varphi)$ is finite. For every formula $\psi(x) = \mathsf{A}\chi(x)$ in $\mathit{sub}_1(\varphi)$ with a free variable x, we reserve a unary predicate $P_\psi(x)$, and for every sentence $\psi = \mathsf{A}\chi$ in $\mathit{sub}_0(\varphi)$ we fix a propositional variable p_ψ. $P_\psi(x)$ and p_ψ are called the *surrogates* of $\psi(x)$ and ψ, respectively. Given a state formula ψ, denote by $\overline{\psi}$ the result of replacing all its maximal subformulas of the form $\mathsf{A}\chi$ with their surrogates. Thus, $\overline{\psi}$ contains neither temporal operators nor path quantifiers at all—it is a $\mathcal{PL}$-formula.

Let x be a variable not occurring in φ and let $\mathit{sub}_x(\varphi) = \{\psi\{x/y\} : \psi(y) \in \mathit{sub}_1(\varphi)\}$. Define a *type for* φ as a subset t of $\mathit{sub}_x(\varphi)$ such that $\psi \wedge \chi \in t$ iff $\psi \in t$ and $\chi \in t$, for every $\psi \wedge \chi \in \mathit{sub}_x\varphi$, and $\neg\psi \in t$ iff $\psi \notin t$, for every $\psi \in \mathit{sub}_x\varphi$. Given a type t for φ and a constant $c \in \mathit{con}(\varphi)$, the pair $\langle t, c\rangle$ is called an *indexed type* for φ. We will not be distinguishing between a type t and the conjunction $\bigwedge t$ of formulas in it.

Suppose that T is a set of types for φ, and $T^{con} = \{\langle t, c\rangle : c \in con(\varphi)\}$ is a set of indexed types such that $\{t : \langle t, c\rangle \in T^{con}\} \subseteq T$. Then the pair $\mathfrak{C} = \langle T, T^{con}\rangle$ is called a *state candidate* for φ. Consider a first-order $\mathcal{PL}$-structure

$$\mathfrak{D} = \left\langle D, P_0^{\mathfrak{D}}, \ldots, c_0^{\mathfrak{D}}, \ldots \right\rangle \tag{1}$$

and suppose that $a \in D$. The set $t^{\mathfrak{D}}(a) = \{\psi \in sub_x(\varphi) : \mathfrak{D} \models \overline{\psi}[a]\}$ is clearly a type for φ. Say that $\mathfrak{D}$ *realizes* a state candidate $\langle T, T^{con}\rangle$ if $T = \{t^{\mathfrak{D}}(a) : a \in D\}$ and $T^{con} = \{\langle t^{\mathfrak{D}}(c^{\mathfrak{D}}), c\rangle : c \in con(\varphi)\}$. A state candidate is *realizable* if some $\mathfrak{D}$ realizes it.

Let $\mathfrak{F} = \langle W, <\rangle$ be an ω-tree. A *state function* for φ over $\mathfrak{F}$ is a map f associating with each $w \in W$ a realizable state candidate $f(w) = \langle T_w, T_w^{con}\rangle$ for φ.

For every subformula $\psi(y)$ of φ, every full branch h in $\mathfrak{F}$, every $w \in h$, and every function r mapping each $w \in W$ to a type for φ, define inductively a 'formula' $cond(\psi, h, r, w)$:

- $cond(\psi, h, r, w)$ is '$\psi\{x/y\} \in r(w)$' if $\psi(y)$ is a state formula;
- if $\psi_1 \wedge \psi_2$ is not a state formula, $cond(\psi_1 \wedge \psi_2, h, r, w)$ is '$cond(\psi_1, h, r, w) \wedge cond(\psi_2, h, r, w)$';
- if $\neg\psi$ is not a state formula, $cond(\neg\psi, h, r, w)$ is '$\neg cond(\psi, h, r, w)$';
- $cond(\psi_1 \mathsf{U} \psi_2, h, r, w)$ is '$\exists v > w\,(v \in h \wedge cond(\psi_2, h, r, v) \wedge \forall u \in (w, v)\, cond(\psi_1, h, r, u))$';
- $cond(\psi_1 \mathsf{S} \psi_2, h, r, w)$ is '$\exists v < w\,(cond(\psi_2, h, r, v) \wedge \forall u \in (v, w)\, cond(\psi_1, h, r, u))$'.

Since every $\mathcal{QPCTL}^s$-formula is built from state formulas using only the Booleans and the temporal operators, this is well-defined. Let f be a state function for φ over $\mathfrak{F} = \langle W, <\rangle$, with $f(w) = \langle T_w, T_w^{con}\rangle$ for $w \in W$. By a *run* in f we mean a function r from W into the set $\bigcup_{w \in W} T_w$ such that

- $r(w) \in T_w$, for all $w \in W$,
- for all $\mathsf{A}\psi \in sub_x\varphi$ and $w \in W$, we have $\mathsf{A}\psi \in r(w)$ iff $\forall h \ni w\ cond(\psi, h, r, w)$.

We call f a *quasimodel* for φ over $\mathfrak{F}$ if the following conditions hold:

- for each $c \in con(\varphi)$, the function r_c defined by $r_c(w) = t$, for $\langle t, c\rangle \in T_w^{con}$, $w \in W$, is a run in f,
- for all $w \in W$ and $t \in T_w$, there exists a run r in f such that $r(w) = t$.

Say that φ is *satisfied* in f if there are $w \in W$ and $t \in T_w$ such that $\varphi \in t$.

THEOREM 7. *A $\mathcal{QPCTL}^s_{\boxed{1}}$-sentence φ is satisfiable in a full tree model based on $\mathfrak{F} = \langle W, <\rangle$ iff it is satisfied in a quasimodel for φ over $\mathfrak{F}$.*

Proof. Similar to the proof of Theorem 14 in [14]. For ($\Rightarrow$), we construct a quasimodel from a model in the obvious way. For ($\Leftarrow$), given a quasimodel f for φ over $\mathfrak{F} = \langle W, <\rangle$, we take a cardinal κ and form a set $\mathcal{R}$ of κ copies of each run in f. Since the language $\mathcal{PL}$ is countable and does not contain equality, it follows from classical model theory that if κ is large enough then each $f(w)$ is realized by a model $\mathfrak{D}_w$ with domain $\mathcal{R}$, and with $t^{\mathfrak{D}_w}(r) = r(w)$ for all $r \in \mathcal{R}$. Let $\mathfrak{M} = \langle \mathfrak{F}, \mathcal{R}, (\mathfrak{D}_w)_{w \in W}\rangle$. We may now check by induction that for all state subformulas $\psi(\overline{x})$ of φ, all $\overline{r}$ in $\mathcal{R}$, and all h, w, we have $\mathfrak{D}_w \models \overline{\psi}[\overline{a}]$ iff $(\mathfrak{M}, h, w) \models \psi[\overline{a}]$ for any $h \ni w$. It follows that φ is satisfied in $\mathfrak{M}$. $\Box$

In the proof of the next theorem we show that the existence of a quasimodel satisfying φ can be expressed as a monadic second-order formula, which yields the following satisfiability criterion:

THEOREM 8. *Let $\mathcal{L} \subseteq \mathcal{QPCTL}^{s}_{\boxed{1}}$ and suppose that there is an algorithm that decides for any $\mathcal{L}$-sentence φ whether an arbitrarily-given state candidate is realizable. Then the satisfiability problem for $\mathcal{L}$-sentences in both bundled and full tree models is decidable.*

Proof. We consider only full tree models, leaving the bundled case to Lemma 1. For each $\psi \in sub_x(\varphi)$, let R_ψ be a unary predicate variable, and for each type t for φ, let

$$\chi_t(x) \quad = \quad \bigwedge_{\psi \in t} R_\psi(x) \;\wedge \bigwedge_{\psi \in sub_x(\varphi)\setminus t} \neg R_\psi(x).$$

The formula $\chi_t(x)$ says that the $R_\psi(x)$ define the type t at x. For any subformula $\psi(y)$ of φ, an individual variable x, and a set variable h, define inductively a formula $\gamma(\psi, h, x)$ of monadic second-order logic by taking $\gamma(\psi, h, x) = R_\psi(x)$ if ψ is a state formula, and for non-state formulas taking:

- $\gamma(\psi_1 \wedge \psi_2, h, x) = \gamma(\psi_1, h, x) \wedge \gamma(\psi_2, h, x)$;
- $\gamma(\neg\psi, h, x) = \neg\gamma(\psi, h, x)$;
- $\gamma(\psi_1 \mathsf{U} \psi_2, h, x) = \exists y \in h\,(x < y \wedge \gamma(\psi_2, h, y) \wedge \forall z(x < z < y \to \gamma(\psi_1, h, z)))$;
- $\gamma(\psi_1 \mathsf{S} \psi_2, h, x) = \exists y\,(y < x \wedge \gamma(\psi_2, h, y) \wedge \forall z(y < z < x \to \gamma(\psi_1, h, z)))$.

Denote by $\beta(h, x)$ a monadic second-order formula saying that h is a full branch containing x. Let Σ be the set of all realizable state candidates for φ (Σ can be constructed effectively), and let P_s ($s \in \Sigma$) be a unary predicate variable. Then

$$\rho = \forall x \bigwedge_{\langle T, T^{con}\rangle \in \Sigma} \big(P_{\langle T, T^{con}\rangle}(x) \to \bigvee_{t \in T} \chi_t(x)\big) \wedge \forall x \bigwedge_{\mathsf{A}\psi \in sub_x(\varphi)} \Big(R_{\mathsf{A}\psi}(x) \leftrightarrow \forall h\,\big(\beta(h, x) \to \gamma(\psi, h, x)\big)\Big).$$

says that the $R_\psi(x)$ define a run through realizable state candidates in Σ defined with the help of the P_s. Finally, we define a monadic second-order sentence σ_φ by:

$$\begin{aligned}
\exists_{s \in \Sigma} P_s \Bigg(\forall x \Big[\bigvee_{s \in \Sigma} P_s(x) \wedge \bigwedge_{\substack{s' \in \Sigma \\ s \neq s'}} \neg P_{s'}(x) \Big] \;\wedge \bigvee_{\substack{\langle T, T^{con}\rangle \in \Sigma \\ \varphi \in \bigcup T}} \exists x\, P_{\langle T, T^{con}\rangle}(x) \\
\wedge \bigwedge_{c \in con(\varphi)} \exists_{\psi \in sub_x(\varphi)} R_\psi \Big[\rho \wedge \forall x \bigwedge_{\substack{\langle T, T^{con}\rangle \in \Sigma \\ \langle t, c\rangle \in T^{con}}} \big(P_{\langle T, T^{con}\rangle}(x) \to \chi_t(x)\big)\Big] \\
\wedge\; \forall x \bigwedge_{\substack{\langle T, T^{con}\rangle \in \Sigma \\ t \in T}} \Big[P_{\langle T, T^{con}\rangle}(x) \to \exists_{\psi \in sub_x(\varphi)} R_\psi(\rho \wedge \chi_t(x))\Big] \Bigg).
\end{aligned}$$

It is not hard to check that $\mathfrak{F} \models \sigma_\varphi$ iff φ is satisfied in a quasimodel over $\mathfrak{F}$. It remains to recall that the monadic second-order theory of countably branching trees is decidable (this can easily be shown by reduction to the monadic second-order theory of two successor functions, which is decidable [21]). A Löwenheim–Skolem–Tarski argument (see Theorem 14 of the Appendix) will show that any satisfiable $\mathcal{QPCTL}^{s}_{\boxed{1}}$-formula has a model based on a countably branching tree. $\square$

As a consequence we obtain

THEOREM 9. *The following fragments are decidable:*

- *the two-variable fragment of* $\mathsf{QPCTL}^* \cap \mathcal{QPCTL}^s_{\boxed{1}}$;
- *the two-variable fragment of* $\mathsf{BQPCTL}^* \cap \mathcal{QPCTL}^s_{\boxed{1}}$;
- *the monadic fragment of* $\mathsf{QPCTL}^* \cap \mathcal{QPCTL}^s_{\boxed{1}}$;
- *the monadic fragment of* $\mathsf{BQPCTL}^* \cap \mathcal{QPCTL}^s_{\boxed{1}}$;
- *the guarded fragment of* $\mathsf{QPCTL}^* \cap \mathcal{QPCTL}^s_{\boxed{1}}$;
- *the guarded fragment of* $\mathsf{QPCTL}^* \cap \mathcal{QPCTL}^p_1$.

A similar construction yields the following theorem answering a question that has puzzled many temporal logicians since [11]:

THEOREM 10. *It is decidable whether an arbitrarily-given propositional* $\mathcal{PCTL}^*$*-formula is satisfiable in a 'non-local' full tree model—where truth values of atoms may depend on the branch of evaluation. The same holds for non-local bundled tree models.*

5 Conclusion

This paper may be regarded as a beginning of systematic research into the computational behavior of first-order branching time logics. We have obtained decidability results of two kinds. The most striking bad news is that the one-variable fragments of logics containing BQCTL^*_F and QCTL^*_F are undecidable, which contrasts with the situation in many other modal and temporal first-order logics (cf. [14, 24]). The good news is that there are still ways of obtaining decidable fragments with non-trivial interaction between first-order quantifiers, path quantifiers and temporal operators. In this paper, we considered the case when the first-order quantifiers are applied only to state formulas, while the path-quantifiers and temporal operators are applicable to formulas with ≤ 1 free variable.

Another way to obtain decidable fragments is to restrict further the formulas that can have free variables. For example, using quasimodels and a mosaic technique one can prove a result similar to Theorem 8 for bundled models and the fragment of $\mathcal{QCTL}^*$ in which $\bigcirc$ may be applied to formulas with at most one free variable, U, E and A are applicable only to sentences, and there are no restrictions on first-order quantification (for a precise formulation see [15]).

Two other promising ways are to allow arbitrary first-order quantification, but to restrict applications of path quantifiers to closed formulas and those of linear temporal operators to formulas with ≤ 1 free variable, and vice versa. However, they remain open for investigation.

Acknowledgements

The work of the first and third authors was partially supported by UK EPSRC grant GR/R45369/01 "Analysis and mechanisation of decidable first-order temporal logics" (the third author was also partially supported by grant 99-01-00968 from the Russian Foundation for Basic Research). The work of the second author was supported by Deutsche Forschungsgemeinschaft (DFG) grant Wo583/3-1. We are grateful to S. Bauer, A. Kurucz, A. Rabinovich, M. Reynolds, C. Stirling, and A. Zanardo for stimulating discussions and comments.

References

[1] H. Andréka, J. van Benthem, and I. Németi. Modal languages and bounded fragments of predicate logic. *J. Philosophical Logic*, 27:217–274, 1998.

[2] R. Berger. *The undecidability of the domino problem*, volume 66 of *Memoirs*. Amer. Math. Soc., Providence, RI, 1966.

[3] E. Börger, E. Grädel, and Yu. Gurevich. *The Classical Decision Problem*. Perspectives in Mathematical Logic. Springer, 1997.

[4] J. Burgess. The unreal future. *Theoria*, 44:157–179, 1978.

[5] J. Burgess. Logic and time. *J. Symbolic Logic*, 44:566–582, 1979.

[6] E. Clarke, E. Emerson. Synthesis of synchronization skeletons for branching time temporal logic. In *Proc. IBM Workshop on Logic of Programs*, pages 52–71, Springer-Verlag, 1981.

[7] E. Emerson. Temporal and modal logic. In J. van Leeuwen, editor, *Handbook of Theoretical Computer Science*, pages 996–1076, 1990.

[8] E. Emerson and J. Halpern. 'Sometimes' and 'not never' revisited: on branching versus linear time. *J. ACM*, 33:151–178, 1979.

[9] D. Gabbay, A. Kurucz, F. Wolter, and M. Zakharyaschev. *Many-Dimensional Modal Logics: Theory and Applications*. Elsevier, 2002. See `http://www.dcs.kcl.ac.uk/staff/mz`.

[10] D. Gabbay, M. Reynolds, and M. Finger. *Temporal Logic: Mathematical Foundations and Computational Aspects, Volume 2*. Oxford University Press, 2000.

[11] Y. Gurevich and S. Shelah. The decision problem for branching time logic. *J. Symbolic Logic*, 50:668–681, 1985.

[12] R. Hirsch and I. Hodkinson. Representability is not decidable for finite relation algebras. *Trans. Amer. Math. Soc.*, 353:1403–1425, 2001.

[13] R. Hirsch, I. Hodkinson, and A. Kurucz. On modal logics between $\mathsf{K} \times \mathsf{K} \times \mathsf{K}$ and $\mathsf{S5} \times \mathsf{S5} \times \mathsf{S5}$. *J. Symbolic Logic*, 2001. (In press.)

[14] I. Hodkinson, F. Wolter, and M. Zakharyaschev. Decidable fragments of first-order temporal logics. *Annals of Pure and Applied Logic*, 106:85–134, 2000.

[15] I. Hodkinson, F. Wolter, and M. Zakharyaschev. Monodic fragments of first-order temporal logics: 2000–2001 A.D. In R. Nieuwenhuis and A. Voronkov, editors, *Logic for Programming, Artificial Intelligence and Reasoning*, number 2250 of LNAI, Springer, 2001, pages 1–23.

[16] R. Ladner and J. Reif. The logic of distributed protocols. In J. Halpern, editor, *Theoretical Aspects of Reasoning about Knowledge*, pp.207–222, Morgan Kaufmann, 1986.

[17] R. Maddux. The equational theory of CA_3 is undecidable. *J. Symbolic Logic*, 45:311–316, 1980.

[18] A. Prior. *Past, Present and Future*. Oxford University Press, 1967.

[19] A. Prior. Now. *Nous*, 2:101–119, 1968.

[20] O. Øhrstrøm and P.F.V. Hasle. *Temporal Logic: From Ancient Ideas to Artificial Intelligence*. volume 57 of *Studies in Logic and Philosophy*. Kluwer Academic Publishers, 1996.

[21] M. O. Rabin. Decidability of second order theories and automata on infinite trees. *American Mathematical Society Transactions*, 141:1–35, 1969.

[22] S. Woelf. Combinations of tense and modality for predicate logic. *J. Philosophical Logic*, 28:371–398, 1999.

[23] M. Wooldridge and M. Fisher. A first-order branching time logic of mutli-agent systems. In B. Neumann, editor, *Proc. 10th European Conference on Artificial Intelligence (ECAI92)*, pages 234 – 238, Chichester, 1992. John Wiley and Sons.

[24] F. Wolter and M. Zakharyaschev. Decidable fragments of first-order modal logics. *J. Symbolic Logic*, 66:1415–1438, 2001.

[25] F. Wolter and M. Zakharyaschev. Axiomatizing the monodic fragment of first-order temporal logic. *Annals of Pure and Applied Logic* (In print.) See `http://www.dcs.kcl.ac.uk/staff/mz`.

[26] A. Zanardo. Branching-time logic with quantification over branches: the point of view of modal logic. *J. Symbolic Logic*, 61:1–39, 1996.

Appendix

1 A monodic $\mathcal{QCTL}^*_F$-sentence with only uncountable models

We take a signature with relation symbols P (unary) and $<$ (binary). Let φ be the conjunction of the following $\mathcal{QCTL}^*_F$-sentences:

1. $\mathsf{A}\exists x\forall y\big(\Diamond^+_F P(y) \to [\Diamond^+_F\Box^+_F\forall z(P(z)\to z>y)\to x>y]\wedge[\Diamond^+_F\Box^+_F\forall z(P(z)\to z<y)\to x<y]\big)$
2. $\mathsf{A}\Box^+_F\exists xP(x)$
3. $\mathsf{A}\Box^+_F\forall x\big(P(x)\to \mathsf{E}\Diamond^+_F\mathsf{A}\Box^+_F\forall y(P(y)\to y<x)\wedge \mathsf{E}\Diamond^+_F\mathsf{A}\Box^+_F\forall y(P(y)\to y>x)\big)$
4. $\forall xyz(\neg(x<x)\wedge(x<y\wedge y<z\to x<z))$

φ is monodic and only uses the reflexive temporal operators $\Box^+_F, \Diamond^+_F$. The idea is that φ forces an embedding of 2^ω full branches into the domain.

LEMMA 11. *φ has a full tree model.*

Proof. Let $\mathfrak{F} = \langle {}^{<\omega}2, \prec\rangle$ be the tree with $\prec$ the ordering of proper initial segment. We identify full branches of $\mathfrak{F}$ with elements of ${}^{\omega}2$. Let $D = {}^{\leq\omega}2$. Define a linear order $<$ on D in the standard way. First, the nodes of $\mathfrak{F}$ are ordered 'left–right' by 'projecting' down onto a horizontal axis, as shown in Figure 1.

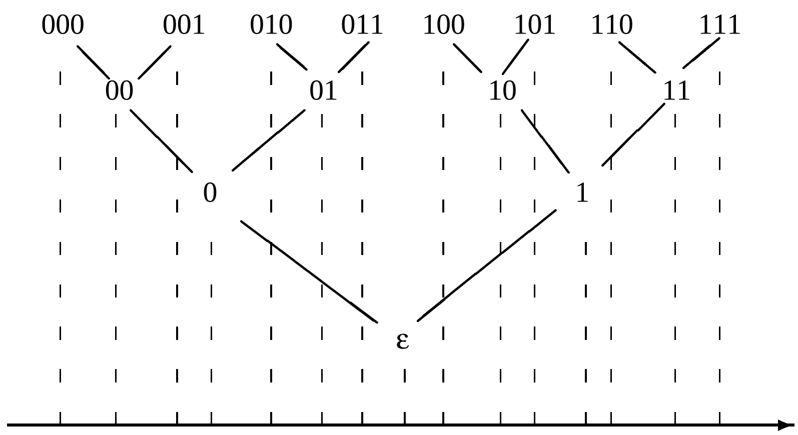

Figure 1: Ordering the nodes of ${}^{<\omega}2$

The full branches h of $\mathfrak{F}$ are then inserted in this ordering in their natural places: h goes to the left of all $t \in {}^{<\omega}2$ such that $t\hat{}0 \in h$, and to the right of all $t \in {}^{<\omega}2$ such that $t\hat{}1 \in h$. This defines $<$ on D.

Now for each $t \in {}^{<\omega}2$ let M_t be the structure with domain D, with $<$ interpreted as above, and with $M_t \models P(d)$ iff $d = t$, for $d \in D$. We obtain a model $\mathfrak{M} = (\mathfrak{F}, D, (M_t : t \in {}^{<\omega}2))$.

Write ε for the root of $\mathfrak{F}$. We check that $(\mathfrak{M}, h, \varepsilon) \models \varphi$ for any h. For sentence 1 above, let h be any branch of $\mathfrak{F}$. We require

$$(\mathfrak{M},h,\varepsilon)\models\exists x\forall y\big(\Diamond^+_F P(y)\to[\Diamond^+_F\Box^+_F\forall z(P(z)\to z>y)\to x>y]\wedge[\Diamond^+_F\Box^+_F\forall z(P(z)\to z<y)\to x<y]\big).$$

We will take 'x' to be $h \in D$. Then for any $d \in D$ with $(\mathfrak{M}, h, \varepsilon) \models \Diamond^+_F P(d)$, we require $(\mathfrak{M}, h, \varepsilon) \models \Diamond^+_F\Box^+_F\forall z(P(z) \to z > d) \to h > d$. Such a d must be in ${}^{<\omega}2$. Suppose that there is $u \in h$ with $M_v \models \forall z(P(z) \to z > d)$ for all $v \in h$ with $v \succeq u$. Choose $v \in h$ higher than d. Since $M_v \models P(v)$, we have $v > d$. By definition of $<$ we have $v > d$ iff $h > d$. So $h > d$, as required. The other side ($x < y$) is similar.

Sentence 2, $\mathsf{A}\Box_F^+\exists x P(x)$, is clearly true, since $M_t \models P(t)$ for any $t \in {}^{<\omega}2$.

For sentence 3, $\mathsf{A}\Box_F^+\forall x\big(P(x) \to \mathsf{E}\Diamond_F^+\mathsf{A}\Box_F^+\forall y(P(y) \to y < x) \wedge \mathsf{E}\Diamond_F^+\mathsf{A}\Box_F^+\forall y(P(y) \to y > x)\big)$, take h, $t \in h$, $d \in D$. If $(\mathfrak{M}, h, t) \models P(d)$ then $d = t$, so we can assume $t = d$. We require

$$(\mathfrak{M}, h, t) \models \mathsf{E}\Diamond_F^+\mathsf{A}\Box_F^+\forall y(P(y) \to y < t), \tag{2}$$

and similarly on the other side with $y > t$. So take g to be any branch containing $t\hat{\ }0$. By definition of $<$, every $u \succeq t\hat{\ }0$ in $\mathfrak{F}$ satisfies $u < t$ in D. Thus, $(\mathfrak{M}, g, t\hat{\ }0) \models \mathsf{A}\Box_F^+\forall y(P(y) \to y < t)$, so $(\mathfrak{M}, g, t) \models \Diamond_F^+\mathsf{A}\Box_F^+\forall y(P(y) \to y < t)$, yielding (2). The second half of sentence 3 is handled similarly, using $t\hat{\ }1$. Sentence 4 is trivially true. $\Box$

LEMMA 12. *Any full tree model of* φ *has uncountable domain.*

Proof. Assume that for some ω-tree $\mathfrak{F} = \langle W, \prec\rangle$, $\mathfrak{M} = (\mathfrak{F}, D, (M_w : w \in W))$ is a full tree model of φ. We can assume that φ is true at the root of $\mathfrak{F}$. By sentence 2, for all $w \in W$ we may pick $d_w \in D$ with $M_w \models P(d_w)$.

We define nodes $w_t \in W$ (for $t \in {}^{<\omega}2$) by induction on $|t|$. We let w_ε be the root of $\mathfrak{F}$, where ε is the empty sequence in ${}^{<\omega}2$. Let $t \in {}^{<\omega}2$ and assume that w_t is defined. By sentence 3, there is a full branch h of $\mathfrak{F}$ with $w_t \in h$ and $(\mathfrak{M}, h, w_t) \models \Diamond_F^+\mathsf{A}\Box_F^+\forall y(P(y) \to y < d_{w_t})$. So there is $w_{t\hat{\ }0} \in h$ with $w_{t\hat{\ }0} \succeq w_t$ and $(\mathfrak{M}, h, w_{t\hat{\ }0}) \models \mathsf{A}\Box_F^+\forall y(P(y) \to y < d_{w_t})$. Similarly, there are $g \ni w_{t\hat{\ }1} \succeq w_t$ with $(\mathfrak{M}, g, w_{t\hat{\ }1}) \models \mathsf{A}\Box_F^+\forall y(P(y) \to y > d_{w_t})$.

For each full branch β of ${}^{<\omega}2$, pick a full branch $\hat{\beta}$ of $\mathfrak{F}$ containing $\{w_t : t \in \beta\}$. Sentence 1 tells us that there is $d_\beta \in D$ with

$$(\mathfrak{M}, \hat{\beta}, w_\varepsilon) \models \forall y\big(\Diamond_F^+P(y) \to [\Diamond_F^+\Box_F^+\forall z(P(z) \to z > y) \to d_\beta > y] \wedge [\Diamond_F^+\Box_F^+\forall z(P(z) \to z < y) \to d_\beta < y]\big).$$

Let β, γ be distinct branches of ${}^{<\omega}2$; we claim that $d_\beta \neq d_\gamma$. Let t be the highest node in $\beta \cap \gamma$. We may assume without loss of generality that $t\hat{\ }0 \in \beta$ and $t\hat{\ }1 \in \gamma$. Now, $(\mathfrak{M}, \hat{\beta}, w_\varepsilon) \models \forall y\big(\Diamond_F^+P(y) \to [\Diamond_F^+\Box_F^+\forall z(P(z) \to z < y) \to d_\beta < y]\big)$. We have $(\mathfrak{M}, \hat{\beta}, w_\varepsilon) \models \Diamond_F^+P(d_{w_t})$ since $M_{w_t} \models P(d_{w_t})$. Hence,

$$(\mathfrak{M}, \hat{\beta}, w_\varepsilon) \models \Diamond_F^+\Box_F^+\forall z(P(z) \to z < d_{w_t}) \to d_\beta < d_{w_t}.$$

By definition of $w_{t\hat{\ }0}$, we have $(\mathfrak{M}, \hat{\beta}, w_{t\hat{\ }0}) \models \Box_F^+\forall y(P(y) \to y < d_{w_t})$. So certainly,

$$(\mathfrak{M}, \hat{\beta}, w_\varepsilon) \models \Diamond_F^+\Box_F^+\forall z(P(z) \to z < d_{w_t}).$$

Hence, $(\mathfrak{M}, \hat{\beta}, w_\varepsilon) \models d_\beta < d_{w_t}$. That is, $M_{w_\varepsilon} \models d_\beta < d_{w_t}$ Similarly, $M_{w_\varepsilon} \models d_\gamma > d_{w_t}$. By sentence 4, $<$ is transitive in M_{w_ε}, so $M_{w_\varepsilon} \models d_\beta < d_\gamma$. Since $<$ is irreflexive, $d_\beta \neq d_\gamma$, as claimed.

Thus, $|D| \geq |\{d_\beta : \beta \text{ a branch of } {}^{<\omega}2\}| = 2^\omega$, proving the lemma. $\Box$

REMARK 13. We can make any model of φ have uncountable branching factor by adding a conjunct $\mathsf{A}\Box_F^+\forall x\mathsf{E}\bigcirc(Q(x) \wedge \forall y(y < x \vee y > x \to \neg Q(y)))$ plus the statement that $<$ is rigid.

2 $\mathcal{QPCTL}^*$-formulas with countable models

On the other hand, if we restrict quantifiers to apply only to state formulas, we get a downward Löwenheim–Skolem–Tarski theorem:

THEOREM 14. *Let* φ *be a* $\mathcal{QPCTL}^s$*-sentence, and suppose that* φ *is satisfiable in a full tree model. Then* φ *is satisfiable in a full tree model with countable tree and domain.*

Proof. Let $\mathfrak{M}$ be any full tree model. We may view $\mathfrak{M} = (\mathfrak{F}, D, I)$ as a three-sorted first-order structure, with sorts for the domain, the tree, and the set of all full branches. Taking a countable elementary substructure of this yields a *bundled* tree model $\mathfrak{N} = (\mathfrak{F}_0, D_0, \mathcal{H}, I_0)$ whose domain D_0, tree $\mathfrak{F}_0$, and bundle $\mathcal{H}$ are countable. It is easy to translate $\mathcal{QPCTL}^*$-formulas to three-sorted first-order formulas with the same meaning. It follows that for any $\overline{a}$ in D_0, $h \in \mathcal{H}$, $w \in h$, and any $\mathcal{QPCTL}^*$-formula $\psi(\overline{x})$, we have $(\mathfrak{M}, h, w) \models \psi(\overline{a})$ iff $(\mathfrak{N}, h, w) \models \psi(\overline{a})$.

Now let $\overline{\mathfrak{N}} = (\mathfrak{F}_0, D_0, I_0)$ be the full tree model based on $\mathfrak{N}$. We claim that for all formulas $\varphi(\overline{x})$ as in the formulation of the theorem, all full branches g of $\mathfrak{F}_0$ (g may not be in $\mathcal{H}$), all $w \in g$, and all $\overline{a}$ in D_0, we have

$$(\mathfrak{M}, g, w) \models \varphi(\overline{a}) \iff (\overline{\mathfrak{N}}, g, w) \models \varphi(\overline{a}).$$

The proof is by induction on φ. The atomic, boolean, and temporal cases are easy and we omit them. Consider the case $\mathsf{E}\varphi(\overline{x})$ and inductively assume the result for φ. If $(\overline{\mathfrak{N}}, g, w) \models \mathsf{E}\varphi(\overline{a})$, it is easily seen that $(\mathfrak{M}, g, w) \models \mathsf{E}\varphi(\overline{a})$. Conversely, if $(\mathfrak{M}, g, w) \models \mathsf{E}\varphi(\overline{a})$, pick $h \in \mathcal{H}$ containing w. Clearly, $(\mathfrak{M}, h, w) \models \mathsf{E}\varphi(\overline{a})$, so $(\mathfrak{N}, h, w) \models \mathsf{E}\varphi(\overline{a})$. So there is $h' \in \mathcal{H}$ with $(\mathfrak{N}, h', w) \models \varphi(\overline{a})$. Thus, $(\mathfrak{M}, h', w) \models \varphi(\overline{a})$. Inductively, $(\overline{\mathfrak{N}}, h', w) \models \varphi(\overline{a})$. So $(\overline{\mathfrak{N}}, g, w) \models \mathsf{E}\varphi(\overline{a})$, as required.

Finally consider the case $\exists x \varphi(x, \overline{y})$, for a state formula φ for which we assume the result inductively. If $(\overline{\mathfrak{N}}, g, w) \models \exists x \varphi(x, \overline{a})$, then $(\overline{\mathfrak{N}}, g, w) \models \varphi(b, \overline{a})$ for some $b \in D_0$. Inductively, $(\mathfrak{M}, g, w) \models \varphi(b, \overline{a})$, so $(\mathfrak{M}, g, w) \models \exists x \varphi(x, \overline{a})$. Conversely, suppose that $(\mathfrak{M}, g, w) \models \exists x \varphi(x, \overline{a})$. Pick $h \in \mathcal{H}$ containing w. Then as $\exists x \varphi(x, \overline{y})$ is a state formula, $(\mathfrak{M}, h, w) \models \exists x \varphi(x, \overline{a})$. So $(\mathfrak{N}, h, w) \models \exists x \varphi(x, \overline{a})$, whence $(\mathfrak{N}, h, w) \models \varphi(b, \overline{a})$ for some $b \in D_0$. Then $(\mathfrak{M}, h, w) \models \varphi(b, \overline{a})$. Since φ is a state formula, $(\mathfrak{M}, g, w) \models \varphi(b, \overline{a})$. Inductively, $(\overline{\mathfrak{N}}, g, w) \models \varphi(b, \overline{a})$, so $(\overline{\mathfrak{N}}, g, w) \models \exists x \varphi(x, \overline{a})$ as required. ☐

3 Diagrams for Theorem 4

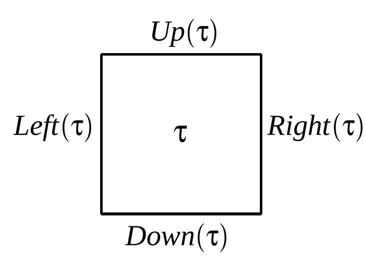

Figure 2: a tile τ

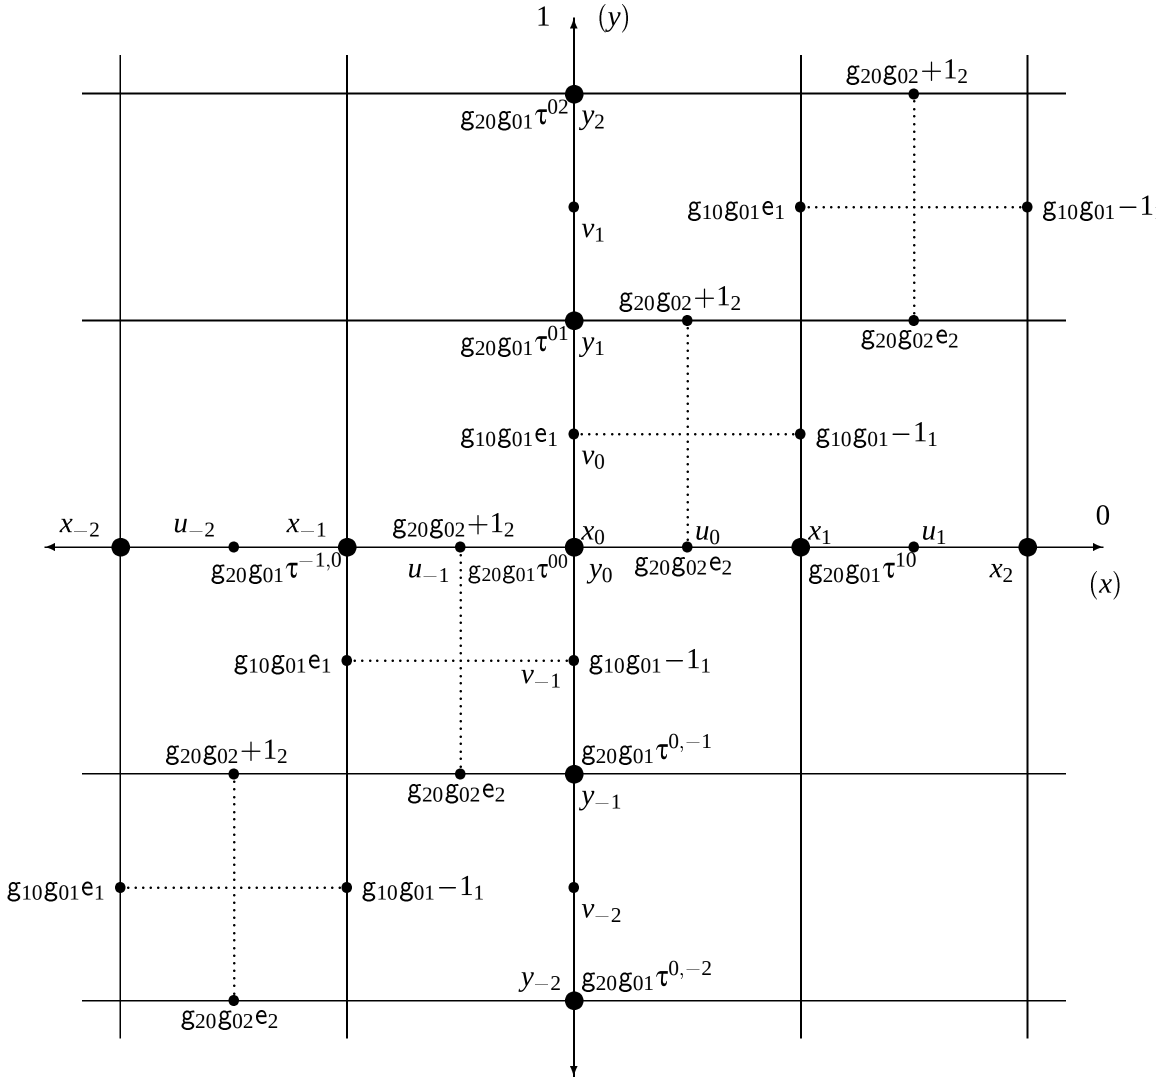

Figure 3: The defined points

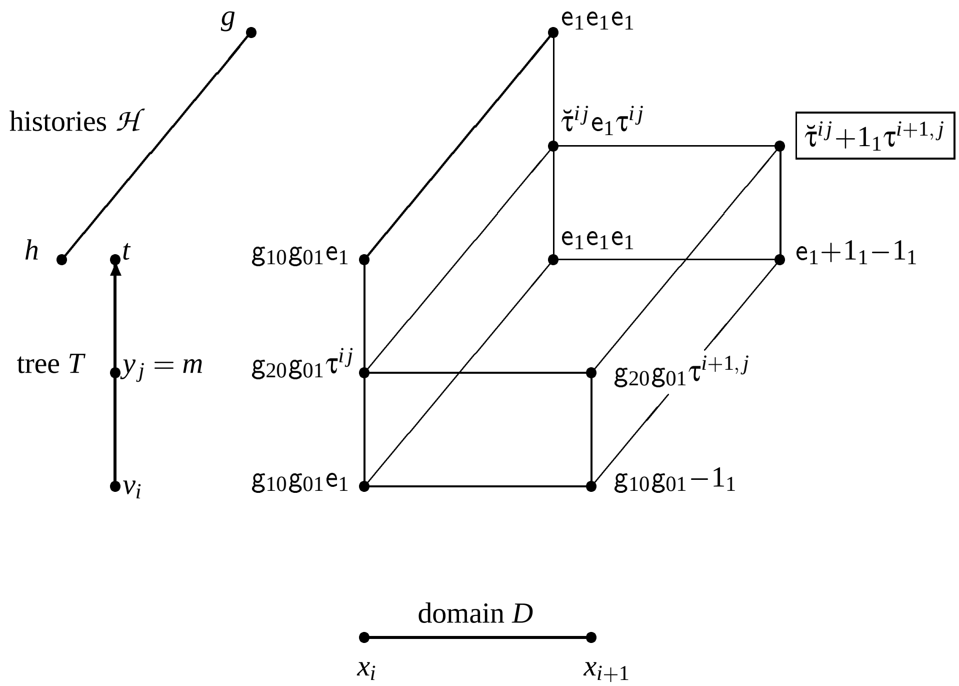

Figure 4: Checking $\mathit{Right}(\tau^{ij}) = \mathit{Left}(\tau^{i+1,j})$.