A Decidable Temporal Logic of Repeating Values*

Stéphane Demri[1], Deepak D'Souza[2], and Régis Gascon[1]

[1] LSV, ENS Cachan, CNRS, INRIA
{demri,gascon}@lsv.ens-cachan.fr
[2] Dept. of Computer Science & Automation,
Indian Institute of Science, Bangalore, India
deepakd@csa.iisc.ernet.in

Abstract. Various logical formalisms with the freeze quantifier have been recently considered to model computer systems even though this is a powerful mechanism that often leads to undecidability. In this paper, we study a linear-time temporal logic with past-time operators such that the freeze operator is only used to express that some value from an infinite set is repeated in the future or in the past. Such a restriction has been inspired by a recent work on spatio-temporal logics. We show decidability of finitary and infinitary satisfiability by reduction into the verification of temporal properties in Petri nets. This is a surprising result since the logic is closed under negation, contains future-time and past-time temporal operators and can express the nonce property and its negation. These ingredients are known to lead to undecidability with a more liberal use of the freeze quantifier. The paper contains also insights about the relationships between temporal logics with the freeze operator and counter automata.

1 Introduction

Temporal logic with freeze. In logical languages, the freeze mechanism allows to store a value in a register and to test later the value in the register with a current value. This operator is useful to compare values at distinct states of Kripke-like structures. The freeze quantifier has found applications in real-time logics [Hen90], in hybrid logics [Gor96,ABM01], in modal logics with predicate λ-abstraction [Fit02] and for the specification of computations of systems with unboundedly many locations as resources [LP05]. Although it is known that the freeze operator can lead to undecidability (even with only equality on data [LP05,DLN07]), many decidable temporal logics have a freeze mechanism, sometimes implicitly, see e.g. [AH94,LMS02,KV06]. Recent developments have shown the ubiquity of the freeze operator [LP05,tCF05,DLN07,Laz06,Seg06] and its high expressive power as witnessed by the Σ^1_1-completeness results shown in [DLN07].

The need to design decidable fragments of simple linear-time temporal logic LTL with the freeze quantifier stems from [DLN07,Laz06] and most known decidable fragments in [DLN07,Laz06] does not allow unrestricted use of negation. Still, finitary and infinitary satisfiability for Boolean combinations of safety formulae (with a unique register) is decidable [Laz06]. Potential applications range from the verification of infinite-state systems [Hen90,DLN07] to querying XML

* Work supported by the Indo-French project "Timed-DISCOVERI" (P2R/RNP scheme). This paper is an extended version of [DDG07].

documents or more modestly data strings [BMS+06,Seg06]. In the paper, we are interested in studying fragments of LTL with the freeze operator that are decidable in the finitary and infinitary cases, that allow unrestricted use of negation (by contrast to the flat fragments in [DLN07]) and that allow all standard past-time operators (by constrast to what is done in [BMS+06,DL06]). These are strong requirements. Even in terms of expressive power, the fragment newly shown decidable in the paper can express the "nonce property" and its negation (all the values of a variable are different at every position). Moreover, in [WZ00, Sect. 7], the authors advocate the need to consider infinitary disjunction of the form $\bigvee_{i>0} x = \mathsf{X}^i y$ where $\mathsf{X}^i y$ refers to the value of y at the i^{th} next position. This states that a future value of y is equal to the current value of x. Our fragment can express this property, with the formula $x = \Diamond y$, as well as the dual one: $\bigwedge_{i>0} x = \mathsf{X}^i y$ can be expressed by the formula $\neg(x \neq \Diamond y)$. In the paper we introduce the constraint logic CLTL$(\mathbb{N}, =)$ with atomic formulae $x = \Diamond y$ and past-time operators X^{-1} and S. This logic is denoted by CLTL$^{\Diamond}$. Hence, in CLTL$^{\Diamond}$, the freeze quantifier is only used to specify that some values are repeated. Even though CLTL$^{\Diamond}$ does not enjoy first-order completeness, see e.g. [Rab06], it satisfies interesting computational properties as shown below.

Our contribution. We show that finitary and infinitary satisfiability for CLTL$^{\Diamond}$ with temporal operators $\{\mathsf{X}, \mathsf{X}^{-1}, \mathsf{S}, \mathsf{U}\}$ is decidable. We provide a uniform proof for the finite and infinite cases based on some substantial extension of the automaton-based approach for (constraint) LTL from [VW94,DD07]. The possibility to compare two values at unbounded distance requires a special class of counter automata for which finitary and infinitary nonemptiness is shown decidable. To do so, we take advantage of a deep result from [Jan90] establishing that verifying fairness properties based on the temporal operator GF ("always eventually") in Petri nets is decidable. By contrast model-checking for full LTL over Petri nets is undecidable [HR89] (see also [Esp94] with linear-time mu-calculi). Observe that infinitary CLTL$^{\Diamond}$ is the first decidable fragment of CLTL$_1^{\downarrow}(\mathbb{N}, =)$ [DLN07] with an unrestricted use of negation and that contains all the temporal operators from $\{\mathsf{X}, \mathsf{X}^{-1}, \mathsf{S}, \mathsf{U}\}$. A nice by-product of our technique is that the extensions with temporal operators definable in Monadic Second Order Logic (MSOL) or with $x = \Diamond^{-1} y$ ("a value of y in the past is equal to the current value of x") are also decidable. Finally, we show that finitary and infinitary satisfiability for CLTL$^{\Diamond}$ restricted to one variable is PSPACE-complete.

Structure of the paper. In Section 2, we introduce the temporal logic with repeating values, we state a few results about its expressive power and its known decidable fragments. Section 3 presents the automata-based approach with symbolic models in order to solve the satisfiability problem. The characterization of satisfiable symbolic models is provided in Section 4. Section 5 presents how satisfiable symbolic models can be recognized by a specific class of counter automata for which nonemptiness is decidable, whence the decidability of the temporal logic with repeating values. PSPACE-completeness is established in presence of a

unique variable as well as the extension of decidability with repeating constraints for past values. Concluding remarks and open problems are given in Section 6.

2 Preliminaries

2.1 Temporal Logic with Repeating Values

Let VAR $= \{x_1, x_2, \ldots\}$ be a countably infinite set of variables. The formulae of the logic CLTL$^{\Diamond}$ are defined as follows:

$$\phi ::= x = \mathsf{X}^i y \mid x = \Diamond y \mid \phi \wedge \phi \mid \neg\phi \mid \mathsf{X}\phi \mid \phi \mathsf{U} \phi \mid \mathsf{X}^{-1}\phi \mid \phi \mathsf{S} \phi$$

where $x, y \in$ VAR and $i \in \mathbb{N}$. Formulae of the form either $x = \mathsf{X}^i y$ or $x = \Diamond y$ are said to be atomic and an expression of the form $\mathsf{X}^i x$ (i next symbols followed by a variable) is called a term. Given a set of temporal operators definable from those in $\{\mathsf{X}, \mathsf{X}^{-1}, \mathsf{S}, \mathsf{U}\}$ and $k \geq 0$, we write CLTL$_k^{\Diamond}(\mathcal{O})$ to denote the fragment of CLTL$^{\Diamond}$ restricted to formulae with temporal operators from $\mathcal{O}$ and with at most k variables.

A valuation is a map VAR $\to \mathbb{N}$ and a model σ is a non-empty sequence of valuations either finite or infinite. All the subsequent developments can be equivalently done with the domain $\mathbb{N}$ replaced by an infinite set D since only equality tests are performed. We write $|\sigma|$ to denote the length of σ. The satisfaction relation is defined inductively as follows where σ is a model and $0 \leq i \leq |\sigma| - 1$:

- $\sigma, i \models x = \mathsf{X}^j y$ iff $i + j \leq |\sigma| - 1$ and $\sigma(i)(x) = \sigma(i + j)(y)$,
- $\sigma, i \models x = \Diamond y$ iff there exists $j > 0$ s.t. $i+j \leq |\sigma|-1$ and $\sigma(i)(x) = \sigma(i+j)(y)$,
- $\sigma, i \models \phi \wedge \phi'$ iff $\sigma, i \models \phi$ and $\sigma, i \models \phi'$, $\quad \sigma, i \models \neg\phi$ iff $\sigma, i \not\models \phi$,
- $\sigma, i \models \mathsf{X}\phi$ iff $i + 1 \leq |\sigma| - 1$ and $\sigma, i + 1 \models \phi$,
- $\sigma, i \models \mathsf{X}^{-1}\phi$ iff $i > 0$ and $\sigma, i - 1 \models \phi$,
- $\sigma, i \models \phi \mathsf{U} \phi'$ iff there is $i \leq j \leq |\sigma| - 1$ s.t. $\sigma, j \models \phi'$ and for every $i \leq l < j,\ \sigma, l \models \phi$.
- $\sigma, i \models \phi \mathsf{S} \phi'$ iff there is $0 \leq j \leq i$ s.t. $\sigma, j \models \phi'$ and for $j \leq l < i,\ \sigma, l \models \phi$.

We write $\sigma \models \phi$ if $\sigma, 0 \models \phi$. We shall use the standard abbreviations about the temporal operators (G, F, F^{-1}, ...) and Boolean operators ($\vee$, $\Rightarrow$, ...). We use the notation $\mathsf{X}^i x = \mathsf{X}^j y$ as an abbreviation for $\mathsf{X}^i(x = \mathsf{X}^{j-i} y)$ (when $i \leq j$).

The *finitary [resp. infinitary] satisfiability problem* consists in checking whether given a formula ϕ, there is a finite [resp. infinite] model such that $\sigma \models \phi$. It is known that finitary satisfiability for LTL can be easily reduced in logspace to infinitary satisfiability by introducing for instance an additional propositional variable p and by requiring that $p \mathsf{U} \mathsf{G} \neg p$ holds true. In that way, p holds true at every state of a prefix and p does not hold on the complement suffix. The same principle does not apply to reduce finitary satisfiability for CLTL$^{\Diamond}$ to infinitary satisfiability even by introducing additional variables in order to simulate a propositional variable. This is due to the additional atomic formulae of the form $x = \Diamond y$. That is why we distinguish the two problems in this paper.

We note that a constraint of the form x *diff* $\Diamond y$ ("the value of x differs from some future value of y") can be expressed in CLTL$^{\Diamond}$:

$$x \textit{ diff } \Diamond y \;\Leftrightarrow\; (\neg(x = \mathsf{X}y) \wedge \mathsf{X}\top) \vee ((x = \mathsf{X}y) \wedge \mathsf{X}(y = \mathsf{X}y)\mathsf{U}((y \neq \mathsf{X}y) \wedge \mathsf{X}\top))) \quad (1)$$

With infinite models, the conjunct $\mathsf{X}\top$ can be deleted. We could also introduce constraints of the form $x = \mathsf{X}^{-i}y$ but this can be expressed in the language using the equivalence $x = \mathsf{X}^{-i}y \Leftrightarrow \mathsf{X}^{-i}\top \wedge \mathsf{X}^{-i}(y = \mathsf{X}^i x)$. Similarly, CLTL$^{\Diamond}$ can express whether a variable is a nonce by the formula $\mathsf{G}\neg(x = \Diamond x)$. The formula below states a valid property when x and y are nonces:

$$(\mathsf{G}\neg(x = \Diamond x) \wedge \mathsf{G}\neg(y = \Diamond y)) \Rightarrow \mathsf{G}(x = y \Rightarrow \neg(x = \Diamond y)).$$

Other properties witnessing the high expressive power of CLTL$^{\Diamond}$ can be found in [LP05, Sect.3] about systems of pebbles evolving in time.

Apart from the above-mentioned problems, we also introduce the model-checking problem for constraint automata over CLTL$^{\Diamond}$ specifications. Constraint automata are finite-state automata with alphabet made of Boolean combinations of atomic formulae. More precisely, a constraint automaton is a tuple $\langle Q, I, F, \Sigma, \delta\rangle$ such that

- Q is a finite set of states,
- $I, F \subseteq Q$ are respectively the set of initial and final states,
- Σ is a finite set of Boolean combination of constraints of the form either $x = \mathsf{X}^i y$ or $x = \Diamond y$ where $x, y \in \mathrm{VAR}$ and $i \in \{0, 1\}$,
- $\delta \subseteq Q \times \Sigma \times Q$.

A run of $\mathcal{A}$ is an infinite sequence of the form $q_0 \xrightarrow{\varphi_0} q_1 \xrightarrow{\varphi_1} q_2 \xrightarrow{\varphi_2} \ldots$ such that $\langle q_i, \varphi_i, q_{i+1}\rangle \in \delta$ for every $i \in \mathbb{N}$. We say that an infinite model σ *realizes* a run of $\mathcal{A}$ if for every $i \in \mathbb{N}$ we have $\sigma, i \models \varphi_i$.

The model-checking problem takes as input a CLTL$^{\Diamond}$ formula ϕ and a constraint automaton $\mathcal{A}$ and answers positively iff there is an infinite model σ that both realizes a run of $\mathcal{A}$ and satisfies ϕ. This problem can also be reduced to the infinitary satisfiability problem since we can encode the execution of a constraint automaton into a CLTL$^{\Diamond}$ formula. In the rest of the paper, we restrict ourselves to satisfiability but all the decidability results can be lifted to model-checking.

2.2 Known extensions of CLTL$^{\Diamond}$

In this section, we recall the definition of a few known extensions of CLTL$^{\Diamond}$ which is useful for future comparisons. The logic CLTL$^{\Diamond}$ is clearly a fragment of the logic CLTL$^{\downarrow}(\mathbb{N}, =)$ introduced in [DLN07] and restricted to one register. An equivalent logic of CLTL$^{\downarrow}(\mathbb{N}, =)$ is denoted by CLTL$^{\downarrow}$ in this paper and it is defined as follows. We consider an additional set of registers $\mathrm{REG} = \{r_1, r_2, \ldots\}$ and the formulae of CLTL$^{\downarrow}$ are defined as those of CLTL$^{\Diamond}$ except that:

- we allow only atomic formulae of the form $r = x$ where $r \in \mathrm{REG}$ and $x \in \mathrm{VAR}$,

– we add the inductive clause $\downarrow_{r=x} \phi$.

The satisfaction relation is parameterized by a register assignment $\rho : \mathrm{REG} \to \mathbb{N}$ with $\sigma, i \models_\rho \downarrow_{r=x} \phi$ iff $\sigma, i \models_{\rho[r \mapsto \sigma(i)(x)]} \phi$ and $\sigma, i \models_\rho r = x$ iff $\rho(r) = \sigma(i)(x)$. Consequently, the atomic formula $x = \Diamond y$ in CLTL$^{\Diamond}$ can be naturally encoded in CLTL$^{\downarrow}$ by $\downarrow_{r=x} \mathsf{XF}(r = y)$ and $x = \mathsf{X}^i y$ by $\downarrow_{r=x} \mathsf{X}^i(r = y)$.

We write $\mathrm{CLTL}^{\downarrow}_{(k,k')}(\mathcal{O})$ to denote the fragment of CLTL$^{\downarrow}$ restricted to the temporal operators from $\mathcal{O}$ with at most k variables and k' registers. Following the notation from [DL06], for $\alpha \geq 0$ we write $\mathrm{LTL}^{\downarrow}_{\alpha}(\sim, \mathcal{O})$ to denote the fragment $\mathrm{CLTL}^{\downarrow}_{(1,\alpha)}(\mathcal{O})$ restricted to atomic formulae of the form $r = x$.

2.3 A decidable fragment of finitary satisfiability

It is shown in [DLN07] that CLTL$^{\downarrow}$ is strictly more expressive than its freeze-free fragment. The same argument applies to show that CLTL$^{\Diamond}$ is strictly more expressive than its fragment without atomic formulae of the form $x = \Diamond y$. Observe also that CLTL$^{\Diamond}$ is neither a fragment of the pure-future safety fragment in [Laz06] (where occurrences of U formulae are never in the scope of an even number of negations) nor a fragment of the flat fragment of CLTL$^{\downarrow}$. Unlike these fragments, CLTL$^{\Diamond}$ contains past-time operators and negation can be used without any restriction. Infinitary satisfiability for safety $\mathrm{LTL}^{\downarrow}_1(\sim, \mathsf{X}, \mathsf{U})$ is EXPSPACE-complete [Laz06], for full $\mathrm{LTL}^{\downarrow}_1(\sim, \mathsf{X}, \mathsf{U})$ is Π^0_1-complete and, finitary and infinitary satisfiability for flat CLTL$^{\downarrow}$ are PSPACE-complete. By contrast, in this paper we show that finitary and infinitary satisfiability for CLTL$^{\Diamond}$ (with full past-time temporal operators) are decidable problems. By taking advantage of [DLN07,DL06], it is already possible to establish decidability of *finitary* satisfiability for *strict* fragments of CLTL$^{\Diamond}$.

Theorem 1. **(I)** *Finitary satisfiability for* $\mathrm{CLTL}^{\Diamond}(\mathsf{X}, \mathsf{U})$ *is decidable.*
(II) *Finitary satisfiability for* $\mathrm{CLTL}^{\Diamond}(\mathsf{X}, \mathsf{X}^{-1}, \mathsf{F}, \mathsf{F}^{-1})$ *is decidable.*

Showing similar results in the infinitary case with the same approach of the proof seems difficult since infinitary satisfiability for $\mathrm{LTL}^{\downarrow}_1(\sim, \mathsf{X}, \mathsf{F})$ is Π^0_1-complete. In the paper we shall show much stronger results: finitary and infinitary satisfiability for full CLTL$^{\Diamond}$ even augmented with MSOL definable temporal operators are decidable. In the rest of the paper, we systematically treat the finitary and infinitary cases simultaneously. However, we provide the full technical details for the infinitary case only and we sketch the main ideas for the finitary case. This latter case cannot be reduced in the obvious way to the infinitary case but its solution is close to the one for the infinitary case.

3 Automata-based Approach with Symbolic Models

In this section we explain how satisfiability can be solved using symbolic models which are abtractions of concrete CLTL$^{\Diamond}$ models. We provide the outline of our automata-based approach, and consider the technical details in Sects. 4 and 5.

3.1 Symbolic Models

Let ϕ be a CLTL$^{\Diamond}$ formula with k variables $\{x_1, \ldots x_k\}$ and we write l (the "X-length" of ϕ) to denote the maximal i such that a term of the form $\mathsf{X}^i x$ occurs in ϕ. In order to define the set of atomic formulae that are helpful to determine the satisfiability status of ϕ, we introduce the set of constraints Ω_k^l that contains constraints of the form either $\mathsf{X}^i x = \mathsf{X}^j y$ or $\mathsf{X}^i(x = \Diamond y)$ and their negation with $x, y \in \{x_1, \ldots x_k\}$ and $i, j \in \{0, \ldots, l\}$. Models are abstracted as sequences of frames that are defined as maximally consistent subsets of Ω_k^l.

An *l-frame* is a set $fr \subseteq \Omega_k^l$ that is maximally consistent in that it satisfies the conditions below:

(F1) For every constraint $\varphi \in \Omega_k^l$, either φ or $\neg\varphi$ belongs to fr but not both.
(F2) For all $i \in \{0, \ldots, l\}$ and $x \in \{x_1, \ldots x_k\}$, $\mathsf{X}^i x = \mathsf{X}^i x \in fr$.
(F3) For all $i, j \in \{0, \ldots, l\}$ and $x, y \in \{x_1, \ldots x_k\}$, $\mathsf{X}^i x = \mathsf{X}^j y \in fr$ iff $\mathsf{X}^j y = \mathsf{X}^i x \in fr$.
(F4) For all $i, j, j' \in \{0, \ldots, l\}$ and $x, y, z \in \{x_1, \ldots, x_k\}$, $\{\mathsf{X}^i x = \mathsf{X}^j y, \mathsf{X}^j y = \mathsf{X}^{j'} z\} \subseteq fr$ implies $\mathsf{X}^i x = \mathsf{X}^{j'} z \in fr$.
(F5) For all $i, j \in \{0, \ldots, l\}$ and $x, y \in \{x_1, \ldots x_k\}$ such that $\mathsf{X}^i x = \mathsf{X}^j y \in fr$:
- if $i = j$, then for every $z \in \{x_1, \ldots, x_k\}$ we have $\mathsf{X}^i(x = \Diamond z) \in fr$ iff $\mathsf{X}^j(y = \Diamond z) \in fr$;
- if $i < j$ then $\mathsf{X}^i(x = \Diamond y) \in fr$, and for $z \in \{x_1, \ldots, x_k\}$, $\mathsf{X}^i(x = \Diamond z) \in fr$ iff either $\mathsf{X}^j(y = \Diamond z) \in fr$ or there exists $i < j' \leq j$ such that $\mathsf{X}^i x = \mathsf{X}^{j'} z \in fr$.

Conditions (F2)–(F4) simply encode that equality is an equivalence relation.

For an l-frame fr, $x \in \{x_1, \ldots, x_k\}$ and $i \in \{0, \ldots, l\}$, we define

- the set of future obligations for x at level i in fr as $\Diamond_{fr}(x, i) \stackrel{\text{def}}{=} \{y \mid \mathsf{X}^i(x = \Diamond y) \in fr\}$,
- the equivalence class of x at level-i in fr as $[(x, i)]_{fr} \stackrel{\text{def}}{=} \{y \mid \mathsf{X}^i x = \mathsf{X}^i y \in fr\}$.

An l-frame fr can be represented as an annotated undirected graph G_{fr} which has vertices (x, i) for $x \in \{x_1, \ldots, x_k\}$ and $i \in \{0, \ldots, l\}$, and an edge between (x, i) and (y, j) iff the constraint $\mathsf{X}^i x = \mathsf{X}^j y$ belongs to fr. A vertex (x, i) in the graph is annoted with an "open" arc labelled by the set of future obligations $\Diamond_{fr}(x, i)$ for that vertex. Fig. 1 shows an example of a 3-frame over the variables $\{x, y, z\}$. For convenience we avoid showing transitively inferable edges.

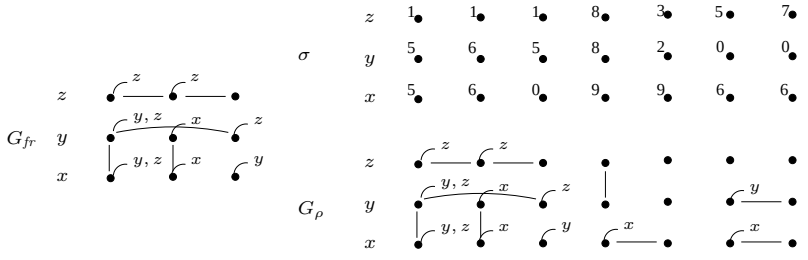

Fig. 1. Example 3-frame graph, concrete model σ, and its induced 3-frame graph G_ρ.

We denote by $\texttt{Frame}_k^l$ the set of such frames built w.r.t. k and l. We say that a model σ satisfies a frame fr at position i (denoted $\sigma, i \models fr$) iff $\sigma, i \models \varphi$ for every constraint φ in fr.

Since a frame can be viewed as a set of constraints about $l+1$ consecutive positions, for finitary satisfiability we need to add an information about the possibility to end the model before the end of the current window of length $l+1$. This can be done with $\mathcal{O}(l)$ bits and then the conditions (F1)–(F5) need to be updated accordingly in order to take into account this possibility. Note that this method allows to handle the particular case where there exists a model whose size is smaller than the X-length of the formula.

Lemma 1. *For all models σ with k variables and $0 \leq i \leq |\sigma| - 1$, there exists a unique frame $fr \in \texttt{Frame}_k^l$ such that $\sigma, i \models fr$.*

Proof. We treat only the infinitary case, the finitary case being similar with the proper notion of frame. We can easily show that $\{\{\sigma \mid \sigma \models fr\} \mid fr \in \texttt{Frame}_k^l\}$ is a partition of the set of models $\mathbb{N} \to (\{x_1, \ldots, x_k\} \to \mathbb{N})$. So the i^{th} prefix of σ has a unique frame abstraction. □

A pair of l-frames $\langle fr, fr' \rangle$ is said to be *one-step consistent* $\overset{\text{def}}{\Leftrightarrow}$

- for all $0 < i, j \leq l$, $\mathsf{X}^i x = \mathsf{X}^j y \in fr$ iff $\mathsf{X}^{i-1} x = \mathsf{X}^{j-1} y \in fr'$,
- for all $0 < i \leq l$, $\mathsf{X}^i(x = \Diamond y) \in fr$ iff $\mathsf{X}^{i-1}(x = \Diamond y) \in fr'$.

A *symbolic model* of X-length l is a (finite or infinite) sequence of l-frames ρ such that for $0 \leq i < |\rho| - 1$, $\langle \rho(i), \rho(i+1) \rangle$ is one-step consistent. We define the symbolic satisfaction relation $\rho, i \models_{\text{symb}} \phi$, for a formula ϕ of X-length l and a symbolic model ρ of X-length l, as done for CLTL$^\Diamond$ except that for atomic formulas φ we have: $\rho, i \models_{\text{symb}} \varphi \overset{\text{def}}{\Leftrightarrow} \varphi \in \rho(i)$. We say a model σ *realizes* a symbolic model ρ (or equivalently that ρ *admits* a model σ) $\overset{\text{def}}{\Leftrightarrow}$ for every $0 \leq i \leq |\sigma| - 1$, we have $\sigma, i \models \rho(i)$.

A symbolic model ρ of X-length l can also be represented as an annotated graph G_ρ in a similar manner to l-frames. Thus the vertices of G_ρ are of the form (x, i) with an edge between (x, i) and (y, j) with $0 \leq j - i \leq l$ iff there was an edge between $(x, 0)$ and $(y, j - i)$ in the frame graph $G_{\rho(i)}$. The annotations for future obligations are added similarly. Fig. 1 shows the graph representation of a symbolic model ρ of X-length 3, and a model it admits. By a *path* p in G_ρ we will mean as usual a (finite or infinite) sequence of vertices $v_0, v_1 \ldots$ in G_ρ such that each v_i, v_{i+1} is connected by an edge in G_ρ. We call p a *forward* path if each v_{i+1} is at a level strictly greater than v_i.

3.2 Automata for Symbolic Models

In order to check whether a CLTL$^\Diamond$ formula is satisfiable we use Lemma 2 below based on the approach developed in [DD07].

Lemma 2. *A* CLTL$^\Diamond$ *formula ϕ of* X*-length l is satisfiable iff there exists a symbolic model ρ of* X*-length l such that $\rho \models_{\text{symb}} \phi$ and ρ admits a model.*

Proof. Only the infinite case is explicitly presented. Suppose that ϕ is satisfiable and let $\sigma \models \phi$. From σ we can build a symbolic model ρ such that for every $i \geq 0$, we have $\sigma, i \models \rho(i)$. By Lemma 1, ρ is unique. By definition of a frame, if σ satisfies an atomic constraint $\varphi \in \Omega_k^l$ at position i then $\varphi \in \rho(i)$ (consequence of **(F1)**) and so $\rho, i \models_{\text{symb}} \varphi$. Using this property we can prove by induction on the structure of ϕ that $\rho \models_{\text{symb}} \phi$.

Conversely, suppose there exists σ and ρ such that $\rho \models_{\text{symb}} \phi$ and $\sigma \models \rho$. By definition we have $\sigma, i \models \rho(i)$ for every $i \geq 0$ and $\rho, i \models_{\text{symb}} \varphi$ implies $\varphi \in \rho(i)$. So for every atomic formula $\varphi \in \Omega_k^l$ if $\rho, i \models_{\text{symb}} \varphi$ then $\sigma, i \models \varphi$. By induction on the structure of ϕ, we get $\sigma \models_{\text{symb}} \phi$. □

In order to take advantage of Lemma 2, we use the automaton-based approach from [VW94]. We build an automaton $\mathcal{A}_\phi$ as the intersection of two automata $\mathcal{A}_{\text{symb}}$ and $\mathcal{A}_{\text{sat}}$ such that the language recognized by $\mathcal{A}_{\text{symb}}$ is the set of symbolic models satisfying ϕ and the language recognized by $\mathcal{A}_{\text{sat}}$ is the set of symbolic models that are realized by some models.

We define the automaton $\mathcal{A}_{\text{symb}}$ by adapting the construction from [VW94] for LTL. We define $cl(\phi)$ the closure of ϕ as usual, and an atom of ϕ is a maximally consistent subset of $cl(\phi)$. For the infinitary case, $\mathcal{A}_{\text{symb}}$ is the generalized Büchi automaton $(Q, Q_0, \rightarrow, F)$ such that:

- Q is the set of atoms of ϕ and $Q_0 = \{At \in Q \mid \phi \in At,\ \mathsf{X}^{-1}\top \notin At\}$,
- $At \xrightarrow{fr} At'$ iff
  (atomic constraints) for every atomic formula φ in At, $\varphi \in fr$,
  (one step) for every $\mathsf{X}\psi \in cl(\phi)$, $\mathsf{X}\psi \in At$ iff $\psi \in At'$, and for every $\mathsf{X}^{-1}\psi \in cl(\phi)$, $\psi \in At$ iff $\mathsf{X}^{-1}\psi \in At'$
- let $\{\psi_1 \mathsf{U} \phi_1, \ldots, \psi_r \mathsf{U} \phi_r\}$ be the set of until formulae in $cl(\phi)$. We pose $F = \{F_1, \ldots, F_r\}$ where for every $i \in \{1, \ldots, r\}$, $F_i = \{At \in Q : \psi_i \mathsf{U} \phi_i \notin At \text{ or } \phi_i \in At\}$.

For the finitary case, the finite-state automaton $\mathcal{A}_{\text{symb}}$ accepting finite words is defined as above except that F is a set of states At such that

- no atomic formula of the form $x = \Diamond y$ occurs in At,
- no formula of the form either $\mathsf{X}\phi$ or $x = \mathsf{X}^i y$ with $i > 0$ occurs in At.

Moreover, such final states can only be reached when the frame labelling the last transition contains proper information about the end of the model.

In the next section, we explain how one can build the automaton $\mathcal{A}_{\text{sat}}$ that recognizes the set of satisfiable symbolic models. Since $\mathcal{A}_\phi$ is the automaton recognizing the intersection of the languages accepted by $\mathcal{A}_{\text{symb}}$ and $\mathcal{A}_{\text{sat}}$, the following result is a direct consequence of Lemma 2.

Theorem 2. *A* CLTL$^\Diamond$ *formula ϕ is satisfiable iff the language recognized by $\mathcal{A}_\phi$ is nonempty.*

Note that we separate the temporal logic part and the constraint part by defining two different automata. This allows to extend the decidability results to any extension of LTL that induces an ω-regular class of models. We only need to change the definition of $\mathcal{A}_{\text{symb}}$.

4 Characterization of Satisfiable Symbolic Models

In order to determine whether a symbolic model ρ is "satisfiable" (i.e. it admits a model), we introduce counters that remember the satisfaction of constraints $x = \Diamond y$. If $x = \Diamond y_1 \wedge \cdots \wedge x = \Diamond y_n$ needs to be satisfied at the current position, then we shall increment a counter indexed by $\{y_1, \ldots, y_n\}$ that remembers this set of obligations. In a finite model, all the obligations need to be fulfilled before the last position whereas in an infinite model either no more unsatisfied obligations arise after a point, or they are essentially fulfilled infinitely often. The exact conditions will be spelt out soon.

4.1 Counting Sequence

For each $X \in \mathcal{P}^+(\{x_1, \ldots, x_k\})$ (set of non-empty subsets of $\{x_1, \ldots, x_k\}$), we introduce a counter that keeps track of the number of obligations that need to be satisfied by X. We identify the counters with elements of $\mathcal{P}^+(\{x_1, \ldots, x_k\})$. A counter valuation $\boldsymbol{c}$ is a map $\boldsymbol{c} : \mathcal{P}^+(\{x_1, \ldots, x_k\}) \to \mathbb{N}$. For instance, we write $\boldsymbol{c}(\{x, y\})$ to denote the value of the counter $\{x, y\}$, which will stand for the number of obligations to repeat a distinct value in x and y.

We will define a canonical sequence of counter valuations along a symbolic model. We introduce some definitions first. For an l-frame fr and $X \in \mathcal{P}^+(\{x_1, \ldots, x_k\})$, we define a *point of increment* for X in fr to be an equivalence class of the form $[(x, 0)]_{fr}$ such that $\Diamond_{fr}(x, 0) = X$ and $(x, 0)$ is not connected by a forward edge to a node in fr (i.e. there is no edge between $(x, 0)$ and (y, j) for any $j \in \{1, \ldots l\}$). A *point of decrement* for X in fr is defined to be an equivalence class of the form $[(x, l)]_{fr}$ such that $\Diamond_{fr}(x, l) \cup [(x, l)]_{fr} = X$, and (x, l) is not connected by a backward edge to another node in fr (i.e. there is no edge between (x, l) and (y, j) for any $j \in \{0, \ldots l-1\}$). Let u^+_{fr} denote a counter valuation which records the number of points of increment for each counter X, in fr. Similarly let u^-_{fr} denote the counter valuation which records the number of points of decrement for each counter X in fr.

Now let ρ be a symbolic model of X-length l. We carry over the notations for the set of future obligations and the equivalence class for x at level i to symbolic models as well. Thus $\Diamond_\rho(x, i)$ is equal to $\Diamond_{\rho(i)}(x, 0)$ and $[(x, i)]_\rho$ is $[(x, 0)]_{\rho(i)}$. For $X \in \mathcal{P}^+(\{x_1, \ldots, x_k\})$, a *point of increment* for X in ρ is an equivalence class of the form $[(x, i)]_\rho$ such that $[(x, 0)]_{\rho(i)}$ is a point of increment for X in the frame $\rho(i)$. Similarly, a *point of decrement* for X in ρ is an equivalence class of the form $[(x, i)]_\rho$ such that $i \geq l + 1$ and $[(x, l)]_{\rho(i-l)}$ is a point of decrement for X in the frame $\rho(i - l)$.

We can now define a canonical counter valuation sequence α along ρ, called the *counting sequence* along ρ, which counts the number of "unsatisfied" points of increments for each counter X. Let $\dot{+}$ denote the "proper addition" of integers, defined by $n \dot{+} m = \max(0, n + m)$. We define α inductively for each $X \in \mathcal{P}^+(\{x_1, \ldots, x_k\})$ and $0 \leq i < |\rho|$ as: $\alpha(0)(X) = 0$; and $\alpha(i + 1)(X) = \alpha(i)(X) \dot{+} (u^+_{\rho(i)}(X) - u^-_{\rho(i+1)}(X))$.

4.2 Characterising Satisfiable Symbolic Models

We characterize satisfiable symbolic models using their counting sequences.

Lemma 3. *A finite symbolic model ρ is satisfiable (i.e. admits a model) iff the final value of the counting sequence α along ρ has value 0 for each counter X (i.e. $\alpha(|\rho|-1)(X) = 0$) and in the last frame fr of ρ, there are no "unsatisfied"obligations (where by an unsatisfied obligation in fr we mean a node (x,i) in G_{fr} with a variable $y \in \Diamond_{fr}(x,i)$, but no edge from (x,i) to (y,j) for any $j > i$).*

An infinite symbolic model ρ is satisfiable iff the following conditions are satisfied:

(C1) There does not exist an infinite forward path p in ρ and a counter X, such that every node in the path has future obligation X, and there is a variable y in X which is never connected by a forward edge from a node in p (i.e. no node in p is connected by a forward edge to a node of the form (y,i)).

(C2) In the counting sequence along ρ, each counter X satisfies one of the conditions:

- *(a) there is a point after which the value of counter X is always zero and after which we never see a point of increment for X, or,*
- *(b) infinitely often we see a point of decrement for X which is connected by a forward path to a point of increment of the form $[u]_\rho$ with $\Diamond_\rho(u) \subset X$ (where '$\subset$' denotes "strict subset"), or,*
- *(c) for each $x \in X$, we infinitely often see a point of decrement for X, which is connected by a forward path to an x node (i.e a node of the form (x,i)).*

Proof. Let ρ be a symbolic model of X-length l, which admits a concrete model σ. We show that ρ satisfies the conditions above.

Consider a point of increment $[(x,i)]_\rho$ for a counter X. Then in the concrete model σ, the value $\sigma(i)(x)$ subsequently repeats in all the variables in X. Let (y,j) be the *first* time this happens. So $y \in X$ and $\sigma(j)(y) = \sigma(i)(x)$. We claim that $[(y,j)]_\rho$ must be a point of decrement for X. This is true since by the choice of (y,j), it cannot be connected to any node to the left of it, and clearly $\Diamond_\rho(x,i) = [(y,j)]_\rho \cup \Diamond_\rho(y,j)$. Further, the correspondence between points of increment and points of decrement described above is injective. If not, let $[(x,i)]_\rho$ and $[(x',i')]_\rho$ be two distinct points of increment with the same corresponding point of decrement $[(y,j)]_\rho$. Without loss of generality, we assume $i \leq i'$. If $i = i'$, it would mean $\sigma(i)(x) = \sigma(i)(x')$ (since they both have the same value as (y,j) by assumption), which would contradict the fact that $[(x,i)]_\rho$ and $[(x',i')]_\rho$ were assumed to be distinct. If $i < i'$, then (y,j) could not have been the first repeat for (x,i) since (x',i') is one such repeat and it occurs strictly before (y,j).

Now if ρ was a finite sequence, then clearly the value of each counter X is 0 in $\alpha(|\rho|-1)$. This is because of the fact that by the above correspondence, each point of increment for X is cancelled out by a unique point of decrement for X. Furthermore, the last frame in ρ clearly cannot have any unsatisfied obligations. This proves that the conditions of Lemma 3 are satisfied for the case of finite symbolic models.

Consider now the case when ρ is an infinite symbolic model. We show that conditions (C1) and (C2) are satisfied. Let p be an infinite forward path from a vertex u in G_ρ and let $y \in \Diamond_\rho(u)$. Since $y \in \Diamond_\rho(u)$, it must be the case that the value of u in the concrete model σ repeats at some future point in a y-node, say (y, j). Now the path p must pass through a node v in the l-frame $\rho(j-l)$ to which (y, j) belongs. Since the value of v must be same as that of u, which in turn is same as that of (y, j), there must be an edge between v and (y, j) in $\rho(j-l)$. This proves that ρ satisfies the condition (C1).

To see that condition (C2) is satisfied, let X be any counter. Two cases arise: either we have only finitely many points of decrement for X in ρ, or there are infinitely many. For the first case, let i be the level at which the last point of decrement corresponding to a point of increment for X occurs. Then it is clear that $\alpha(i-l)(X) = 0$. Further, by the choice of i, we never see a point of increment for X after level i, and the value of X in α stays 0. Thus in this case condition C2(a) is satisfied.

For the case when there are infinitely many points of increment for X, suppose that X satisfies neither C2(b) nor C2(c). Then there must be a variable $y \in X$ and a level i after which we never see a point of decrement for X which is connected by a forward path to a point of increment for X with future obligation strictly smaller than X, nor a point of decrement for X which is connected by a forward path to a y-node. Consider any point of increment $[u_0]_\rho$ for X after level i. Let its value in the concrete model be m. In the concrete model, m must subsequently repeat in a y-node. Let this node be (y, j). Now for $[u_0]_\rho$ there is a corresponding point of decrement $[v_0]_\rho$ for X (obtained as above by taking the first node where the value m repeats). Note that there cannot exist an infinite forward path from v_0, since otherwise by an argument similar to the one for C1 above, we would have a forward path from v_0 to (y, j), contradicting our assumption. So there is a maximal forward path (possibly of length 0) from v_0 to a node u_1, which (again by our assumptions) must necessarily be such that $[u_1]_\rho$ forms a point of increment for X. This argument can be repeated to construct a sequence of nodes $u_0, v_0, u_1, v_1, \ldots$ such that each $[u_i]_\rho$ and $[v_i]_\rho$ are respectively points of increment and decrement for X, and there is a forward path from each v_i to each u_{i+1}. This is shown below:

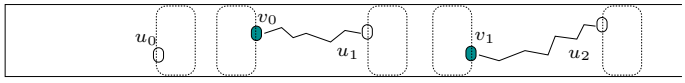

It is also clear that by construction, all the nodes above (as well as the nodes in the paths between the v_i's and u_{i+1}'s) have the value m in the concrete model σ. Now consider the node (y, j). Clearly it cannot lie between any u_i and v_i. Thus it must lie between some v_i and u_{i+1}, and must be connected by a forward edge from a node in the path between them. This contradicts our assumption that after level i, no point of decrement for X was connected by a forward path to a y-node.

Thus for the case when ρ is infinite, we have shown that ρ must satisfy the conditions of Lemma 3.

For the converse direction, let ρ be a symbolic model of X-length l satisfying the conditions of Lemma 3. We will show that ρ admits a concrete model by first describing an augmented graph G'_ρ which is obtained from G_ρ by adding additional edges, and then describing a labelling procedure for G_ρ which respects the edges in G_ρ.

For the case when ρ is finite, the augmented graph G'_ρ is obtained from G_ρ by adding edges (which we call *augmented* edges) as follows:

> From each level i going from $l+1$ upto $|\rho|-1$, for each counter X, and for each point of decrement $[v]_\rho$ for X at level i, if there is a point of increment $[u]_\rho$ for X at a level less than $i-l$ for which an augmented edge has not been added (we call this an "unmatched" point of increment), add augmented edges between every node in $[u]_\rho$ and every node in $[v]_\rho$.

Here is (the only) way of adding augmented edges according to the procedure above in the example symbolic model of Fig. 1:

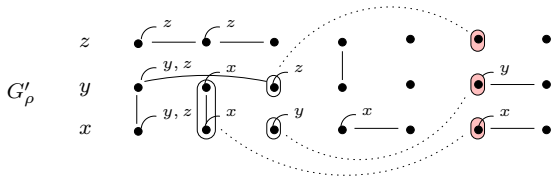

We note that when the procedure has completed for level i, the number of unmatched points of increment for any counter X is precisely $\alpha(i)(X)$, the value for the counter X at position i in the counting sequence α for ρ. Since $\alpha(|\rho|-1)(X)=0$ for each X, it follows that at the end of the procedure above, we will have no unmatched points of increment.

For the case when ρ is infinite, we add the augmented edges in a slightly different way. If a counter X is such that a point of increment for it occurs only finitely often in G_ρ, we add the augmented edges between points of increment and points of decrement for X in the same way as the procedure above for the case of finite models. By condition C2(a), every point of increment for X will be matched. If X is a counter for which points of increment occur infinitely often, then by conditions C2(b) and C2(c) two cases are possible: either there are infinitely many points of decrement for X which are connected by a forward path to a point of increment with a strictly smaller set of obligations than X, or for each $x \in X$, we infinitely often see a point of decrement for X which is connected by a forward path to an x-node. In the former case, we assign matches by proceeding from left to right, adding augmented edges from each point of increment for X to a subsequent point of decrement for X from which there is a forward path to a point of increment with a strictly smaller set of obligations than X. In the latter case, suppose $X = \{z_1, \ldots, z_m\}$. We assign matches by cycling through the z_i's repeatedly: thus, we proceed from left to right, and assign to the first point of increment for X, a point of decrement for X that is connected to a z_1-node, to the next point of increment for X a point of decrement for X that is connected by a forward path to z_2 node, and so on till z_m; and keep

repeating this process. Thus, this process of matching points of increment covers all points of increment, and so every point of increment has an augmented edge to a subsequent point of decrement.

We now describe a way of labelling the nodes of G_ρ with natural numbers. We use a natural ordering on nodes in G_ρ, given by $(x_m, i) \prec (x_p, j)$ iff $i < j$ or, $i = j$ and $m < p$.

> We label the first vertex $(x_1, 0)$ by 0. The remaining vertices are labelled in order according to the following rule: If L is the portion of the graph already labelled, and u is the next vertex to be labelled:
> 1. if there is a path in G'_ρ from u to a vertex v in L, give u the same label as v.
> 2. else, label u by $n + 1$ where n is the maximum label used so far in L.

We note that the labelling above is deterministic, in the sense that in step 1 u can only be assigned a single value. If not, consider the first point that a vertex u had a path to two vertices v and v' with distinct labels. Without loss of generality, say v was labelled before v'. Then there is a path from v' to v in G'_ρ (via u), and hence v' must have been labelled with the same value as v.

The labelling above thus gives us a concrete model σ, and we claim that ρ is in fact the symbolic model of X-length l induced by σ. For this it is sufficient to argue that the labelling σ respects all the edges of G_ρ (i.e. the normal as well as annotated edges).

Before we do this let us first observe a useful property of G'_ρ.

Claim (1). Let $u = (x, i)$ and $v = (y, j)$ be distinct vertices which are connected by a path in G'_ρ. Then

1. if the level of u is strictly less than the level of v (i.e. $i < j$), we have a *forward* path from u to v in G'_ρ.
2. if u and v are at the same level (i.e. $i = j$), we have an edge between u and v.

Proof (of claim). We proceed by induction on the length of the shortest path between u and v. In fact we show that the shortest path must be a forward path in the case of $i < j$, and a single edge in the case of $i = j$. If the shortest path between u and v is of length 1, then if $i < j$ we have a forward edge from u to v, and if $i = j$, an edge at level i between u and v.

For the induction step, let us assume it holds for nodes connected by a shortest path of length m or less, and suppose the shortest path between u and v is of length $m + 1$. Consider the case when $i < j$. Let the first node on this shortest path after u be w. Now w cannot be to the left of u: for if this was the case, there must be either an original edge (i.e. an edge of G_ρ) from u to w, or an augmented edge. Suppose it was an original edge, then by induction hypothesis we have a forward path from w to v. The first edge in this forward path from w clearly cannot be an augmented edge, since a forward augmented edge must begin at a point of increment, and w cannot be a point of increment since it is connected by a forward edge to u. Neither can the first edge in the forward path

from w be an original edge, since then we must have an edge between u and the target vertex of this edge, which gives us a strictly shorter path from u to v. For the case when have an augmented edge from u to w, the forward path from w to v must pass through this or an "equivalent" edge, i.e. an edge from w to u' where $u' \in [u]_\rho$. In either case, we have a contradiction to the fact that we had started out with a shortest path from u to v.

Similarly, the first edge from u to w cannot be an edge at the same level, since this would again contradict the fact that it was part of the shortest path from u to v.

Hence it must be a forward edge from u to w, and using the induction hypothesis we obtain a forward path from u to v.

The case when $i = j$ is handled similarly. This completes the proof of our claim. □

Here are some more properties of G'_ρ which are easily verified:

Claim (2).

1. If we have forward path from a node u to an x-node in G'_ρ then we must have $x \in \Diamond_\rho(u)$.
2. If we have a forward path p from u to v in G'_ρ, and $x \in \Diamond_\rho(u)$, then either some node in p is connected by a forward edge to an x-node, or $x \in \Diamond_\rho(v)$. □

We can now prove that the labelling σ of G_ρ is an edge-respecting one. We first argue that σ respects the normal edges in G_ρ. Let (u, v) be an edge in G_ρ, with say $u \prec v$. Then by the deterministic property of the labelling procedure, v would be given the same value as u in step 1. Further, if u and v were in the same l-frame (i.e. their levels differ by at most $l + 1$) and there was no edge between them in G_ρ, then we argue that they would be given different labels. Suppose to the contrary that u and v were given the same label m. Then it is easy to see that both u and v must be connected to a vertex w which was the first vertex to be labelled m. Thus there is a path between u and v, and by Claim 1 we must have a forward path from u to v. Since u and v lie in a common l-frame, this means that they must be connected by an edge in G_ρ.

Let us now consider the annotated edges in G_ρ and show that the labelling σ respects these edges also. Let u be a node in G_ρ with a variable $x \in \Diamond_\rho(u)$. Let us first consider the case when ρ is finite. Consider a maximal forward path from u in G'_ρ, and let w be the final node in this path. Then w cannot have any future obligations: for otherwise w would either be a point of increment, in which case we could have continued the path by an augmented edge, contradicting the assumption that the path was maximal; or w would lie in the final l-frame of ρ, contradicting the assumption that ρ satisfies the conditions of Lemma 3. So in particular $x \notin \Diamond_\rho(w)$. Thus by Claim 2, we have that there must be a vertex v in the path which is connected by a forward edge to an x-node. This x-node

would have been labelled m, and we have a required x-node which satisfies the obligation x of u.

For the case when ρ is infinite, let $\Diamond_\rho(u) = X$ (with $x \in X$). We argue by induction on the size of X, that the x-obligation of u is satisfied. When X is a singleton, i.e. $X = \{x\}$, suppose u was not connected by a forward path to an x-node in G'_ρ. Then it must be the case (using Claim 2) that we have an infinite forward path from u in G'_ρ, along which all nodes have future obligation exactly $\{x\}$. If this path uses only finitely many augmented edges, we have a contradiction of the assumption that ρ satisfies the condition C1 of Lemma 3. If the path uses infinitely many augmented edges, we must have infinitely many points of increment for $\{x\}$ along the path, and by the way we added the augmented edges in G'_ρ, these augmented edges are to points of decrement for $\{x\}$ which are necessarily connected by a forward path to an x-node. For the induction step, suppose X had more than 1 element, and suppose once again that u was not connected by a forward path to any x-node in G'_ρ. Then again, there must be an infinite forward path from u in G'_ρ along which the obligation x is preserved. Since the obligations along a forward path can only decrease (or stay the same), it must be the case that after a point the set of future obligations remains at some X' with $x \in X'$. Again, if this path had only finitely many augmented edges, it would contradict condition C1. Otherwise, it has infinitely many augmented edges (and hence points of increment for X') and by the way augmented edges were added, it must be the case that these edges are to points of decrement for X' from which there is a forward path to a point of increment for X'' with $X'' \subset X'$. By our induction hypothesis, there is a path from these points of increment for X'' to a y-node for each $y \in X''$. Since $x \in X''$, we have a path from u to an x-node, contradicting our assumption. Thus, it cannot be the case that there is no forward path from u tox an x-node. Thus there is a forward path from u to an x-node, say v. By our labelling procedure, v would also be given the same value as u. Thus the obligation x in $\Diamond_\rho(u)$ is satisfied.

We now argue that for any node u in G_ρ that if $x \notin \Diamond_\rho(u)$ then there is no x-node at a level greater than that of u which is given the same label. Suppose some such x-node v was given the same label, say m, as u. Then, as similarly observed before, both u and v must be connected by a path in G'_ρ (via the first vertex labelled m). By Claim 1, there is a forward path from u to v. By Claim 2(1), we must have $x \in \Diamond_\rho(u)$.

This completes the proof of the fact that the labelling σ respects the edges of G_ρ, and hence ρ is the symbolic model induced by σ.

With this we have finally completed the proof of Lemma 3. □

In Sect. 5 we show that we can check these conditions on counting sequences using counter automata with a decidable nonemptiness problem.

5 Decidability

We introduce a class of counter automata with a disjunctive variant of generalized Büchi acceptance condition in which, along any run, a zero test is performed at most once for each counter.

5.1 Simple Counter Automata

A simple counter automaton $\mathcal{A}$ is a tuple $\langle \Sigma, C, Q, \mathcal{F}, I, \rightarrow \rangle$ such that

- Σ is a finite alphabet, C is a finite set of counters,
- Q is a finite set of locations, $I \subseteq Q$ is the set of initial locations,
- $\mathcal{F} = \{F_0, F_1, \ldots, F_K\}$ for some $K \geq 0$ where $F_i \subseteq \mathcal{P}(Q)$ for each $i = 1 \ldots K$,
- $\rightarrow$ is a finite subset of $Q \times \mathcal{P}(C) \times \mathbb{Z}^C \times \Sigma \times Q$.

Elements of $\rightarrow$ are also denoted by $q \xrightarrow{Y,up,a} q'$ where Y is interpreted as zero tests on all counters in Y. A configuration $\langle q, \boldsymbol{c} \rangle$ is an element of $Q \times \mathbb{N}^C$ and, $\langle q, \boldsymbol{c} \rangle \rightarrow \langle q', \boldsymbol{c}' \rangle$ iff there is a transition $q \xrightarrow{Y,up,a} q'$ in $\mathcal{A}$ s.t. for $c \in Y$, $\boldsymbol{c}(c) = 0$ and for $c \in C$, $\boldsymbol{c}'(c) = \boldsymbol{c}(c) + up(c)$. As usual, a run is a sequence of configurations ruled by the transitions of $\mathcal{A}$. An infinite run is accepting iff there exists a set $F \in \mathcal{F}$ such that every set $Y \in F$ is visited infinitely often. Elements of Σ^ω labeling accepting runs define the language accepted by $\mathcal{A}$. In order to accept finite words, we suppose that $\mathcal{F}$ defines a single set of final states and a finite run is accepting iff it ends at a final state with all the counters equal to zero.

However, we require additional conditions on the control graph of $\mathcal{A}$ to be declared as simple. We require that there is a partition $\{Q_0, \ldots, Q_K\}$ of Q and corresponding sets of counters $C_0, C_1, \ldots, C_K$ with $C_0 = \emptyset$ such that $I \subseteq Q_0$ and for $i \in \{1, \ldots, K\}$, a transition from a location in Q_0 to a location in Q_i can be fired only if the counters of C_i are equal to zero and all the transitions from a location in Q_i go to another location of Q_i. Moreover every transition from a location of Q_i does not modify the value of the counters in C_i. As a consequence, when we enter in the component made of the locations of Q_i the counters in C_i are equal to zero forever. Finally, for $i \in \{0, \ldots, K\}$, $F_i \subseteq \mathcal{P}(Q_i)$. Let us summarize the conditions:

1. $Q = Q_0 \uplus \cdots \uplus Q_K$ and $I \subseteq Q_0$,
2. $\mathcal{F} = \{F_0, F_1, \ldots, F_K\}$ where each $F_i \subseteq \mathcal{P}(Q_i)$,
3. there exist $K+1$ sets of counters $C_0, \ldots, C_K \subseteq C$ with $C_0 = \emptyset$ such that the transition relation $\rightarrow \subseteq Q \times \mathcal{P}(C) \times \mathbb{Z}^C \times \Sigma \times Q$ verifies the conditions below: for all $i, i' \in \{0, \ldots, K\}$, $q \in Q_i$ and $q' \in Q_{i'}$, the transitions from q to q' are of the form $q \xrightarrow{Y,up,a} q'$ where
   (a) $i \neq i'$ implies $i = 0$ and $Y = C_{i'}$,
   (b) $i = i'$ implies $Y = \emptyset$,
   (c) for $c \in C_{i'}$, $up(c) = 0$.

In the sequel we consider simple counter automata with $C = \mathcal{P}^+(\{x_1, \ldots, x_k\})$, $K = 2^k - 1$ and each set F_i contains sets of states reached by decrementing the counters in C_i. Lemma 4 below states that simplicity implies decidability of the nonemptiness problem thanks to [Jan90].

Lemma 4. *The nonemptiness problem for simple counter automata is decidable.*

Proof. We focus on the infinite case only. We reduce this problem to the problem $\mathbb{P}_{\text{temp}}$ shown decidable in [Jan90]. Let us briefly recall a fragment of the problem $\mathbb{P}_{\text{temp}}$ that consists in checking fairness conditions in Petri nets. Let $N = \langle S, T, W, M_0 \rangle$ be a Petri net [Pet81] where S is a set of places, T is a set of transitions, $W : (S \times T) \cup (T \times S) \to \mathbb{N}$ is a weight function and M_0 is an initial marking. The fragment of the language $L(Q', \mathsf{GF})$ [Jan90] we consider here is the following:

$$\psi ::= s = i \ \mid \ \psi \vee \psi \ \mid \ \psi \wedge \psi \ \mid \ \mathsf{GF}\psi,$$

where $s \in S$ and $i \in \mathbb{N}$. A formula "$s = i$" states that the number of tokens in the place s is i. As expected, "$\mathsf{GF}\psi$" states that infinitely often ψ holds true. In full generality, the problem $\mathbb{P}_{\text{temp}}$ takes as input a formula ϕ in $L(Q', \mathsf{GF})$ and an initial marking M_0 for a Petri net N and checks whether there is a an infinite execution from M_0 that satisfies ψ. This problem is decidable by [Jan90].

Consider a simple counter automaton $\mathcal{A}$ (we use the previous notations). We can build a Petri net $N_{\mathcal{A}}$ that simulates $\mathcal{A}$ apart from the zero tests. As usual, for every location q in $\mathcal{A}$, we introduce a place s_q in $N_{\mathcal{A}}$ and for every counter $c \in C$, we introduce a place s_c. Similarly, for every transition in $\mathcal{A}$, say $q \xrightarrow{Y,up,a} q'$, we consider a transition in $N_{\mathcal{A}}$ that consumes a token in s_q, produces a token in $s_{q'}$ and produces [resp. consumes] $up(c)$ counters in the place s_c when $up(c) \geq 0$ [resp. when $up(c) < 0$]. The zero tests Y will be taken into account separately in the $L(Q', \mathsf{GF})$ formula below. An initial marking contains one token in some place s_{q_0} for some initial location q_0 and no token in places of the form s_c. From this marking, we obtain markings where a unique token belongs to a place of the form s_q.

We claim that checking the nonemptiness of $\mathcal{A}$ is equivalent to verify whether the formula below holds true in $N_{\mathcal{A}}$

$$(\bigstar) \quad \bigvee_{0 \leq i \leq K} \big(\mathsf{GF}(\bigwedge_{c \in C_i} s_c = 0) \wedge (\bigwedge_{Y \in F_i} \mathsf{GF}(\bigvee_{q \in Y} s_q = 1) \big).$$

Indeed, if this property holds true there is $i \in \{0, \ldots, K\}$ such that all the counters of C_i are equal to zero infinitely often and each set of places corresponding to the set of locations of F_i is visited infinitely often. Since $N_{\mathcal{A}}$ is obtained by translation of a simple counter automaton, whenever we enter a subcomponent Q_i with $i > 0$, the counters of C_i are not modified anymore. For $j \geq 0$, if $s_c = j$ for a marking of $N_{\mathcal{A}}$ where $s_q = 1$ also holds for some location of Q_i, then $s_c = j$ always holds in the future. So the conjunction implies that the counters of C_i are equal to zero before entering the subcomponent Q_i and they remain equal to zero afterwards. The second part expresses the acceptance condition. Then it is obvious that $(\bigstar)$ holds iff $\mathcal{A}$ has an infinite accepting run.

Finite case can be solved by using the decidable reachability problem for Petri nets [Kos82]. □

5.2 Automata recognizing satisfiable counting sequences

Now, we can build a simple counter automaton $\mathcal{A}_k^l$ recognizing the set of satisfiable symbolic models of X-length l. We describe the construction for the infinite case and the automaton that recognizes finite satisfiable symbolic models can be defined similarly.

The simple counter automaton $\mathcal{A}_k^l$ is defined to be the intersection of the automata $\mathcal{A}^1$ and $\mathcal{A}^2$ which check conditions (C1) and (C2) respectively. Automaton $\mathcal{A}^1$ is a Büchi automaton and is easy to define. We focus on defining the counter automaton $\mathcal{A}^2$. We define $\mathcal{A}^2 = \langle \Sigma, C, Q, F, \{s\}, \rightarrow \rangle$, where $\Sigma = \mathtt{Frame}_k^l$, $C = \mathcal{P}^+(\{x_1, \ldots, x_k\})$, $Q = \{s\} \cup \mathtt{Frame}_k^l \cup \bigcup_{Z \subseteq C} Q_{\mathcal{A}_Z}$, where $Q_{\mathcal{A}_Z}$ is the set of states of the automaton $\mathcal{A}_Z$ which we define below, $\rightarrow$ is given by

$$s \xrightarrow{\emptyset, up_0, fr} fr$$
$$fr \xrightarrow{\emptyset, up, fr'} fr'$$
$$fr \xrightarrow{Z, up, fr'} s_{\mathcal{A}_Z}$$

where up_0 is the zero update (i.e. $up_0(X) = 0$ for each $X \in C$), $u^+_{fr}(X) \leq up(X) \leq u^+_{fr}(X) - u^-_{fr'}(X)$ for each $X \in C$, and $s_{\mathcal{A}_Z}$ is the start state of automaton $\mathcal{A}_Z$. Moreover, we require in the last rule that for every counter $X \in Z$ (i.e. every counter that is tested to zero) we have $up(X) = 0$.

The Büchi automaton $\mathcal{A}_Z$ is given by:

$$\mathcal{A}_Z = \mathcal{A}^{2a}_{C \setminus Z} \cap \bigcup_{X \in Z} (\mathcal{A}^{2b}_X \cup \mathcal{A}^{2c}_X)$$
$$\mathcal{A}^{2c}_X = \bigcap_{x \in X} \mathcal{A}_{X,x}$$

where the automata are defined as follows.

- The automaton $\mathcal{A}^{2a}_{C \setminus Z}$ accepts symbolic models in which there are no points of increment for any X in $C \setminus Z$.
- The automaton $\mathcal{A}_{X,x}$ checks condition C2(c) for X and a variable $x \in X$. The automaton $\mathcal{A}_{X,x}$ is the complement of the Büchi automaton $\mathcal{B}_{X,x}$ which accepts symbolic models in which there is a point after which we never see an x-node reachable by a forward path from a point of decrement for X. The automaton $\mathcal{B}_{X,x}$ has states of the form (fr, S) where fr is a frame and S is a subset of nodes in fr. We have a transition from (fr, S) to (fr', S') iff S' is the set of nodes in fr' which are either a point of decrement for X in fr' or are connected by a forward edge to a node in S in fr'. The automaton nondeterministically moves to a second copy where it allows the above transitions only if S' does not contain a node of the form (x, i). All states in the second copy are final.
- The automaton $\mathcal{A}^{2b}_X$ checks that infinitely often there is a point of decrement for X that is connected by a forward path to a point of increment of the form $[u]_\rho$ with $\Diamond_\rho(u) \subset X$. Its construction is similar to that of $\mathcal{A}_{X,x}$.

Observe that the automaton $\mathcal{A}_Z$ does not any essential use of the counters, namely they remain unchanged.

We can easily check that all the properties of simple counter automata are verified by this construction. For the finite case, $\mathcal{A}_k^l$ is similar to $\mathcal{A}^2$ above, except that a word is accepted when the run ends with all the counters equal to zero.

Lemma 5. *Let ρ be a symbolic model of X-length l. Then ρ is accepted by $\mathcal{A}_k^l$ iff ρ is satisfiable.*

We are now in position to state the main result of the paper.

Theorem 3. *Finitary and infinitary satisfiability for* CLTL$^{\Diamond}$ *is decidable.*

Proof. Let ϕ be a CLTL$^{\Diamond}$ formula over k variables with X-length l. Let $\mathcal{A}_\phi$ be the simple counter automaton built as the intersection of $\mathcal{A}_{\text{symb}}$, $\mathcal{A}_k^l$ and $\mathcal{A}_{lc}$ that accepts sequences in which consecutive frames are one-step consistent. Synchronization between $\mathcal{A}_{\text{symb}} \cap \mathcal{A}_{lc}$ and $\mathcal{A}_k^l$ is done as follows: for all $q_{\text{symb}}, q'_{\text{symb}} \in \mathcal{A}_{\text{symb}} \cap \mathcal{A}_{lc}$ and $q, q' \in \mathcal{A}_k^l$, $\langle q_{\text{symb}}, q\rangle \xrightarrow{X,up,fr} \langle q'_{\text{symb}}, q'\rangle$ iff $q_{\text{symb}} \xrightarrow{fr} q'_{\text{symb}}$ is a transition of $\mathcal{A}_{\text{symb}} \cap \mathcal{A}_{lc}$ and $q \xrightarrow{X,up,fr} q'$ is a transition of $\mathcal{A}_k^l$. Since $\mathcal{A}_{\text{symb}}$ have no counters, the automaton $\mathcal{A}_\phi$ is a simple counter automaton once we combine the acceptance conditions. We can check whether the language accepted by $\mathcal{A}_\phi$ is non-empty (Lemma 4). Thus, by Theorem 2, checking the satisfiability of ϕ is decidable.

According to the acceptance condition of $\mathcal{A}_k^l$ in the finite case, finitary satisfiability reduces to the reachability problem in Petri nets. □

Theorems 2 and 3 entail that this decidability result can be extended to any extension of LTL as soon as the temporal operators as definable in MSOL, see for instance [GK03].

Corollary 1. *Finitary and infinitary satisfiability for* CLTL$^{\Diamond}$ *augmented with MSOL definable temporal operators is decidable.*

5.3 A PSPACE fragment of CLTL$^{\Diamond}$

In this section, we consider the fragment CLTL$_1^{\Diamond}$ with a unique variable x. The models are sequences of natural numbers and the only counter in counting sequences α is $\{x\}$ (we identify $\alpha(i)(\{x\})$ with $\alpha(i)$). Given a symbolic model ρ over the alphabet $\mathtt{Frame}_1^l$ and the counting sequence α along ρ, for every $0 \leq i < |\rho|$, $\alpha(i+1) = \alpha(i) \dot{+} (u_i^+ - u_{i+1}^-)$ with $u_i^+, u_{i+1}^- \in \{0,1\}$. By Lemma 3, when ρ is satisfiable, in the counting sequence along ρ either the unique counter remains equal to zero after a finite number of steps or it is decremented infinitely often. Moreover, the value of the unique counter in the counting sequence is nicely bounded unlike in general with strictly more than one variable.

Lemma 6. *Let ρ be a symbolic model and α be the counting sequence along ρ. For $0 \leq i < |\rho|$, $\alpha(i) \leq l$.*

Proof. Let ρ be a symbolic model and α be the counting sequence along ρ. We show that for $0 \leq i < |\rho| - 1$, the number of equivalence classes $\rho(i)$ is bounded by $l + 1 - \alpha(i)$. We proceed by induction on i.

The number of equivalence classes in $\rho(0)$, denoted by $\sharp(\rho(0))$, is bounded by $l+1$ since there are $l+1$ terms in a frame. Since $\alpha(0) = 0$, $\sharp(\rho(0)) \leq l+1-\alpha(0)$ holds true. Now we suppose that $\sharp(\rho(i)) \leq l + 1 - \alpha(i)$ and we consider the different cases.

Case $\alpha(i+1) = \alpha(i) + 1$.
Necessarily, $u_i^+ = 1$ and $u_i^- = 0$. Hence, there is no constraint of the form $x = \mathsf{X}^j x$ with $j > 0$ in $\rho(i)$ and $\{x\}$ is an equivalence class of $\rho(i)$. Since $u_i^- = 0$, there exists a constraint in $\rho(i+1)$ of the form $\mathsf{X}^j x = \mathsf{X}^l x$ with $j < l$. The term $\mathsf{X}^l x$ in $\rho(i+1)$ does not add any new equivalence class compared to $\rho(i)$. Remember that ρ is one-step consistent and, the term $\mathsf{X}^{j+1} x$ in $\rho(i+1)$ and the term $\mathsf{X}^{j+1} x$ in $\rho(i)$ refer to the same value. As a consequence we have $\sharp(\rho(i)) = \sharp(\rho(i+1)) + 1$ and $\sharp(\rho(i+1)) \leq l + 1 - \alpha(i+1)$.

Case $\alpha(i+1) = \alpha(i) - 1$.
Necessarily, $u_i^+ = 0$ and $u_i^- = 1$. Following a reasoning similar to the previous cases, we obtain $\sharp(\rho(i)) = \sharp(\rho(i+1)) - 1$ and the inequality holds at position $i+1$.

Case $\alpha(i+1) = \alpha(i)$ and $u_i^+ = u_i^- = 1$.
We can show that $\sharp(\rho(i)) = \sharp(\rho(i+1))$ which implies that the inequality holds at position $i+1$.

Case $\alpha(i+1) = \alpha(i)$ and $u_i^+ = 0$.
Two cases need to be distinguished depending whether $u_i^- = 0$ or $u_i^- = 1$. If $u_i^- = 0$ (conditions for decrementation are not meet), then we obtain as in the previous case that $\sharp(\rho(i)) = \sharp(\rho(i+1))$ and we can conclude. Otherwise, $u_i^- = 1$ and this entails that $\alpha(i+1) = \alpha(i) = 0$. Indeed if $u_i^+ - u_i^- = -1$ and $\alpha(i+1) = \alpha(i)$, this means that the counter cannot be decremented because it is equal to zero. In this case, $\sharp(\rho(i)) = \sharp(\rho(i+1)) - 1$ and $\sharp(\rho(i+1)) \leq l+1-\alpha(i+1)$ because the number of equivalence classes is always bounded by $l+1$.

Since the number of equivalence classes is strictly positive, we have $l + 1 - \alpha(i) > 0$ and therefore $\alpha(i) \leq l$ for every position i of α. □

Boundedness entails the possibility to use automata without counters.

Lemma 7. *The set of satisfiable symbolic models over the alphabet* $\mathtt{Frame}_1^l$ *can be recognized by a standard Büchi automaton* $\mathcal{A}_1^l$ *for the infinite case, or by a finite-state automaton for the finite case.*

Proof. We know that ρ is satisfiable iff the counting sequence α over ρ satisfies the conditions (C1) and (C2). Again, we treat below only the infinite case. Since the only counter is $\{x\}$, the conjunction of (C1) and (C2) is equivalent to: either there is a position after which the value of the counter $\{x\}$ is always zero and

after which we never see a point of increment for $\{x\}$ or we infinitely often see a decrement point for the counter $\{x\}$. By Lemma 6, the value of the counter $\{x\}$ in α is bounded by l.

The Büchi automaton defined below is obtained by simplifying the previous construction with a bounded counter.

- The set of states Q is equal to $\mathtt{Frame}_1^l \times \{0, \ldots, l\} \times \{\mathtt{dec}, \neg\mathtt{dec}, \mathtt{zero}\}$. The second component encodes the value of the counter and the third one encodes whether the counter is equal to zero from now on ($\mathtt{zero}$), or a transition decrementing the counter has been just fired ($\mathtt{dec}$) or the last transition has not decremented the counter and the first condition does not hold ($\neg\mathtt{dec}$).
- For all $fr \in \mathtt{Frame}_1^l$, $i \in \{0, \ldots, l\}$ and $\theta \in \{\mathtt{dec}, \neg\mathtt{dec}\}$ we have $\langle fr, i, \theta\rangle \xrightarrow{fr} \langle fr', i', \theta'\rangle$ iff $\langle fr, fr'\rangle$ is one-step consistent and
  - if $u^+_{fr} = 1$ and $u^-_{fr'} = 0$, then $i' = i + 1$ and $\theta' = \neg\mathtt{dec}$,
  - if $u^+_{fr} = 1$ and $u^-_{fr'} = 1$, then $i' = i$ and $\theta' = \mathtt{dec}$,
  - if $u^+_{fr} = 0$ and $u^-_{fr'} = 1$, then either $i' = i - 1 = 0$ and $\theta' \in \{\mathtt{dec}, \mathtt{zero}\}$ or $i' = i - 1 > 0$ and $\theta' = \mathtt{dec}$,
  - if $u^+_{fr} = u^-_{fr'} = 0$ then either $i' = i = 0$ and $\theta' \in \{\neg\mathtt{dec}, \mathtt{zero}\}$ or $i' = i > 0$ and $\theta' = \neg\mathtt{dec}$.
- for every $fr \in \mathtt{Frame}_k^l$, we have $\langle fr, 0, \mathtt{zero}\rangle \xrightarrow{fr} \langle fr', 0, \mathtt{zero}\rangle$ iff $\langle fr, fr'\rangle$ is one-step consistent and $u^+_{fr} = u^-_{fr'}$.
- The set of final states is $\{\langle fr, i, \theta\rangle \mid \theta \in \{\mathtt{dec}, \mathtt{zero}\}\}$.

The transition relation simulates the behaviour of the counter in the counting sequence α along ρ. By construction, when a location of the form $\langle fr, 0, \mathtt{zero}\rangle$ is reached, all the future locations are of the form $\langle fr', 0, \mathtt{zero}\rangle$. Otherwise, the acceptance condition ensures that the run has to visit infinitely often states where the counter $\{x\}$ is decremented. □

The automaton $\mathcal{A}_1^l$ has an exponential size and can be built in polynomial space in l. Checking nonemptiness for this automaton can be done in non deterministic logarithmic space which allows to establish Theorem 4 below.

Theorem 4. *Finitary and infinitary satisfiability for* $\mathrm{CLTL}_1^{\Diamond}$ *is* PSPACE*-complete.*

Proof. The finitary case is analogous to the infinitary case and we treat below the latter one. Satisfiability for $\mathrm{CLTL}(\mathbb{N}, =)$ is PSPACE-hard (propositional variables can be encoded by constraints of the form $x = y$) and $\mathrm{CLTL}(\mathbb{N}, =)$ restricted to a single variable is also PSPACE-hard by using [DG07, Lemma 2]. Since $\mathrm{CLTL}(\mathbb{N}, =)$ restricted to a single variable is a fragment of $\mathrm{CLTL}_1^{\Diamond}$, we obtain the PSPACE lower bound.

In order to show the PSPACE upper bound, we recall the standard analysis. Given a $\mathrm{CLTL}_1^{\Diamond}$ formula ϕ, the automaton $\mathcal{A}_\phi$ is defined as the intersection $\mathcal{A}_{\mathrm{symb}}$ and $\mathcal{A}_1^l$. The construction of $\mathcal{A}_{\mathrm{symb}}$ described in Section 3.2 and the construction of $\mathcal{A}_1^l$ in the proof of Lemma 7 can both be performed in polynomial space in the size of ϕ even if their size is exponential. Synchronizing these two Büchi automata w.r.t. the alphabet of frames can be done in PSPACE, and the resulting

automaton $\mathcal{A}_\phi$ is a Büchi automaton. We recall that $\mathcal{A}_\phi$ is satisfiable iff the language recognized by $\mathcal{A}_\phi$ is nonempty (see Theorem 2). Since checking the nonemptiness of a Büchi automaton can be done in NLOGSPACE, we get a whole procedure in nondeterministic polynomial space which is equivalent to PSPACE by Savitch's theorem. □

The models for $\mathrm{CLTL}^{\Diamond}_1$ corresponds to models of $\mathrm{LTL}^{\downarrow}(\sim, \mathsf{X}, \mathsf{U})$. Therefore, finitary and infinitary satisfiability for $\mathrm{LTL}^{\downarrow}_1(\sim, \mathsf{X}, \mathsf{X}^{-1}, \mathsf{U}, \mathsf{S})$ restricted to formulae such that the freeze operator is restricted to subformulae of the form $\downarrow_{r=x} \mathsf{XF}(r = x)$ and $\downarrow_{r=x} \mathsf{X}^i(r = x)$ is decidable in polynomial space (r is the unique register and x the unique variable).

5.4 Repeating values in the past is still decidable

In this section we explain why we can allow the constraints of the language to state properties about past repetitions of a value without loosing decidability. Let $\mathrm{CLTL}^{\Diamond,\Diamond^{-1}}$ be the extension of $\mathrm{CLTL}^{\Diamond}$ with atomic formulae of the form $x = \Diamond^{-1} y$. The satisfaction relation is extended as follows:

$\sigma, i \models x = \Diamond^{-1} y$ iff there exists $j > 0$ s.t. $x = \sigma(i - j)(y)$ and $0 \leq i - j$.

Similarly to what is done in Section 2.1, x *diff* $\Diamond^{-1} y$ can be omitted since it can be defined from $x = \Diamond^{-1} y$ (a variant of the equivalence (1)).

In order to deal with satisfiability for $\mathrm{CLTL}^{\Diamond,\Diamond^{-1}}$ we need to extend the symbolic representation of models. In addition of the conditions(F1)–(F5) defined in Sect. 3.1, a frame fr has to verify the following property (the finitary case admits a similar update):

(F6) for all $i, j \in \{0, \ldots, l\}$ and $x, y \in \{x_1, \ldots x_k\}$, if $\mathsf{X}^i x = \mathsf{X}^j y$ is in fr then
- if $i = j$, then for every $z \in \{x_1, \ldots, x_k\}$ we have $\mathsf{X}^i(x = \Diamond^{-1} z) \in fr$ iff $\mathsf{X}^j(y = \Diamond^{-1} z) \in fr$ (we extend the notion of frames);
- if $i > j$ then $\mathsf{X}^i(x = \Diamond^{-1} y) \in fr$, and for every $z \in \{x_1, \ldots, x_k\}$, $\mathsf{X}^i(x = \Diamond^{-1} z) \in fr$ iff either $\mathsf{X}^j(y = \Diamond^{-1} z) \in fr$ or there exists $i > j' \geq j$ such that $\mathsf{X}^i x = \mathsf{X}^{j'} z$ is in fr.

We pose $\Diamond^{-}_{fr}(\mathsf{X}^i x) \stackrel{\text{def}}{=} \{y \mid \mathsf{X}^i(x = \Diamond^{-} y) \in fr\}$. Since we need to deal with past obligations, a counter is a pair $\langle X_p, X_f \rangle$ in $\mathcal{P}^+(\{x_1, \ldots, x_k\}) \times \mathcal{P}^+(\{x_1, \ldots, x_k\})$ where X_p is for past obligations and X_f for future obligations. We update the notion of counter valuations accordingly. A value n for $\langle X_p, X_f \rangle$ is the number of values that occurred in a past state of every variable of X_p and that have to be repeated in a future state of every variable in X_f.

We extend some earlier definitions. For an l-frame fr and counter $\langle X_p, X_f \rangle$, we define a *point of increment* for $\langle X_p, X_f \rangle$ in fr to be an equivalence class of the form $[(x, 0)]_{fr}$ such that $\Diamond_{fr}(x, 0) = X_f$, $(x, 0)$ is not connected by a forward edge to a node in fr *and* $[(x, 0)]_{fr} \cup \Diamond^{-1}_{fr}(x, 0) = X_p$. A *point of decrement* for $\langle X_p, X_f \rangle$ in fr is defined to be an equivalence class of the form $[(x, l)]_{fr}$ such that $\Diamond_{fr}(x, l) \cup [(x, l)]_{fr} = X_f$, (x, l) is not connected by a backward edge to another node in fr *and* $\Diamond^{-1}(x, l) = X_p$. Let u^{+}_{fr} denote a counter valuation which records the number

of points of increment for each counter $\langle X_p, X_f\rangle$, in fr. Similarly let u^-_{fr} denote the counter valuation which records the number of points of decrement for each counter $\langle X_p, X_f\rangle$ in fr. We can now define a canonical counter valuation sequence α along ρ, called the *counting sequence* along ρ, which counts the number of "unsatisfied" points of increments for each counter $\langle X_p, X_f\rangle$ with $X_p \neq \emptyset$. We define α inductively: $\alpha(0)(\langle X_p, X_f\rangle) = 0$ and for $0 \leq i < |\rho|$, $\alpha(i+1)(\langle X_p, X_f\rangle) = \alpha(i)(\langle X_p, X_f\rangle) + (u^+_{\rho(i)}(\langle X_p, X_f\rangle) - u^-_{\rho(i+1)}(\langle X_p, X_f\rangle))$. Note that decrementations are this time compulsory and we allow $\alpha(i)(\langle X_p, X_f\rangle)$ to be negative (but not in the acceptance condition). Though we need more counters, dealing with past repeating values, does not introduce real complications. This is analogous to the passage from LTL to LTL with past-time operators since past is finite and information about past can be accumulated smoothly.

Lemma 8. *A symbolic model ρ for the logic* $\mathrm{CLTL}^{\Diamond,\Diamond^{-1}}$ *is satisfiable iff the counting sequence along ρ satisfies the conditions from Lemma 3 for the future part of the counters and for every $0 \leq i < |\rho|$, $\alpha(i)(\langle X_p, X_f\rangle) \geq 0$.*

The proof is similar to the proof of Lemma 3. In order to define an augmented graph, we use counters of the form $\langle X_p, X_f\rangle$ instead of X. Moreover, the condition $\alpha(i)(\langle X_p, X_f\rangle) \geq 0$ guarantees that in the procedure to add augmented edges, for each point of decrement for $\langle X_p, X_f\rangle$ with $X_p \neq \emptyset$, there is an unmatched point of increment for $\langle X_p, X_f\rangle$. Indeed, in the proof of Lemma 3, it was possible to find points of decrement that do not create augmented edges.

As a consequence, we can easily update the construction of $\mathcal{A}_\phi$ in order to deal with past repeating values. The definition of the automata $\mathcal{A}_{\text{symb}}$ and $\mathcal{A}^l_k$ are just extended by considering the new definition for frames. By way of example, in the definition of $\mathcal{A}^2$, the update function *up* for each transition satisfies: for every counter $\langle X_p, X_f\rangle$ with $X_p \neq \emptyset$, $up(\langle X_p, X_f\rangle) = u^+_{fr}(\langle X_p, X_f\rangle) - u^-_{fr'}(\langle X_p, X_f\rangle)$. It is also worth observing that the automaton $\mathcal{A}_\phi$ obtained by synchronization of these automata still belongs to the class of simple counter automata and the decidability result also holds for $\mathrm{CLTL}^{\Diamond,\Diamond^{-1}}$ satisfiability problem.

Theorem 5. (I) *Finitary and infinitary satisfiability for* $\mathrm{CLTL}^{\Diamond,\Diamond^{-1}}$ *is decidable.* **(II)** *Finitary and infinitary satisfiability for* $\mathrm{CLTL}_1^{\Diamond,\Diamond^{-1}}$ *is* PSPACE*-complete.*

Proof. (I) Proof similar to the proof of Theorem 3.
(II) In presence of a unique variable x, the unique counter of the form $\langle X_p, X_f\rangle$ with $X_p \neq \emptyset$ is $\langle\{x\}, \{x\}\rangle$. Lemma 6 still holds true for this extension and the proof disregards the case $\alpha(i) = \alpha(i+1)$ with $u^+_i = 0$ and $u^-_i = 1$. Consequently, Lemma 7 and Theorem 4 can be adapted accordingly.

6 Concluding Remarks

We have shown that satisfiability for $\mathrm{CLTL}^{\Diamond}$ with operators in $\{\mathsf{X}, \mathsf{X}^{-1}, \mathsf{S}, \mathsf{U}\}$ is decidable by reduction into the verification of fairness properties in Petri nets [Jan90]. The proof is uniform for the finitary and infinitary cases and it

can be extended to atomic constraints of the form $x = \Diamond^{-1} y$ and to any set of MSOL definable temporal operators. Moreover, satisfiability for CLTL$^{\Diamond}$ restricted to one variable is shown PSPACE-complete. Hence, we have defined and studied a well-designed decidable fragment of LTL with the freeze quantifier answering some question from [WZ00] and circumventing some undecidability results from [DL06]. Finally, as done also in [DL06,Laz06,BMS$^+$06], we show relationships between fragments of LTL with freeze and counter automata. Our connection is all the more interesting because we deal with the finitary and infinitary cases while preserving decidability.

The main question left open by our work is the complexity of satisfiability for CLTL$^{\Diamond}$ and more precisely we do not know whether CLTL$^{\Diamond}$ satisfiability has elementary complexity. Similarly, are there natural fragments of CLTL$^{\Diamond}$ that are of lower complexity, for instance the one involved in Theorem 1? Another promising extension consists in considering other concrete domains as $\langle \mathbb{R}, <, = \rangle$ and to allow atomic formulae of the form $x < \Diamond y$. The decidability status of such a variant is still open, even if in absence of the restricted use of the freeze quantifier, PSPACE-completeness is known [DD07]. Finally, it would be interesting to investigate branching-time extensions.

Acknowledgements: We are grateful to Petr Jančar (TU Ostrava) for pointing us to [Jan90] in order to solve the nonemptiness problem for simple counter automata and for suggesting the proof of Lemma 4 and Ranko Lazić (U. of Warwick) for remarks about a preliminary version.

References

[ABM01] C. Areces, P. Blackburn, and M. Marx. Hybrid logics: characterization, interpolation and complexity. *The Journal of Symbolic Logic*, 66(3):977–1010, 2001.

[AH94] R. Alur and Th. Henzinger. A really temporal logic. *Journal of the Association for Computing Machinery*, 41(1):181–204, 1994.

[BMS$^+$06] M. Bojańczyk, A. Muscholl, Th. Schwentick, L. Segoufin, and C. David. Two-variable logic on words with data. In *LICS'06*, pages 7–16. IEEE, 2006.

[DD07] S. Demri and D. D'Souza. An automata-theoretic approach to constraint LTL. *Information and Computation*, 205(3):380–415, 2007.

[DDG07] S. Demri, D. D'Souza, and R. Gascon. Decidable temporal logic of repeating values. In *LFCS'07*, volume 4514 of *Lecture Notes in Computer Science*, pages 180–194. Springer, 2007.

[DG07] S. Demri and R. Gascon. The effects of bounding syntactic resources on presburger ltl (extended abstract). In *TIME'07*, 2007. to appear.

[DL06] S. Demri and R. Lazić. LTL with the freeze quantifier and register automata. In *LICS*, pages 17–26. IEEE, 2006.

[DLN07] S. Demri, R. Lazić, and D. Nowak. On the freeze quantifier in constraint LTL: decidability and complexity. *Information and Computation*, 205(1):2–24, 2007.

[Esp94] J. Esparza. On the decidability of model checking for several μ-calculi and Petri nets. In *CAAP'94*, volume 787 of *Lecture Notes in Computer Science*, pages 115–129. Springer, 1994.

[Fit02] M. Fitting. Modal logic between propositional and first-order. *Journal of Logic and Computation*, 12(6):1017–1026, 2002.

[GK03] P. Gastin and D. Kuske. Satisfiability and model checking for MSO-definable temporal logics are in PSPACE. In *CONCUR'03*, volume 2761 of *Lecture Notes in Computer Science*, pages 222–236. Springer, 2003.

[Gor96] V. Goranko. Hierarchies of modal and temporal logics with references pointers. *Journal of Logic, Language, and Information*, 5:1–24, 1996.

[Hen90] Th. Henzinger. Half-order modal logic: how to prove real-time properties. In *PODC'90*, pages 281–296. ACM, 1990.

[HR89] R.R. Howell and L.E. Rosier. Problems concerning fairness and temporal logic for conflict-free Petri nets. *Theoretical Computer Science*, 64:305–329, 1989.

[Jan90] P. Jančar. Decidability of a temporal logic problem for Petri nets. *Theoretical Computer Science*, 74(1):71–93, 1990.

[Kos82] R. Kosaraju. Decidability of reachability in vector addition systems. In *STOC'82*, pages 267–281, 1982.

[KV06] O. Kupferman and M. Vardi. Memoryful Branching-Time Logic. In *LICS'06*, pages 265–274. IEEE, 2006.

[Laz06] R. Lazić. Safely freezing LTL. In *FST&TCS'06*, volume 4337 of *Lecture Notes in Computer Science*, pages 381–392. Springer, 2006.

[LMS02] F. Laroussinie, N. Markey, and Ph. Schnoebelen. Temporal logic with forgettable past. In *LICS'02*, pages 383–392. IEEE, 2002.

[LP05] A. Lisitsa and I. Potapov. Temporal logic with predicate λ-abstraction. In *TIME'05*, pages 147–155. IEEE, 2005.

[Pet81] J.L. Peterson. *Petri Net Theory and the Modelling of Systems*. Prentice-Hall, 1981.

[Rab06] A. Rabinovich. Decidability and expressive power of real time logics. In *FORMATS'06*, volume 4202 of *Lecture Notes in Computer Science*, page 32. Springer, 2006. Invited talk.

[Seg06] L. Segoufin. Automata and logics for words and trees over an infinite alphabet. In *CSL'06*, volume 4207 of *Lecture Notes in Computer Science*, pages 41–57. Springer, 2006.

[tCF05] B. ten Cate and M. Franceschet. On the complexity of hybrid logics with binders. In *CSL'05*, volume 3634 of *Lecture Notes in Computer Science*, pages 339–354. Springer, 2005.

[VW94] M. Vardi and P. Wolper. Reasoning about infinite computations. *Information and Computation*, 115:1–37, 1994.

[WZ00] F. Wolter and M. Zakharyaschev. Spatio-temporal representation and reasoning based on RCC-8. In *KR'00*, pages 3–14. Morgan Kaufmann, 2000.

A Proof of Theorem 1

(I) Finitary satisfiability for $\mathrm{CLTL}^{\Diamond}(\mathsf{X},\mathsf{U})$ can be easily reduced to finitary satisfiability for $\mathrm{CLTL}^{\downarrow}_{(\omega,1)}(\mathsf{X},\mathsf{U})$. By [DLN07, Proposition 4], this latter problem can be reduced to finitary satisfiability for $\mathrm{LTL}^{\downarrow}_{1}(\sim,\mathsf{X},\mathsf{U})$ that is decidable by [DL06, Corollary 13].

(II) Let ϕ be a $\mathrm{CLTL}^{\Diamond}(\mathsf{X},\mathsf{X}^{-1},\mathsf{F},\mathsf{F}^{-1})$ formula over the variables $\{x_1,\ldots,x_k\}$ and l be the maximal i such that a term of the form $\mathsf{X}^i x$ occurs in ϕ. This formula can expressed into an equivalent formula of $\mathrm{CLTL}^{\downarrow}_{(\omega,1)}(\mathsf{X},\mathsf{X}^{-1},\mathsf{F},\mathsf{F}^{-1})$ using the equivalence below:

(X) $x = \mathsf{X}^i y \Leftrightarrow \downarrow_{r=x} \mathsf{X}^i(r=y)$,
($\Diamond$) $x = \Diamond y \Leftrightarrow \downarrow_{r=x} \mathsf{XF}(r=y)$.

Let $N = 3k(l+1)$ and $\mathcal{O}_N$ be the set of temporal operators below:

$$\mathcal{O}_N = \{\mathsf{X},\mathsf{X}^2,\ldots\mathsf{X}^N,\mathsf{X}^{N+1}\mathsf{F},\mathsf{X}^{-1},\mathsf{X}^{-2},\ldots\mathsf{X}^{-N},\mathsf{X}^{-(N+1)}\mathsf{F}^{-1}\}.$$

Using a proof technique from [DG07], we build a formula ϕ' in the simple fragment of $\mathrm{CLTL}^{\downarrow}_{(1,1)}(\mathcal{O}_N)$ such that ϕ is satisfiable iff ϕ' is. The simple fragment of $\mathrm{CLTL}^{\downarrow}_{(1,1)}(\mathcal{O}_N)$ is defined as the fragment such that every temporal operator $\mathrm{O} \in \mathcal{O}_N$ is in the direct scope of a freeze operator $\downarrow_{r=x} \mathrm{O}\phi$ and there is no other occurrences of the freeze operator. Finitary satisfiability for the simple fragment of $\mathrm{CLTL}^{\downarrow}_{(1,1)}(\mathcal{O}_N)$ is decidable [BMS+06,DL06].

The idea is to encode one state from a k-variable model into $3k$ states in 1-variable models. Only one state over three encodes values. Intermediate states are used to know when a sequence of $3k$ states corresponds to a state in of the k-variable model. For instance, if $k=2$, then the 2-variable model below

$$\begin{pmatrix} y_1^0 \\ y_2^0 \end{pmatrix} \begin{pmatrix} y_1^1 \\ y_2^1 \end{pmatrix} \ldots$$

is encoded as the 1-variable model

$$y_1^0 = \circ \neq \circ \neq y_2^0 \neq \circ \neq \circ \neq y_1^1 = \circ \neq \circ \neq y_2^1 \neq \circ \neq \circ \ldots$$

where $\circ$ denotes arbitrary values satisfying the mentioned relations with its neighbors (each occurrence of $\circ$ corresponds to a possibly distinct value). The beginning of the encoding of some state from the 2-variable model satisfies that two consecutive values of x_1 are identical. More generally, in the 1-variable model, $x = \mathsf{X}x$ holds true when the current position starts the encoding of a position in the k-variable model and access to the value of $\mathsf{X}^i x_j$ in the k-variable model is done via the term $\mathsf{X}^{3ik+3j}x$. We can impose that $x = \mathsf{X}x$ every $3k$ states and the model is of length multiple of $3k$ by using the formula below:

$$\phi_{3k} \stackrel{\text{def}}{\equiv} (x = \mathsf{X}x) \wedge \bigwedge_{0<i<3k} \mathsf{X}^i(x \neq \mathsf{X}x) \wedge$$

$$\mathsf{G}\big(((x = \mathsf{X}x) \wedge \mathsf{X}^{3k+1}\top) \Leftrightarrow \mathsf{X}^{3k}(x = \mathsf{X}x)\big) \wedge \mathsf{G}((x = \mathsf{X}x) \Rightarrow \bigwedge_{0<i<3k} \mathsf{X}^i\top)$$

which is equivalent to

$$\downarrow_{r=x} \mathsf{X}(r = x) \wedge \bigwedge_{0<i<3k} \mathsf{X}^i \downarrow_{r\neq x} \mathsf{X}(r = x) \wedge \neg\mathsf{F}\big((\downarrow_{r=x} \mathsf{X}(r = x) \wedge \mathsf{X}^{3k+1}\top$$
$$\wedge \mathsf{X}^{3k} \downarrow_{r=x} \mathsf{X}(r \neq x)) \vee ((\downarrow_{r=x} \mathsf{X}(r \neq x) \vee \neg\mathsf{X}^{3k+1}\top) \wedge \mathsf{X}^{3k} \downarrow_{r=x} \mathsf{X}(r = x)))\big) \wedge$$
$$\neg\mathsf{F}(\downarrow_{r=x} \mathsf{X}(r = x) \wedge (\bigvee_{0<i<3k} \neg\mathsf{X}^i\top))$$

This formula can be expressed in the simple fragment using the equivalences

($\star$) $\mathsf{F}\phi \Leftrightarrow \phi \vee \mathsf{X}\phi \vee \cdots \vee \mathsf{X}^N\phi \vee \mathsf{X}^{N+1}\mathsf{F}\phi$ (and similarly with F^{-1}) ,
($\star\star$) $\mathrm{O}(\downarrow_{r=x} \phi_1 \otimes \cdots \otimes \downarrow_{r=x} \phi_m) \Leftrightarrow \downarrow_{r=x} \mathrm{O}(\downarrow_{r=x} \phi_1 \otimes \cdots \otimes \downarrow_{r=x} \phi_m)$
for all $\mathrm{O} \in \mathcal{O}_N$ and $\otimes \in \{\wedge, \vee\}$.

Formally, we define a map f over the set of CLTL$^{\Diamond}$ formulae such that:

(X) $f(x_m = \mathsf{X}^i x_n) \stackrel{\text{def}}{\equiv} f(\downarrow_{r=x_m} \mathsf{X}^i(r = x_n)) = \downarrow_{r=\mathsf{X}^{3m}x} \mathsf{X}^{3ik}(r = \mathsf{X}^{3n}x)$,
($\Diamond$) $f(x_m = \Diamond x_n) \stackrel{\text{def}}{\equiv} f(\downarrow_{r=x_m} \mathsf{XF}(r = x_n)) = \downarrow_{r=\mathsf{X}^{3m}x} \mathsf{X}^{3k}\mathsf{F}(\mathsf{X}^{-3n}(x = \mathsf{X}x) \wedge (r = x))$,

- f is homomorphic for the Boolean operators,
- $f(\mathsf{X}\phi) = \mathsf{X}^{3k} f(\phi)$,
- $f(\mathsf{F}\phi) = \mathsf{F}((x = \mathsf{X}x) \wedge f(\psi))$,
- $f(\mathsf{X}^{-1}\phi) = \mathsf{X}^{-3k} f(\phi)$,
- $f(\mathsf{F}^{-1}\phi) = \mathsf{F}^{-1}((x = \mathsf{X}x) \wedge f(\psi))$.

Finally, we can prove that any formula obtained by applying the map f is equivalent to a formula in the simple fragment of $\mathrm{CLTL}^{\downarrow}_{(1,1)}(\mathcal{O}_N)$. We develop the cases (X) and ($\Diamond$):

(X) $\downarrow_{r=\mathsf{X}^{3m}x} \mathsf{X}^{3ik}(r = \mathsf{X}^{3n}x)$
$\equiv \mathsf{X}^{3m} \downarrow_{r=x} \mathsf{X}^{3(ik-m)}(r = \mathsf{X}^{3n}x)$
$\equiv \downarrow_{r=x} \mathsf{X}^{3m} \downarrow_{r=x} \mathsf{X}^{3ik-m+n}(r = x)$.

($\Diamond$) $\downarrow_{r=\mathsf{X}^{3m}x} \mathsf{X}^{3k}\mathsf{F}(\mathsf{X}^{-3n}(x = \mathsf{X}x) \wedge (r = x))$
$\equiv \mathsf{X}^{3m} \downarrow_{r=x} \mathsf{X}^{3(k-m)}\mathsf{F}(\mathsf{X}^{-3n} \downarrow_{r=x} \mathsf{X}(r = x) \wedge (r = x))$.
$\equiv \downarrow_{r=x} \mathsf{X}^{3m} \downarrow_{r=x} \mathsf{X}^{3(k-m)}\mathsf{F}(\downarrow_{r=x} \mathsf{X}^{-3n} \downarrow_{r=x} \mathsf{X}(r = x) \wedge (r = x))$.

We can see that these formulae have one register, one free variable and all the temporal operators are directly under the scope of a freeze quantifier. For the temporal operators $\mathcal{O}_N$ we need to use ($\star$) and ($\star\star$) to find an equivalent formula in the simple fragment of $\mathrm{CLTL}^{\downarrow}_{(1,1)}(\mathcal{O}_N)$.

The formula ϕ is satisfiable iff $\phi_{3k} \wedge f(\phi)$ is satisfiable. □