

Prompt Interval Temporal Logic

Dario Della Monica¹, Angelo Montanari², Aniello Murano¹, and Pietro Sala³

¹ Università degli Studi di Napoli “Federico II”, Napoli, Italy,

dario.dellamonica@unina.it, murano@na.infn.it

² University of Udine, Udine, Italy angelo.montanari@uniud.it

³ University of Verona, Verona, Italy, pietro.sala@univr.it

Abstract. Interval temporal logics are expressive formalisms for temporal representation and reasoning, which use time intervals as primitive temporal entities. They have been extensively studied for the past two decades and successfully applied in AI and computer science. Unfortunately, they lack the ability of expressing promptness conditions, as it happens with the commonly-used temporal logics, e.g., LTL: whenever we deal with a liveness request, such as “something good eventually happens”, there is no way to impose a bound on the delay with which it is fulfilled. In the last years, such an issue has been addressed in automata theory, game theory, and temporal logic. In this paper, we approach it in the interval temporal logic setting. First, we introduce PROMPT-PNL, a prompt extension of the well-studied interval temporal logic PNL, and we prove the undecidability of its satisfiability problem; then, we show how to recover decidability (NEXPTIME-completeness) by imposing a natural syntactic restriction on it.

1 Introduction

Interval temporal logics provide a powerful framework suitable for reasoning about time. Unlike classic temporal logics, such as Linear Temporal Logic (LTL) [21] and the like, they use time intervals, instead of time points, as primitive temporal entities. Such a distinctive feature turns out to be very useful in various Computer Science and AI application domains, ranging from hardware and real-time system verification to natural language processing, from constraint satisfaction to planning [1,2,10,20,22,23]. As concrete applications, we mention TERENCE [14], an adaptive learning system for poor comprehenders and their educators (based on *Allen’s interval algebra* IA [1]), and RISMA [17], an algorithm to analyze behavior and performance of real-time data systems (based on Halpern and Shoham’s *modal logic of Allen’s relations* HS [15]).

A fundamental class of properties that can be expressed in (both interval- and point-based) temporal logics is that of *liveness* properties, which allow one to state that something “good” will eventually happen. However, a limitation that is common to most temporal logics is the lack of support for *promptness*: it is not possible to bound the delay with which a liveness request is fulfilled, despite the fact that this is desirable for many practical applications (see [16] for a convincing argument). To overcome such a shortcoming, a whole body of work has been recently devoted to the study of promptness. In [4,16], the authors extend LTL with the ability of bounding the delay with which a temporal request is satisfied. In [3], the use of prompt accepting conditions in the context of ω -regular automata is explored by introducing *prompt-Büchi automata*, whose

accepting condition imposes the existence of a bound on the number of non-accepting states in between two consecutive occurrences of accepting ones. Prompt extensions of LTL have also been investigated outside the realm of closed systems. Two-player turn-based games with perfect information have been explored in the prompt LTL setting in [24]. In [9], the authors lift the prompt semantics to ω -regular games, under the parity winning condition, by introducing *finitary parity* games. They make use of the concept of *distance* between positions in a play that refers to the number of edges traversed in the game arena; the classical parity winning condition is then reformulated to take into consideration only those states occurring with a bounded distance. Such an idea has been generalised to deal with more involved prompt parity conditions [13,19]. In the field of formal languages, promptness comes into play in [6], where ω B-regular languages and their automata counterpart, known as ω B-automata, are studied. Intuitively, ω B-regular languages extend ω -regular ones with the ability of bounding the distance between occurrences of sub-expressions in consecutive ω -iterations, within each word of the language. Finally, an extension of alternating-time epistemic temporal logic with prompt-eventuality has been recently investigated in [5].

In this paper, we show that interval temporal logics can be successfully provided with a support for prompt-liveness specifications by lifting the work done in [4,16] to the interval-based setting.

In [4], the language of LTL is enriched with *parameterized* versions of temporal modalities F (*eventually*) and U (*until*), as well as of the dual modalities G (*globally*) and R (*release*). The resulting logic, called PLTL, features the following parameterized modalities: $F_{\leq x}$, $F_{> y}$, $G_{\leq y}$, $G_{> x}$, $U_{\leq x}$, $U_{> y}$, $R_{\leq y}$, and $R_{> x}$, where $x \in X$, $y \in Y$, and X and Y are two disjoint sets of *bounding variables*. Intuitively, a formula $F_{\leq x}\phi$ is true if ϕ is satisfied within x time units, according to the valuation of x (the other parameterized modalities have an analogous interpretation). Thus, PLTL models are LTL models, i.e., words over the powerset of the set of atomic propositions, enriched with a valuation for the bounding variables in $X \cup Y$. The satisfiability problem for PLTL is PSPACE-complete, as for LTL. The assumption that X and Y are disjoint is crucial in retaining decidability. In [16], the authors introduce the logic PROMPT-LTL, which restricts PLTL in three ways: (i) a parameterized version is introduced for the modality F only (parameterized versions of modalities G , U , and R are not included); (ii) only upper bounds appear in parameterized modalities, i.e., no subscript of the form $>_x$ occurs; (iii) there is only one bounding variable. The restriction imposed by PROMPT-LTL is less strong than it looks like: as shown in [4], operator $F_{\leq x}$, along with the classic LTL constructs, is enough to define operators $G_{> x}$, $U_{\leq x}$, $R_{> x}$ (i.e., all the operators involving in their subscript variables in X). As PROMPT-LTL enriches LTL with the ability of limiting the amount of time a fulfillment of an existential request (corresponding to a liveness property) can be delayed, it can be thought of as an extension of LTL with *prompt liveness*. In [16], it is shown that reasoning about PROMPT-LTL is not harder than reasoning about LTL, with respect to a series of basic problems, including satisfiability (PSPACE-complete).

In the present paper, we show how to extend the logic PNL of temporal neighborhood (a well-known fragment of HS whose satisfiability problem is NEXPTIME-complete [8]), with the ability of expressing prompt-liveness properties. Following

the approach of [16], we introduce ‘prompt’ versions (i.e., upper bounds only) of all modalities of PNL. The resulting modality templates are as follows: the *prompt-right-adjacency* $\langle A_x \rangle$ and the *prompt-left-adjacency* $\langle \bar{A}_x \rangle$, capturing prompt-liveness in the future and in the past, respectively, as well as the dual modalities $[A_x]$ and $[\bar{A}_x]$. Intuitively, a modality $\langle A_x \rangle$ (for some upper bound x) forces the existence of an event starting exactly when the current one terminates and ending within an amount of time bounded above by the value of x . Similarly, $\langle \bar{A}_x \rangle$ forces the existence of an event ending exactly when the current one begins and starting at most x time units before the beginning of the current one. Modalities $[A_x]$ and $[\bar{A}_x]$ express dual properties in the standard way, namely, $[A_x]\psi$ stands for $\neg\langle A_x \rangle\neg\psi$ and $[\bar{A}_x]\psi$ stands for $\neg\langle \bar{A}_x \rangle\neg\psi$. We name the proposed logic PROMPT-PNL (Section 2).

We first prove that the future fragment of PROMPT-PNL (PROMPT-RPNL), involving the future modalities $\langle A \rangle$, $[A]$, $\langle A_x \rangle$, and $[A_x]$ only, is expressive enough to encode the *finite colouring problem*, known to be undecidable [18]. Undecidability of PROMPT-RPNL (and PROMPT-PNL) immediately follows (Section 3). Notably, unlike LTL, PNL is strictly more expressive than its future fragment RPNL (see [12]); such a separation result holds between PROMPT-PNL and PROMPT-RPNL as well. Our undecidability result hinges on the unrestricted use of bounding variables within prompt modalities, which allows one to somehow establish tight bounds for the length of intervals. We show that decidability can be recovered by using two disjoint sets of bounding variables, one for existential modalities and the other for universal ones. Formulas of the resulting logic, which we name PROMPT^dPNL, enjoy some useful monotonicity property, i.e., the truth of a formula $\langle A_x \rangle\psi$ under a certain interpretation $\sigma(x)$ of the bounding variable x implies its truth under every interpretation σ' , with $\sigma'(x) \geq \sigma(x)$. This allows us to prove a small (pseudo-)model property for PROMPT^dPNL, from which we conclude that the satisfiability problem for PROMPT^dPNL is NEXPTIME-complete (Section 4). Due to lack of space, most of the proofs are omitted (see [11] for full proofs).

2 The logic PROMPT-PNL

Let us start with some basic notions of interval-based temporal logics. A *linear order* \mathbb{D} is a pair $\langle D, < \rangle$, where D is a set, called *domain*, whose elements are referred to as *points*, and $<$ is a strict total order over D . A (strongly) *discrete* linear order is a linear order such that there are only finitely many points in between any two points. In the rest of the paper, we tacitly assume every domain to be discrete. For the sake of simplicity, we identify the domain of a linear order with the linear order itself, e.g. we write “ $d \in \mathbb{D}$ ” instead of “ $d \in D$ ”. Let $d \in \mathbb{D}$. The *successors* (resp., *predecessors*) of d in \mathbb{D} are the points $d' \in \mathbb{D}$ such that $d < d'$ (resp., $d' < d$); the *immediate successor* (resp., *immediate predecessor*) of d in \mathbb{D} , denoted by $\text{succ}_{\mathbb{D}}(d)$ (resp., $\text{pred}_{\mathbb{D}}(d)$), is (if any) the point $d' \in \mathbb{D}$ such that d' is a successor (resp., predecessor) of d in \mathbb{D} and no point $d'' \in \mathbb{D}$ exists with $d < d'' < d'$ (resp., $d' < d'' < d$). Note that $\text{succ}_{\mathbb{D}}(d)$ (resp., $\text{pred}_{\mathbb{D}}(d)$) is defined unless d is the greatest (resp., least) element in \mathbb{D} . Given a linear order \mathbb{D} and two points $a, b \in \mathbb{D}$, with $a < b$, we denote by $[a, b]$ an *interval* (over \mathbb{D}). The set of intervals over a linear order \mathbb{D} is denoted by $\mathbb{I}(\mathbb{D})$. An *interval structure* (over

a countable set \mathcal{AP} of *atomic propositions*) is a pair $\langle \mathbb{D}, V \rangle$, where \mathbb{D} is a linear order and $V : \mathbb{I}(\mathbb{D}) \rightarrow 2^{\mathcal{AP}}$ is a *valuation function*, which assigns to each interval over \mathbb{D} the set of atomic proposition that are true over it. Given a linear order \mathbb{D} and $a, b \in \mathbb{D}$, we denote by $\mathbb{D}^{\geq a}$ (resp., $\mathbb{D}^{>a}$, $\mathbb{D}^{\leq a}$, $\mathbb{D}^{<a}$, $\mathbb{D}^{[a,b]}$, $\mathbb{D}^{]a,b[}$, $\mathbb{D}^{[a,b[}$, $\mathbb{D}^{]a,b]}$) the set of elements $d \in \mathbb{D}$ such that $d \geq a$ (resp., $d > a$, $d \leq a$, $d < a$, $a \leq d \leq b$, $a < d < b$, $a \leq d < b$, $a < d \leq b$). For instance, we denote by $\mathbb{R}^{>0}$ the set of positive reals.

Syntax and semantics. Let \mathcal{AP} (*atomic propositions*) and X (*bounding variables*) be two countable sets. Formulas of PROMPT-PNL in negation normal form are defined as follows:

$$\begin{array}{l} \varphi ::= p \mid \varphi \wedge \varphi \mid \langle A \rangle \varphi \mid \langle \bar{A} \rangle \varphi \mid \langle A_x \rangle \varphi \mid \langle \bar{A}_x \rangle \varphi \\ \quad \mid \neg p \mid \varphi \vee \varphi \mid [A] \varphi \mid [\bar{A}] \varphi \mid [A_x] \varphi \mid [\bar{A}_x] \varphi \end{array}$$

where $p \in \mathcal{AP}$ and $x \in X$. We also use other standard Boolean connectives, e.g., \rightarrow , and logical constants \top and \perp , which are defined in the usual way. We denote by PROMPT-RPNL the PROMPT-PNL fragment obtained by excluding past modalities $\langle \bar{A} \rangle$, $[\bar{A}]$, $\langle \bar{A}_x \rangle$, and $[\bar{A}_x]$, and we write PROMPT-(R)PNL when we refer to both formalisms. In the following, we will take the liberty of writing PROMPT-(R)PNL formulas not in negation normal form when useful.

PROMPT-(R)PNL models are interval structures enriched with a valuation function for bounding variables in X and a metric over the underlying domain. Formally, a model for PROMPT-(R)PNL (over \mathcal{AP} and X) is a quadruple $\langle \mathbb{D}, V, \sigma, \delta \rangle$, where $\langle \mathbb{D}, V \rangle$ is an interval structure (\mathbb{D} is the *domain* of the model), $\sigma : X \rightarrow \mathbb{R}^{>0}$ is a *valuation function for bounding variables*, and $\delta : \mathbb{D} \times \mathbb{D} \rightarrow \mathbb{R}^{>0}$ is a metric over \mathbb{D} (i.e., the pair (\mathbb{D}, δ) is a metric space) satisfying the additional properties: for every $d, d', d'' \in \mathbb{D}$ (i) if $d < d' < d''$, then $\delta(d, d'') = \delta(d, d') + \delta(d', d'')$, (ii) if d has infinitely many successors in \mathbb{D} , then the set $\{\delta(d, \bar{d}) \mid d < \bar{d}\}$ is not bounded above, and (iii) if d has infinitely many predecessors in \mathbb{D} , then the set $\{\delta(\bar{d}, d) \mid \bar{d} < d\}$ is not bounded above. For a model $M = \langle \mathbb{D}, V, \sigma, \delta \rangle$, we let $\mathbb{D}_M = \mathbb{D}$, $V_M = V$, $\sigma_M = \sigma$, and $\delta_M = \delta$, that is, \mathbb{D}_M , V_M , σ_M , and δ_M denote the four components of M . A PROMPT-(R)PNL model is finite (resp., infinite) if so is its domain.

The *truth value* of a PROMPT-PNL formula over a model and an interval in it is inductively defined as follows:

- $M, [a, b] \models p$ if and only if $p \in V_M([a, b])$, for every $p \in \mathcal{AP}$;
- $M, [a, b] \models \neg p$ if and only if $p \notin V_M([a, b])$, for every $p \in \mathcal{AP}$;
- $M, [a, b] \models \varphi_1 \wedge \varphi_2$ if and only if $M, [a, b] \models \varphi_1$ and $M, [a, b] \models \varphi_2$;
- $M, [a, b] \models \varphi_1 \vee \varphi_2$ if and only if $M, [a, b] \models \varphi_1$ or $M, [a, b] \models \varphi_2$;
- $M, [a, b] \models \langle A \rangle \varphi$ if and only if there is $c \in \mathbb{D}_M^{>b}$ such that $M, [b, c] \models \varphi$;
- $M, [a, b] \models [A] \varphi$ if and only if for all $c \in \mathbb{D}_M^{>b}$ it holds $M, [b, c] \models \varphi$;
- $M, [a, b] \models \langle \bar{A} \rangle \varphi$ if and only if there is $c \in \mathbb{D}_M^{\leq a}$ such that $M, [c, a] \models \varphi$;
- $M, [a, b] \models [\bar{A}] \varphi$ if and only if for all $c \in \mathbb{D}_M^{\leq a}$ it holds $M, [c, a] \models \varphi$;
- $M, [a, b] \models \langle A_x \rangle \varphi$ if and only if there is $c \in \mathbb{D}_M^{>b}$, with $\delta_M(b, c) \leq \sigma_M(x)$, such that $M, [b, c] \models \varphi$, for every $x \in X$;
- $M, [a, b] \models [A_x] \varphi$ if and only if for all $c \in \mathbb{D}_M^{>b}$, with $\delta_M(b, c) \leq \sigma_M(x)$, it holds $M, [b, c] \models \varphi$, for every $x \in X$;
- $M, [a, b] \models \langle \bar{A}_x \rangle \varphi$ if and only if there is $c \in \mathbb{D}_M^{\leq a}$, with $\delta_M(c, a) \leq \sigma_M(x)$, such that $M, [c, a] \models \varphi$, for every $x \in X$;

- $M, [a, b] \models [\overline{A}_x]\varphi$ if and only if for all $c \in \mathbb{D}_M^{\leq a}$, with $\delta_M(c, a) \leq \sigma_M(x)$, it holds $M, [c, a] \models \varphi$, for every $x \in X$.

The truth value of a PROMPT-RPNL formula is obtained, as expected, by restricting to the relevant clauses only.

In PNL, modalities $\langle L \rangle$ and $\langle \overline{L} \rangle$, corresponding to Allen’s relations *later* and *before*, are definable as: $\langle L \rangle\varphi \equiv \langle A \rangle\langle A \rangle\varphi$ and $\langle \overline{L} \rangle\varphi \equiv \langle \overline{A} \rangle\langle \overline{A} \rangle\varphi$. Additionally, in PROMPT-PNL it is possible to define the ‘prompt’ counterparts of modalities $\langle L \rangle$ and $\langle \overline{L} \rangle$ as: $\langle L_x \rangle\varphi \equiv \langle A_x \rangle\langle A_x \rangle\varphi$ and $\langle \overline{L}_x \rangle\varphi \equiv \langle \overline{A}_x \rangle\langle \overline{A}_x \rangle\varphi$. The resulting semantic interpretation for $\langle L_x \rangle$ and $\langle \overline{L}_x \rangle$ is as follows:

- $M, [a, b] \models \langle L_x \rangle\varphi$ if and only if there is $[c, d] \in \mathbb{I}(\mathbb{D}_M)$ such that $b < c$, $\delta_M(b, c) \leq \sigma_M(x)$, $\delta_M(c, d) \leq \sigma_M(x)$, and $M, [c, d] \models \varphi$;
- $M, [a, b] \models \langle \overline{L}_x \rangle\varphi$ if and only if there is $[c, d] \in \mathbb{I}(\mathbb{D}_M)$ such that $d < a$, $\delta_M(d, a) \leq \sigma_M(x)$, $\delta_M(c, d) \leq \sigma_M(x)$, and $M, [c, d] \models \varphi$.

Intuitively, a modality $\langle L_x \rangle$, for some bounding variable x , requires the existence of an event starting and ending within a bounded amount of time after the termination of the current one (modalities $\langle \overline{L}_x \rangle$ impose an analogous constraint in the past). Obviously, only $\langle L_x \rangle$ is definable in PROMPT-RPNL ($\langle \overline{L}_x \rangle$ is not).

The *globally-in-the-future* modality $[G]$ is defined as $[G]\psi \equiv \psi \wedge [A]\psi \wedge [A][A]\psi$, for every PROMPT-PNL formula ψ ; analogously the *prompt-globally-in-the-future* modality $[G_x]$ is defined as $[G_x]\psi \equiv \psi \wedge [A_x]\psi \wedge [A][A_x]\psi$, for every PROMPT-PNL formula ψ and $x \in X$. Given a PROMPT-(R)PNL model M , modalities $[G]$ and $[G_x]$ induce the sets $\mathcal{G}_M^{[a,b]} = \{[a, b]\} \cup \{[c, d] \in \mathbb{I}(\mathbb{D}_M) \mid b \leq c\}$ and $\mathcal{G}_M^{[a,b],x} = \{[a, b]\} \cup \{[c, d] \in \mathbb{I}(\mathbb{D}_M) \mid b \leq c \text{ and } \delta_M(c, d) \leq \sigma_M(x)\}$. We omit the subscript M when it is clear from the context. For every PROMPT-(R)PNL model M , $[a, b] \in \mathbb{I}(\mathbb{D}_M)$, and PROMPT-PNL formula ψ , it holds that $M, [a, b] \models [G]\psi$ if and only if $M, [c, d] \models \psi$ for every $[c, d] \in \mathcal{G}_M^{[a,b]}$ and $M, [a, b] \models [G_x]\psi$ if and only if $M, [c, d] \models \psi$ for every $[c, d] \in \mathcal{G}_M^{[a,b],x}$. Finally, for a model M and $[a, b] \in \mathbb{I}(\mathbb{D}_M)$, we define the *length* of $[a, b]$ (in M) as the value $\delta_M(a, b)$ and, for every $p \in \mathcal{AP}$, if $M, [a, b] \models p$, then we say that $[a, b]$ is a p -interval (in M).

The satisfiability problem. A PROMPT-(R)PNL formula φ is *satisfiable* if, and only if, there exist a PROMPT-(R)PNL model M and an interval $[x, y]$ in M such that $M, [x, y] \models \varphi$. Moreover, a satisfiable formula is said to be *finitely satisfiable* if there exists a finite model for it; otherwise it is *non-finitely satisfiable*. The *satisfiability* (resp., *finite satisfiability*) problem for PROMPT-(R)PNL consists in deciding whether a given PROMPT-(R)PNL formula is satisfiable (resp., finitely satisfiable).

3 Undecidability of PROMPT-RPNL

We prove the undecidability of the satisfiability problem for the logic PROMPT-RPNL (and thus for PROMPT-PNL as well), by a reduction from the *finite coloring problem* (FCP) [18]. An instance of FCP (aka *finite tiling problem*) is a tuple $\Delta = \langle C, H, V, c_i, c_f \rangle$, where C is a finite, non-empty set of *colours*, $H, V \subseteq C \times C$ are total binary relations over the set of colours C , and $c_i, c_f \in C$ are distinguished colours. A solution to Δ is a pair $\langle \mathcal{C}, (K, L) \rangle$, where $K, L \in \mathbb{N}$ and $\mathcal{C} : \{0, \dots, K\} \times \{0, \dots, L\} \rightarrow C$ is a *colouring function* such that $\mathcal{C}(0, 0) = c_i$, $\mathcal{C}(K, L) = c_f$, and, in addition,

- $(\mathcal{C}(i, j), \mathcal{C}(i + 1, j)) \in H$, for each $i < K$ and $j \leq L$ (horizontal constraint), and
- $(\mathcal{C}(i, j), \mathcal{C}(i, j + 1)) \in V$, for each $i \leq K$ and $j < L$ (vertical constraint).

FCP consists in establishing whether there are two natural numbers K and L , and a colouring of the plane $\{0, \dots, K\} \times \{0, \dots, L\}$ such that horizontal and vertical constraints are fulfilled, and bottom-left and top-right colours are given. CFP is undecidable [18, Proposition 7.2]. We encode CFP by means of a PROMPT-RPNL formula. The different aspects of the problem are encoded by means of (blocks of) formulas and the correctness of such partial encodings is testified by the corresponding lemmas below. Clearly, the conjunction of all these formulas is satisfiable if and only if CFP admits a solution. In what follows, we fix an interval model $M = \langle \mathbb{D}, V, \sigma, \delta \rangle$.

For every $d \in \mathbb{D}$ and $x \in X$, we define $\lfloor \sigma \rfloor_d(x) = \max\{\delta(d, d') \in \mathbb{R}^{>0} \mid d' \in \mathbb{D}^{>d} \text{ and } \delta(d, d') \leq \sigma(x)\}$. It clearly holds that $\lfloor \sigma \rfloor_d(x) \leq \sigma(x)$ and, for every $d' \in \mathbb{D}^{\geq d}$, we have that $\delta(d, d') \leq \sigma(x)$ implies $\delta(d, d') \leq \lfloor \sigma \rfloor_d(x)$. For every $x \in X$, there is exactly one point $d' \in \mathbb{D}^{\geq d}$ such that $\delta(d, d') = \lfloor \sigma \rfloor_d(x)$; we call such a point the *x-canonical successor* of d . The length of an interval $[d, d'] \in \mathbb{I}(\mathbb{D})$, where d' is the *x-canonical successor* of d , is said to be *x-canonical*, for every $x \in X$.

Let *succ-upperbound* be the formula $[G](\langle A \rangle \top \rightarrow \langle A_s \rangle \top)$, where $s \in X$.

Lemma 1. *If $M, [a, b] \models \text{succ-upperbound}$ for some $[a, b]$, then for every $c \in \mathbb{D}^{\geq b}$ that is not the greatest element in \mathbb{D} it holds $\delta(c, \text{succ}_{\mathbb{D}}(c)) \leq \lfloor \sigma \rfloor_c(s)$. Moreover, let c' be the *x-canonical successor* of c . If c' is not the greatest element in \mathbb{D} , then $\lfloor \sigma \rfloor_c(x) + \sigma(s) > \sigma(x)$, for every $x \in X$.*

Let *less-than*(x, y) be the formula $[G](\langle A \rangle \top \rightarrow \langle A_y \rangle \text{aux}_{x,y}) \wedge [G][A_x] \neg \text{aux}_{x,y}$ (it is a parametric formula to be instantiated with some $x, y \in X$).

Lemma 2. *If $M, [a, b] \models \text{less-than}(x, y)$ for some $[a, b]$, then $\sigma(x) < \lfloor \sigma \rfloor_c(y)$ holds for every $c \in \mathbb{D}^{\geq b}$, unless c is the greatest element in \mathbb{D} .*

Let \exists -*last* be the conjunction of the following formulas:

$$\neg \text{last} \wedge \langle A \rangle \langle A \rangle \text{last} \wedge [G](\langle A \rangle \text{last} \rightarrow \bigwedge_{p \in \mathcal{AP}} [A](\neg p \wedge [A] \neg p)) \quad (1)$$

$$[G](\langle A \rangle \text{last} \rightarrow [A] \neg \langle A \rangle \text{last}) \quad (2)$$

$$[G](\langle \text{last} \rightarrow \langle A \rangle \text{unique} \rangle \wedge (\langle A \rangle \text{unique} \rightarrow [A] \neg \langle A \rangle \text{unique})) \quad (3)$$

Lemma 3. *If $M, [a, b] \models \exists$ -last for some $[a, b]$, then there is exactly one last-interval in $\mathcal{G}^{[a,b]}$, say it $[c, d]$. Moreover, it holds $c > b$ and there is no p -interval starting in c or after it, for every $p \in \mathcal{AP} \setminus \{\text{last}\}$.*

Let $a \in \mathbb{D}$ and $[c, d] \in \mathbb{I}(\mathbb{D})$ be the unique last-interval (see Lemma 3). Given $p \in \mathcal{AP}$, a *p-chain starting at a* (or, simply, *p-chain*) is a finite sequence of p -intervals $[a_0, b_0], [a_1, b_1], \dots, [a_m, b_m]$ such that $a = a_0$, $b_m = c$, and $b_i = a_{i+1}$ for every $i \in \{0, 1, \dots, m-1\}$. Let *chain*(p, x) be the parametric formula, to be instantiated with some $p \in \mathcal{AP}$ and $x \in X$, defined as the conjunction of the following ones:

$$\text{succ-upperbound} \wedge \exists\text{-last} \quad (4)$$

$$\neg p \wedge \langle A_x \rangle p \wedge [G](\langle p \wedge \neg \langle A \rangle \text{last} \rangle \rightarrow \langle A_x \rangle p) \quad (5)$$

$$[G](\mathbf{p} \rightarrow \mathbf{p}_1 \vee \mathbf{p}_2) \quad (6)$$

$$[G](\langle A \rangle \mathbf{p}_i \rightarrow [A_x][A] \mathbf{p}_i^+) \quad i \in \{1, 2\} \quad (7)$$

$$[G_x](\langle A \rangle \mathbf{p}_i^+ \rightarrow \mathbf{p}_i^-) \quad i \in \{1, 2\} \quad (8)$$

$$[G](\mathbf{p}_i \rightarrow \neg \langle A \rangle \mathbf{p}_i^-) \quad i \in \{1, 2\} \quad (9)$$

$$[G](\langle A \rangle \mathbf{p} \rightarrow [A_x](\neg \mathbf{p} \rightarrow [A] \neg \mathbf{p})) \quad (10)$$

Lemma 4. *If $M, [a, b] \models \text{chain}(\mathbf{p}, x)$ for some $[a, b]$, then there is a finite \mathbf{p} -chain starting at b whose intervals have x -canonical length. Moreover, no other \mathbf{p} -interval exists in $\mathcal{G}^{[a, b], x}$ besides the ones in such a \mathbf{p} -chain.*

We now provide an encoding of a finite plane $\{0, \dots, K\} \times \{0, \dots, L\}$, for some $K, L \in \mathbb{N}$. The idea is to use a \mathbf{u} -chain whose intervals are either tile-intervals, encoding some point of the finite plane, or $*$ -intervals, which are used as separators between rows of the plane. Let *plane* be the conjunction of the following formulas:

$$\text{less-than}(s, x) \wedge \text{less-than}(x, y) \wedge \text{chain}(\mathbf{u}, x) \wedge \text{chain}(\text{row}, y) \quad (11)$$

$$[G](\langle \mathbf{u} \leftrightarrow * \vee \text{tile} \rangle \wedge \langle * \rightarrow \neg \text{tile} \rangle) \quad (12)$$

$$\langle A \rangle * \wedge [G](\langle * \rightarrow \langle A \rangle \text{tile} \rangle \wedge \langle \mathbf{u} \wedge \langle A \rangle \text{last} \rightarrow \text{tile} \rangle) \quad (13)$$

$$[G](\langle A \rangle \text{row} \rightarrow \langle A \rangle *) \quad (14)$$

$$[G](\langle A \rangle * \rightarrow [A_y](\langle A \rangle * \rightarrow \text{row})) \quad (15)$$

Lemma 5. *If $M, [a, b] \models \text{plane}$ for some $[a, b]$, then there is a finite sequence of points $b = p_0^0 < p_1^1 < \dots < p_{n_1}^1 = p_0^2 < p_1^2 < \dots < p_{n_2}^2 = p_0^3 < \dots < p_{n_{r-1}}^{r-1} = p_0^r < \dots < p_{n_r}^r$, with $r \geq 1$ and $n_i > 1$ for every $i \in \{1, \dots, r\}$ such that: (i) $[p_0^i, p_1^i]$ is a $*$ -interval and its length is x -canonical, for every $i \in \{1, \dots, r\}$; (ii) $[p_j^i, p_{j+1}^i]$ is a tile-interval and its length is x -canonical, for every $i \in \{1, \dots, r\}$ and $j \in \{1, \dots, n_i - 1\}$; (iii) $[p_0^i, p_0^{i+1}]$ is a row-interval and its length is y -canonical, for every $i \in \{1, \dots, r - 1\}$; (iv) $M, [p_{n_r}^r, p^r]$ is the unique last-interval, for some $p^r > p_{n_r}^r$. Moreover, no other $*$ -interval (resp., tile-interval) exists in $\mathcal{G}^{[a, b], x}$.*

The encoding of the finite plane $\{0, \dots, K\} \times \{0, \dots, L\}$ we have obtained so far is incomplete, the problem being that rows (row-intervals) do not necessarily contain the same number of tiles (tile-intervals). In order to overcome such a problem, we introduce below *corr*-intervals, which are used to link the i th tile-interval of a row to the i th tile-interval of the next row (if any) and to the i th tile-interval of the previous row (if any). This will guarantee that each row of our encoding features the same number of tiles.

Let *w-def* be the conjunction of the following formulas:

$$\text{less-than}(x, w) \wedge \text{less-than}(w, y) \quad (16)$$

$$[A_y] \neg \langle A \rangle * \text{-aux} \wedge \langle \langle A_y \rangle (\text{row} \wedge \neg \langle A_x \rangle \text{last}) \rightarrow \langle A \rangle \langle A \rangle * \text{-aux} \rangle \quad (17)$$

$$\langle A_s \rangle \langle A_w \rangle \langle [A](\neg \text{last} \wedge [A] \neg \text{last}) \vee \langle A \rangle * \text{-aux} \rangle \quad (18)$$

Lemma 6. *If $M, [a, b] \models \text{plane} \wedge \text{w-def}$ for some $[a, b]$, then $\sigma(w) < \lfloor \sigma \rfloor_c(y) \leq \sigma(y) < \sigma(w) + \sigma(s)$ for every $c \in \mathbb{D}^{\geq b}$, unless c is the greatest element in \mathbb{D} .*

Let *correspondence* be the conjunction of the following formulas:

$$\text{plane} \wedge w\text{-def} \wedge \text{less-than}(s, z) \wedge \text{less-than}(z, x) \quad (19)$$

$$[G](\langle A_x \rangle u \rightarrow [A_z] \langle A_z \rangle (\text{u-suffix} \wedge (\langle A_x \rangle u \vee \langle A \rangle \text{last}))) \quad (20)$$

$$[G_s] \neg \text{u-suffix} \quad (21)$$

$$[G](\langle \text{row} \wedge \neg \langle A \rangle \text{last} \rangle \rightarrow \text{corr}) \quad (22)$$

$$[G](\langle \langle A_x \rangle \text{tile} \wedge \langle A \rangle \langle A_x \rangle * \rangle \rightarrow \langle A_y \rangle \text{corr}) \quad (23)$$

$$[G_w] \neg \text{corr} \quad (24)$$

$$[G](\text{corr} \rightarrow \langle A \rangle \text{tile}) \quad (25)$$

$$[G](\langle \langle A_x \rangle \langle \text{tile} \wedge \langle A_x \rangle * \rangle \rangle \rightarrow [A_y](\text{corr} \rightarrow \langle A_x \rangle \langle \text{tile} \wedge \langle A_x \rangle * \rangle)) \quad (26)$$

Lemma 7. *If $M, [a, b] \models \text{correspondence}$ for some $[a, b]$, then $[p_j^i, p_j^{i+1}]$ is a *corr-interval*, with $\sigma(w) < \delta(p_j^i, p_j^{i+1}) \leq \sigma(y)$, for every $i \in \{1, \dots, r-1\}$ and $j \in \{0, \dots, n_i-1\}$. Moreover, for every $i \in \{1, \dots, r-1\}$, it holds that $n_i = n_{i+1}$.*

Now, let $\Delta = \langle C, H, V, c_i, c_f \rangle$ be an instance of FCP and let φ_Δ be the conjunction of the following formulas:

$$\text{correspondence} \wedge \langle A_x \rangle c_i \wedge [G_x](\langle \text{tile} \wedge \langle A \rangle \text{last} \rangle \rightarrow c_f) \quad (27)$$

$$[G_x](\text{tile} \leftrightarrow \bigvee_{c \in C} c) \wedge [G](\bigwedge_{c, c' \in C, c \neq c'} \neg(c \wedge c')) \quad (28)$$

$$[G](\langle \langle A_x \rangle \langle \text{tile} \wedge \langle A_x \rangle \text{tile} \rangle \rightarrow \bigvee_{(c, c') \in H} \langle A_x \rangle (c \wedge \langle A_x \rangle c')) \quad (29)$$

$$[G_x](\langle \langle \langle A_x \rangle \text{tile} \wedge \langle A_y \rangle \text{corr} \rangle \rightarrow \bigvee_{(c, c') \in V} (\langle A_x \rangle c \wedge [A_y](\text{corr} \rightarrow \langle A_x \rangle c')) \rangle) \quad (30)$$

Lemma 8. *The formula φ_Δ is satisfiable iff the CFP instance Δ has a positive answer.*

Theorem 1. *The satisfiability problem for PROMPT-RPNL, and thus the one for PROMPT-PNL, is undecidable.*

4 Decidability of PROMPT^d-PNL

In this section, we show how to restrict the use of prompt modalities to get a fragment of PROMPT-PNL with a decidable satisfiability problem.

We define PROMPT^d-PNL as the fragment of PROMPT-PNL obtained by using disjoint sets of bounding variables for existential and universal prompt modalities. Formally, let us partition the set X of bounding variables into sets X_\diamond and X_\square . The syntax of PROMPT^d-PNL is defined as:

$$\begin{aligned} \varphi ::= & p \mid \varphi \wedge \varphi \mid \langle A \rangle \varphi \mid \langle \bar{A} \rangle \varphi \mid \langle A_x \rangle \varphi \mid \langle \bar{A}_x \rangle \varphi \\ & \mid \neg p \mid \varphi \vee \varphi \mid [A] \varphi \mid [\bar{A}] \varphi \mid [A_y] \varphi \mid [\bar{A}_y] \varphi \end{aligned}$$

where $p \in \mathcal{AP}$, $x \in X_\diamond$, and $y \in X_\square$. Since PROMPT^d-PNL is a syntactic restriction of PROMPT-PNL, both formalisms share the same semantics. In particular, a PROMPT-PNL model is a PROMPT^d-PNL model as well. Analogously to the unrestricted case, we define PROMPT^d-RPNL as PROMPT^d-PNL devoid of past modalities $\langle \bar{A} \rangle$, $[\bar{A}]$, $\langle \bar{A}_x \rangle$, and $[\bar{A}_y]$.

PROMPT^dPNL is not closed under negation. For any given PROMPT^dPNL formula ψ , we inductively define $neg(\psi)$ as shown in Table 1 ($neg(\psi)$ is not necessarily a PROMPT^dPNL formula). If ψ is a (non-prompt) PNL formula, then $neg(\psi) \equiv \neg\psi$. Moreover, we define $neg(\sim\psi)$ as ψ and thus we have that $neg(neg(\psi)) \equiv \psi$, for every PROMPT^dPNL formula ψ .

ψ	$neg(\psi)$	ψ	$neg(\psi)$
p	$\neg p$	$\neg p$	p
$\psi_1 \wedge \psi_2$	$neg(\psi_1) \vee neg(\psi_2)$	$\psi_1 \vee \psi_2$	$neg(\psi_1) \wedge neg(\psi_2)$
$\langle A \rangle \psi_1$	$[A]neg(\psi_1)$	$[A]\psi_1$	$\langle A \rangle neg(\psi_1)$
$\langle \bar{A} \rangle \psi_1$	$[\bar{A}]neg(\psi_1)$	$[\bar{A}]\psi_1$	$\langle \bar{A} \rangle neg(\psi_1)$
ψ			$neg(\psi)$
$\langle A_x \rangle \psi_1$ or $\langle \bar{A}_x \rangle \psi_1$ or $[A_y]\psi_1$ or $[\bar{A}_y]\psi_1$			$\sim\psi$

Table 1. Definition of $neg(\psi)$, for a PROMPT^dPNL formula ψ

A close analysis of the proof of the undecidability of PROMPT-(R)PNL reveals that the unrestricted use of bounding variables within prompt modalities allows one to somehow establish tight bounds for the length of intervals, and this ability is crucial to the encoding. We are going to show that decidability can be recovered by not allowing both existential and universal prompt quantification on the same bounding variable. Intuitively, decidability follows from the fact that, when disjoint sets of bounding variables are used within existential and universal prompt modalities, formulas enjoy a monotonicity property, which does not hold for unrestricted PROMPT-(R)PNL formulas.

Let $M = \langle \mathbb{D}, V, \sigma, \delta \rangle$ be a PROMPT-PNL model, $x \in X$, and $r \in \mathbb{R}^{>0}$. We denote by $M_{[x:=r]}$ the model $\langle \mathbb{D}, V, \sigma', \delta \rangle$, where $\sigma'(x) = r$ and $\sigma'(x') = \sigma(x')$ for every $x' \in X$ with $x' \neq x$.

Proposition 1 (monotonicity). *Let ψ be a formula of PROMPT^dPNL, M be a model of PROMPT^dPNL, and $[a, b]$ be an interval in M . If $M, [a, b] \models \psi$, then $M_{[x:=r]}, [a, b] \models \psi$ for all $x \in X_\diamond$ and $r \in \mathbb{R}^{>0}$, with $r \geq \sigma_M(x)$. In a dual fashion, if $M, [a, b] \models \psi$, then $M_{[y:=r]}, [a, b] \models \psi$ for all $y \in X_\square$ and $r \in \mathbb{R}^{>0}$, with $r \leq \sigma_M(y)$.*

Checking that the above monotonicity property holds for PROMPT^dPNL is immediate. To see that it does not hold for PROMPT-PNL, consider the formula $\psi = [A_y]\neg p \wedge \langle A_x \rangle p \wedge [A_x]\neg q \wedge \langle A_z \rangle q$. Clearly, ψ is satisfiable and all of its models are such that the value of x is bounded below by the value of y and above by the value of z .

By Proposition 1, when studying the (finite) satisfiability problem for PROMPT^dPNL we can assume, w.l.o.g., that $|X_\diamond| = |X_\square| = 1$, as every formula ψ , featuring (possibly) more than one bounding variable in X_\diamond or X_\square , can be transformed into an equisatisfiable one ψ' , obtained by replacing two distinguished (chosen randomly) variables $\hat{x} \in X_\diamond$ and $\hat{y} \in X_\square$ for every $x \in X_\diamond$ and $y \in X_\square$, respectively. It is not difficult to check that, due to monotonicity, ψ is (finitely) satisfiable if and only if so is ψ' . Therefore, for the remainder of the section, we set $X_\diamond = \{x\}$ and $X_\square = \{y\}$.

Finite satisfiability. The finite satisfiability problem for PROMPT^dPNL can be reduced to the one for plain PNL, known to be NEXPTIME-complete [8]. Let ψ be a formula of PROMPT^dPNL and let $plain(\psi)$ be the PNL formula obtained from ψ by:

- (i) replacing existential prompt modalities by the corresponding non-prompt versions (i.e., substituting $\langle A \rangle$ for $\langle A_x \rangle$ and $\langle \bar{A} \rangle$ for $\langle \bar{A}_x \rangle$), and
- (ii) replacing all sub-formulas of the forms $[A_y]\psi$ and $[\bar{A}_y]\psi$ by the constant \top .

It is not difficult to show by induction on the structure of ψ that if ψ is finitely satisfiable, so is $plain(\psi)$. On the other hand, if $plain(\psi)$ is finitely satisfiable, then let $M_{plain(\psi)} = \langle \mathbb{D}, V \rangle$ be a PNL model such that $M_{plain(\psi)}, [a, b] \models plain(\psi)$ for some $[a, b] \in \mathbb{I}(\mathbb{D})$. We define $\delta(d, d') = |\{d'' \in \mathbb{D} \mid d < d'' \leq d'\}|$ for every $d, d' \in \mathbb{D}$. Since \mathbb{D} is finite, both $\max_\delta = \max\{\delta(d, d') \mid d, d' \in \mathbb{D} \text{ and } d \neq d'\}$ and $\min_\delta = \min\{\delta(d, d') \mid d, d' \in \mathbb{D} \text{ and } d \neq d'\}$ are well defined, thus we can set $\sigma(x) = \max_\delta$, and $\sigma(y) = \frac{\min_\delta}{2}$. It is possible to show that $M = \langle \mathbb{D}, V, \sigma, \delta \rangle$ is such that $M, [a, b] \models \psi$. Therefore, ψ is finitely satisfiable, too.

Theorem 2. *The finite satisfiability problem for $PROMPT^d\text{-PNL}$ is NEXPTIME-complete.*

In order to deal with formulas that are non-finitely satisfiable, in what follows we show how the search for an infinite model can be reduced to the search for a finite witness for it, within a finite search space. Decidability of the satisfiability problem for $PROMPT^d\text{-PNL}$ immediately follows.

4.1 Prompt labeled interval structures

In this subsection we define labeled interval structures for $PROMPT^d\text{-PNL}$ formulas, which are, intuitively, extended models, where intervals are labeled with sets of sub-formulas (instead of sets of atomic propositions) of the considered formula. From now on, we let φ be a generic $PROMPT^d\text{-PNL}$ formula.

Let $Sub(\varphi)$ be the set of all sub-formulas of φ and let $Sub^\neg(\varphi) = \{neg(\psi) \mid \psi \in Sub(\varphi)\}$. The *closure* of φ , denoted by $Cl(\varphi)$, is the set $Sub(\varphi) \cup Sub^\neg(\varphi) \cup \{\langle A \rangle\varphi, neg(\langle A \rangle\varphi)\}$. Clearly, $|Cl(\varphi)| \leq 2 \cdot |\varphi| + 2$ holds.

A *future temporal request* of φ is a formula in $Cl(\varphi)$ having one of the following forms: $\langle A \rangle\psi, neg(\langle A \rangle\psi), \langle A_x \rangle\psi, neg(\langle A_x \rangle\psi), [A_y]\psi, neg([A_y]\psi)$, for some ψ . Analogously, a *past temporal request* of φ is a formula in $Cl(\varphi)$ having one of the following forms: $\langle \bar{A} \rangle\psi, neg(\langle \bar{A} \rangle\psi), \langle \bar{A}_x \rangle\psi, neg(\langle \bar{A}_x \rangle\psi), [\bar{A}_y]\psi, neg([\bar{A}_y]\psi)$, for some ψ . We denote by $TR_f(\varphi)$ (resp., $TR_p(\varphi)$) the set of future (resp., past) temporal requests of φ . In addition, the set of *temporal requests* of φ , denoted by $TR(\varphi)$, is defined as $TR_f(\varphi) \cup TR_p(\varphi)$.

A φ -*atom* is a subset A of $Cl(\varphi)$ such that, for every $\psi, \psi_1, \psi_2 \in Cl(\varphi)$, (i) $\psi \in A$ if and only if $neg(\psi) \notin A$, and (ii) $\psi_1 \vee \psi_2 \in A$ if and only if $\psi_1 \in A$ or $\psi_2 \in A$. Notice that conditions (i) and (ii) imply $\psi_1 \wedge \psi_2 \in A$ if and only if $\psi_1 \in A$ and $\psi_2 \in A$. We denote the set of φ -atoms by \mathcal{A}_φ .

A *prompt φ -labeled interval structure* ($pLIS_\varphi$) is a 5-tuple $\mathbf{L} = \langle \mathbb{D}, \mathcal{L}, \delta, \mathcal{X}, \mathcal{Y} \rangle$, where (\mathbb{D}, δ) is a metric space, $\mathcal{L} : \mathbb{I}(\mathbb{D}) \rightarrow \mathcal{A}_\varphi$ is a *labeling function* (or simply *labeling*) such that $\varphi \in \mathcal{L}([a, b])$ for some $[a, b] \in \mathbb{I}(\mathbb{D})$, and $\mathcal{X}, \mathcal{Y} \in \mathbb{N}$ are the *existential* and the *universal* bound, respectively. Sometimes, for the sake of brevity, we omit the last three components of the 5-tuple and we denote a $pLIS_\varphi$ as a 2-tuple $\langle \mathbb{D}, \mathcal{L} \rangle$ instead. Moreover, given a $pLIS_\varphi$ $\mathbf{L} = \langle \mathbb{D}, \mathcal{L} \rangle$, we denote by $\mathbb{D}_\mathbf{L}$ its underlying domain \mathbb{D} and by $\mathcal{L}_\mathbf{L}$ the labeling function \mathcal{L} . A $pLIS_\varphi$ \mathbf{L} is *finite* (resp., *infinite*) if so is $\mathbb{D}_\mathbf{L}$.

Given a $\text{pLIS}_\varphi \mathbf{L}$ and a point $d \in \mathbb{D}_{\mathbf{L}}$ we define the set of *future requests of d in \mathbf{L}* , denoted by $\text{f-REQ}^{\mathbf{L}}(d)$, as $\bigcup_{d' \in \mathbb{D}^{<d}} (\mathcal{L}_{\mathbf{L}}(d', d) \cap \text{TR}_f(\varphi))$, the set of *past requests of d in \mathbf{L}* , denoted by $\text{p-REQ}^{\mathbf{L}}(d)$, as $\bigcup_{d' \in \mathbb{D}^{>d}} (\mathcal{L}_{\mathbf{L}}(d, d') \cap \text{TR}_p(\varphi))$, and the set of *requests of d in \mathbf{L}* , denoted by $\text{REQ}^{\mathbf{L}}(d)$, as $\text{f-REQ}^{\mathbf{L}}(d) \cup \text{p-REQ}^{\mathbf{L}}(d)$. We denote by REQ_φ the class of all sets of requests, i.e., $\text{REQ}_\varphi = \{\mathcal{R} \mid \mathcal{R} = \text{REQ}^{\mathbf{L}}(d) \text{ for some } \text{pLIS}_\varphi \mathbf{L} \text{ and } d \in \mathbb{D}_{\mathbf{L}}\}$. We have that $|\text{REQ}_\varphi| \leq 2^{|\text{Cl}(\varphi)|} \leq 2^{2 \cdot |\varphi| + 2}$.

An *existential request* of φ is a temporal request of φ of one the following forms: $\langle A \rangle \psi$, $\langle \bar{A} \rangle \psi$, $\langle A_x \rangle \psi$, $\langle \bar{A}_x \rangle \psi$, $\text{neg}(\langle A_y \rangle \psi)$, and $\text{neg}(\langle \bar{A}_y \rangle \psi)$, for some ψ . A *universal request* of φ is a temporal request of φ that is not an existential one. Let $\mathbf{L} = \langle \mathbb{D}, \mathcal{L}, \delta, \mathcal{X}, \mathcal{Y} \rangle$ be a pLIS_φ and $d \in \mathbb{D}$. We define $\exists\text{-REQ}^{\mathbf{L}}(d) = \{\psi \in \text{REQ}^{\mathbf{L}}(d) \mid \psi \text{ is an existential request of } \varphi\}$ and $\forall\text{-REQ}^{\mathbf{L}}(d) = \text{REQ}^{\mathbf{L}}(d) \setminus \exists\text{-REQ}^{\mathbf{L}}(d)$.

For $\psi \in \exists\text{-REQ}^{\mathbf{L}}(d)$, we say that ψ is *fulfilled in (\mathbf{L}, d) by $d' \in \mathbb{D}$* if, and only if, one of the following holds:

- $\psi = \langle A \rangle \psi'$ for some ψ' and $\psi' \in \mathcal{L}([d, d'])$,
- $\psi = \langle \bar{A} \rangle \psi'$ for some ψ' and $\psi' \in \mathcal{L}([d', d])$,
- $\psi = \langle A_x \rangle \psi'$ for some ψ' , $\psi' \in \mathcal{L}([d, d'])$, and $\delta(d, d') \leq \mathcal{X}$,
- $\psi = \langle \bar{A}_x \rangle \psi'$ for some ψ' , $\psi' \in \mathcal{L}([d', d])$, and $\delta(d', d) \leq \mathcal{X}$,
- $\psi = \text{neg}(\langle A_y \rangle \psi')$ for some ψ' , $\text{neg}(\psi') \in \mathcal{L}([d, d'])$, and $\delta(d, d') \leq \mathcal{Y}$,
- $\psi = \text{neg}(\langle \bar{A}_y \rangle \psi')$ for some ψ' , $\text{neg}(\psi') \in \mathcal{L}([d', d])$, and $\delta(d', d) \leq \mathcal{Y}$.

ψ is *fulfilled in (\mathbf{L}, d)* if and only if there is d' such that ψ is fulfilled in (\mathbf{L}, d) by d' .

For $\psi \in \forall\text{-REQ}^{\mathbf{L}}(d)$, we say that ψ is *fulfilled in (\mathbf{L}, d)* if, and only if, one of the following holds:

- $\psi = [A] \psi'$ for some ψ' and $\psi' \in \mathcal{L}([d, d'])$ for every $d' \in \mathbb{D}^{>d}$,
- $\psi = [\bar{A}] \psi'$ for some ψ' and $\psi' \in \mathcal{L}([d', d])$ for every $d' \in \mathbb{D}^{<d}$,
- $\psi = [A_y] \psi'$ for some ψ' and $\psi' \in \mathcal{L}([d, d'])$ for every $d' \in \mathbb{D}^{>d}$ with $\delta(d, d') \leq \mathcal{Y}$,
- $\psi = [\bar{A}_y] \psi'$ for some ψ' and $\psi' \in \mathcal{L}([d', d])$ for every $d' \in \mathbb{D}^{<d}$ with $\delta(d', d) \leq \mathcal{Y}$,
- $\psi = \text{neg}(\langle A_x \rangle \psi')$ for some ψ' and $\text{neg}(\psi') \in \mathcal{L}([d, d'])$ for every $d' \in \mathbb{D}^{>d}$ with $\delta(d, d') \leq \mathcal{X}$,
- $\psi = \text{neg}(\langle \bar{A}_x \rangle \psi')$ for some ψ' and $\text{neg}(\psi') \in \mathcal{L}([d', d])$ for every $d' \in \mathbb{D}^{>d}$ with $\delta(d', d) \leq \mathcal{X}$.

d is *\exists -fulfilled in \mathbf{L}* if, and only if, every $\psi \in \exists\text{-REQ}^{\mathbf{L}}(d)$ is fulfilled; d is *\forall -fulfilled in \mathbf{L}* if, and only if, every $\psi \in \forall\text{-REQ}^{\mathbf{L}}(d)$ is fulfilled; d is *fulfilled in \mathbf{L}* if, and only if, it is both \exists - and \forall -fulfilled in \mathbf{L} .

An *existentially fulfilling* (resp., *universally fulfilling*, *fulfilling*) pLIS_φ , aka $\exists\text{-pLIS}_\varphi$ (resp., $\forall\text{-pLIS}_\varphi$, $\exists\forall\text{-pLIS}_\varphi$), is a $\text{pLIS}_\varphi \mathbf{L}$ such that every $d \in \mathbb{D}_{\mathbf{L}}$ is \exists -fulfilled (resp., \forall -fulfilled, fulfilled) in it.

Proposition 2. φ is satisfiable if and only if there exists a $\exists\forall\text{-pLIS}_\varphi$, and it is finitely satisfiable if and only if there exists a finite $\exists\forall\text{-pLIS}_\varphi$.

Before showing the decidability of $\text{PROMPT}^d\text{-PNL}$, we prove a result that will later come in handy. A set of requests $\text{REQ}^{\mathbf{L}}(d)$ (for a $\text{pLIS}_\varphi \mathbf{L}$ and $d \in \mathbb{D}_{\mathbf{L}}$) is *consistent* if for each $\psi \in \text{REQ}^{\mathbf{L}}(d)$, we have that $\text{neg}(\psi) \notin \text{REQ}^{\mathbf{L}}(d)$; otherwise, it is *inconsistent*.

Proposition 3. Let \mathbf{L} be a pLIS_φ and $d \in \mathbb{D}_{\mathbf{L}}$. The following properties hold, unless $\text{REQ}^{\mathbf{L}}(d)$ is inconsistent:

- if $\mathbb{D}^{<d} \neq \emptyset$, then $\text{f-REQ}^{\mathbf{L}}(d) = \mathcal{L}_{\mathbf{L}}(d', d) \cap \text{TR}_f(\varphi)$, for any given $d' \in \mathbb{D}^{<d}$, unless $\text{f-REQ}^{\mathbf{L}}(d)$ is inconsistent;
- if $\mathbb{D}^{>d} \neq \emptyset$, then $\text{p-REQ}^{\mathbf{L}}(d) = \mathcal{L}_{\mathbf{L}}(d, d') \cap \text{TR}_p(\varphi)$, for any given $d' \in \mathbb{D}^{>d}$, unless $\text{p-REQ}^{\mathbf{L}}(d)$ is inconsistent.

4.2 A bounded witness for non-finitely satisfiable formulas

Let \mathbf{L} be a pLIS_{φ} and $d \in \mathbb{D}_{\mathbf{L}}$. A set of *essentials* of d (in \mathbf{L}) is any minimal (with respect to set inclusion) set $E \subseteq \mathbb{D}_{\mathbf{L}}$ such that for every $\psi \in \exists\text{-REQ}^{\mathbf{L}}(d)$ there is $d' \in E$ for which ψ is fulfilled in (\mathbf{L}, d) by d' . We denote by $\mathcal{E}_{\mathbf{L}}(d)$ the class containing all sets of essentials of d in \mathbf{L} , i.e., $\mathcal{E}_{\mathbf{L}}(d) = \{E \subseteq \mathbb{D}_{\mathbf{L}} \mid E \text{ is a set of essentials of } d \text{ in } \mathbf{L}\}$. Intuitively, a set of essentials of d is a collection of points that jointly make d \exists -fulfilled in \mathbf{L} . Clearly $\mathcal{E}_{\mathbf{L}}(d) \neq \emptyset$ if and only if d is \exists -fulfilled in \mathbf{L} . We lift this concept to a higher order: a set of *essentials of essentials* (or *2nd-order essentials*) of d (in \mathbf{L}) is any minimal (with respect to set inclusion) set $E_2 \subseteq \mathbb{D}_{\mathbf{L}}$ such that (i) $E_1 \subseteq E_2$ for some $E_1 \in \mathcal{E}_{\mathbf{L}}(d)$ and (ii) for every $d' \in E_1$ there is $E_{d'} \in \mathcal{E}_{\mathbf{L}}(d')$ for which $E_{d'} \subseteq E_2$. We denote by $\mathcal{E}_{\mathbf{L}}^2(d)$ the class containing all sets of 2nd-order essentials of d in \mathbf{L} , i.e., $\mathcal{E}_{\mathbf{L}}^2(d) = \{E \subseteq \mathbb{D}_{\mathbf{L}} \mid E \text{ is a set of 2nd-order essentials of } d \text{ in } \mathbf{L}\}$.

Definition 1 (representative). Let \mathbf{L} be a finite pLIS_{φ} and $d \in \mathbb{D}_{\mathbf{L}}$.

If $d \notin \{\min \mathbb{D}_{\mathbf{L}}, \max \mathbb{D}_{\mathbf{L}}\}$, then a representative of d in \mathbf{L} is a point $e \in \mathbb{D}_{\mathbf{L}}$ such that $\text{REQ}^{\mathbf{L}}(d) = \text{REQ}^{\mathbf{L}}(e)$, e is fulfilled in \mathbf{L} , and so are points in E_2 , for some $E_2 \in \mathcal{E}_{\mathbf{L}}^2(e)$ with $E_2 \cap \{\min \mathbb{D}, \max \mathbb{D}\} = \emptyset$.

If $d = \min \mathbb{D}_{\mathbf{L}}$ (resp., $d = \max \mathbb{D}_{\mathbf{L}}$), then a representative of d in \mathbf{L} is a point $e \in \mathbb{D}_{\mathbf{L}}$ that is a representative of d' in \mathbf{L} for some $d' \in \mathbb{D}_{\mathbf{L}}$, with $\text{p-REQ}^{\mathbf{L}}(d') = \text{p-REQ}^{\mathbf{L}}(d)$ (resp., $\text{f-REQ}^{\mathbf{L}}(d') = \text{f-REQ}^{\mathbf{L}}(d)$).

A convex subset of a domain \mathbb{D} is a subset \mathbb{D}' of \mathbb{D} such that for every $d', d'' \in \mathbb{D}'$ and $d \in \mathbb{D}$, if $d' < d < d''$, then $d \in \mathbb{D}'$. A right-convex (resp., left-convex) subset of a domain \mathbb{D} is a convex subset \mathbb{D}' of \mathbb{D} such that $\max \mathbb{D} \in \mathbb{D}'$ (resp., $\min \mathbb{D} \in \mathbb{D}'$).

Given a pLIS_{φ} \mathbf{L} and $\mathbb{D}' \subseteq \mathbb{D}_{\mathbf{L}}$, we let $\text{request-sets}^{\mathbf{L}}(\mathbb{D}') = \{\mathcal{R} \mid \text{REQ}^{\mathbf{L}}(d) = \mathcal{R} \text{ for some } d \in \mathbb{D}'\}$.

Definition 2 (left- and right-periodic pLIS_{φ}). Let \mathbf{L} be a finite pLIS_{φ} . A left-period for \mathbf{L} is a left-convex subset \mathbb{E} of $\mathbb{D}_{\mathbf{L}}$ such that, for every $d \in \mathbb{E}$, if d is not fulfilled in \mathbf{L} or $d = \min \mathbb{E}$, then there is $d' \in \mathbb{E}^{>d}$ for which the following holds:

- a) d' is a representative of d in \mathbf{L} ;
- b) $\text{request-sets}^{\mathbf{L}}(\mathbb{E} \setminus \{\min \mathbb{E}\})$ is equal to $\text{request-sets}^{\mathbf{L}}(\mathbb{E}^{<d'} \setminus \{\min \mathbb{E}\})$, which is equal to $\text{request-sets}^{\mathbf{L}}(\mathbb{E}^{>d'})$, and there are $d'' \in \mathbb{E}^{<d'} \setminus \{\min \mathbb{E}\}$ and $d''' \in \mathbb{E}^{>d'}$ such that $\text{p-REQ}^{\mathbf{L}}(\min \mathbb{E}) = \text{p-REQ}^{\mathbf{L}}(d'') = \text{p-REQ}^{\mathbf{L}}(d''')$;
- c) every $\langle A_x \rangle \psi \in \text{f-REQ}^{\mathbf{L}}(d')$ is fulfilled in (\mathbf{L}, d') by a point belonging to \mathbb{E} .

A right-period for \mathbf{L} is defined symmetrically.

\mathbf{L} is periodic if, and only if, there exist both a left- and a right-period for it.

Definition 3 (φ -witness). A φ -witness is a finite, periodic \forall - pLIS_{φ} \mathbf{L} , such that every $d \in \mathbb{D}_{\mathbf{L}} \setminus (\mathbb{E} \cup \mathbb{F})$ is fulfilled in \mathbf{L} , where \mathbb{E} and \mathbb{F} are, respectively, a left- and a right-period for \mathbf{L} , with $\mathbb{E} \cap \mathbb{F} = \emptyset$ and $\mathbb{D}_{\mathbf{L}} \setminus (\mathbb{E} \cup \mathbb{F}) \neq \emptyset$.

Lemma 9. *An infinite $\exists\forall$ -pLIS $_{\varphi}$ $\mathbf{L} = \langle \mathbb{D}, \mathcal{L}, \delta, \mathcal{X}, \mathcal{Y} \rangle$ exists if and only if a φ -witness $\mathbf{L}' = \langle \mathbb{D}', \mathcal{L}', \delta', \mathcal{X}', \mathcal{Y}' \rangle$ exists.*

Thanks to the previous lemma, we can reduce the search for an infinite model for a formula to the search for a finite witness. However, since such a finite witness can be arbitrarily large, the search space is still infinite. In what follows, we provide a bound on the size of the finite witness, thus obtaining a finite search space. Decidability of $\text{PROMPT}^d\text{-PNL}$ immediately follows.

$$\text{Let } \mathcal{B}_{\varphi} = |\text{REQ}_{\varphi}| \cdot (2 \cdot |\text{Cl}(\varphi)|^2 + 2 \cdot |\text{Cl}(\varphi)|) + |\text{REQ}_{\varphi}| \cdot |\text{Cl}(\varphi)| + |\text{Cl}(\varphi)|.$$

Lemma 10. *Let $\mathbf{L} = \langle \mathbb{D}, \mathcal{L}, \delta, \mathcal{X}, \mathcal{Y} \rangle$ be a φ -witness, \mathbb{E} and \mathbb{F} being, respectively, a left- and a right-period for it. If $|\mathbb{E}| > \mathcal{B}_{\varphi}$ (resp., $|\mathbb{F}| > \mathcal{B}_{\varphi}$, $|\mathbb{D} \setminus (\mathbb{E} \cup \mathbb{F})| > \mathcal{B}_{\varphi}$), then there is a φ -witness $\mathbf{L}' = \langle \mathbb{D}', \mathcal{L}', \delta', \mathcal{X}', \mathcal{Y}' \rangle$ with $|\mathbb{D}'| = |\mathbb{D}| - 1$.*

The size of a pLIS $_{\varphi}$ \mathbf{L} is the size of the underlying domain $\mathbb{D}_{\mathbf{L}}$. The following corollary immediately follows from the above lemma.

Corollary 1 (small model property). *A φ -witness exists if and only if there is one of size at most $3 \cdot \mathcal{B}_{\varphi} \leq 3 \cdot [2^{2 \cdot |\varphi| + 2} \cdot (2 \cdot (2 \cdot |\varphi| + 2)^2 + 2 \cdot (2 \cdot |\varphi| + 2)) + 2^{2 \cdot |\varphi| + 2} \cdot (2 \cdot |\varphi| + 2) + (2 \cdot |\varphi| + 2)]$.*

Theorem 3. *The satisfiability problem for $\text{PROMPT}^d\text{-PNL}$ is NEXPTIME-complete.*

5 Conclusions

In this paper, we have studied the problem of enriching the well-known propositional logic of temporal neighborhood PNL with support for prompt-liveness specifications. We first proved that the logic obtained from PNL by introducing “prompt” versions of its modalities with no restriction on the use of bounding variables, that we call PROMPT-PNL , is undecidable. Then, we showed that decidability can be recovered by introducing a partition of bounding variables into two classes, one for the existential modalities, the other for the universal ones. The satisfiability problem for the resulting logic, named $\text{PROMPT}^d\text{-PNL}$, is indeed NEXPTIME-complete.

The work done can be further developed in various directions.

First, we are interested in identifying the minimum number of bounding variables that suffice to make PROMPT-PNL undecidable. We believe it possible to prove that when the set of variables is small enough, e.g., when it includes two bounding variables only, the logic is still expressive enough to capture some meaningful promptness conditions and remains decidable.

We also aim at investigating the more powerful setting of *parametric* extensions of PNL. Parametric PNL can be viewed as a natural generalization of PROMPT-PNL , as *parametric modalities* allow one to express both lower and upper bounds on the delay with which a request is fulfilled (PROMPT-PNL only copes with the latter).

Last but not least, we are interested in comparing the expressiveness of the logics PROMPT-PNL and $\text{PROMPT}^d\text{-PNL}$ with that of *metric PNL*, that is, the metric extension of PNL introduced and systematically studied in [7].

Acknowledgements. The authors acknowledge the support from the Italian GNCS project *Logics, automata, and games for auto-adaptive systems*. In addition, Dario Della Monica and Aniello Murano acknowledge the support from the POR Campania project *Strategic reasoning for multi-agent systems*.

References

1. Allen, J.F.: Maintaining knowledge about temporal intervals. *Communications of the ACM* 26(11), 832–843 (1983)
2. Allen, J.F.: Towards a general theory of action and time. *Artificial Intelligence* 23(2), 123–154 (1984)
3. Almagor, S., Hirshfeld, Y., Kupferman, O.: Promptness in ω -regular automata. In: ATVA. LNCS, vol. 6252, pp. 22–36. Springer (2010)
4. Alur, R., Etessami, K., La Torre, S., Peled, D.: Parametric temporal logic for “model measuring”. *ACM Trans. Comput. Log.* 2(3), 388–407 (2001), <http://doi.acm.org/10.1145/377978.377990>
5. Aminof, B., Murano, A., Rubin, S., Zuleger, F.: Prompt Alternating-Time Epistemic Logics. In: Baral, C., Delgrande, J.P., Wolter, F. (eds.) *Proc. of the 15th KR*. pp. 258–267. AAAI Press (2016)
6. Bojańczyk, M., Colcombet, T.: Bounds in ω -regularity. In: *LICS*. pp. 285–296. IEEE Computer Society (2006)
7. Bresolin, D., Della Monica, D., Goranko, V., Montanari, A., Sciavicco, G.: Metric propositional neighborhood interval logics on natural numbers. *Software and Systems Modeling (SoSyM)* 12(2), 245–264 (2013)
8. Bresolin, D., Goranko, V., Montanari, A., Sciavicco, G.: Propositional interval neighborhood logics: Expressiveness, decidability, and undecidable extensions. *Ann. Pure Appl. Logic* 161(3), 289–304 (2009), <http://dx.doi.org/10.1016/j.apal.2009.07.003>
9. Chatterjee, K., Henzinger, T.A., Horn, F.: Finitary winning in ω -regular games. *ACM Transactions on Computational Logic* 11(1) (2009)
10. Della Monica, D., Goranko, V., Montanari, A., Sciavicco, G.: Interval temporal logics: a journey. *Bulletin of the European Association for Theoretical Computer Science* 105, 73–99 (2011)
11. Della Monica, D., Montanari, A., Murano, A., Sala, P.: Prompt interval temporal logic (extended version) (2016), available at http://wpage.unina.it/dario.dellamonica/techrep/promptPNL_ext.pdf
12. Della Monica, D., Montanari, A., Sala, P.: The importance of the past in interval temporal logics: the case of Propositional Neighborhood Logic. In: *Logic Programs, Norms and Action - Essays in Honor of Marek J. Sergot on the Occasion of His 60th Birthday*, LNCS, vol. 7360, pp. 79–102. Springer Berlin / Heidelberg (2012)
13. Fijalkow, N., Zimmermann, M.: Cost-Parity and Cost-Streest Games. In: *FSTTCS. LIPIcs*, vol. 18, pp. 124–135 (2012)
14. Gennari, R., Tonelli, S., Vittorini, P.: An AI-based process for generating games from flat stories. In: *Proc. of the 33rd SGAI*. pp. 337–350 (2013)
15. Halpern, J.Y., Shoham, Y.: A propositional modal logic of time intervals. *J. ACM* 38(4), 935–962 (1991), <http://doi.acm.org/10.1145/115234.115351>
16. Kupferman, O., Piterman, N., Vardi, M.Y.: From liveness to promptness. *Formal Methods in System Design* 34(2), 83–103 (2009)
17. Laban, S., El-Desouky, A.: RISMA: A rule-based interval state machine algorithm for alerts generation, performance analysis and monitoring real-time data processing. In: *Proc. of the EGU General Assembly 2013*. *Geophysical Research Abstracts*, vol. 15 (2013)

18. Lodaya, K., Parikh, R., Ramanujam, R., Thiagarajan, P.: A logical study of distributed transition systems. *Information and Computation* 119(1), 91–118 (1995), <http://www.sciencedirect.com/science/article/pii/S0890540185710784>
19. Mogavero, F., Murano, A., Sorrentino, L.: On promptness in parity games. In: Proc. of the 19th LPAR. LNCS, vol. 8312, pp. 601–618. Springer (2013)
20. Moszkowski, B.: Reasoning about digital circuits. Tech. rep. stan-cs-83-970, Dept. of Computer Science, Stanford University, Stanford, CA (1983)
21. Pnueli, A.: The temporal logic of programs. In: Proceedings of the 18th Annual Symposium on Foundations of Computer Science (FOCS). pp. 46–57. IEEE Computer Society (1977)
22. Pratt-Hartmann, I.: Temporal prepositions and their logic. *Artificial Intelligence* 166(1-2), 1–36 (2005)
23. Zhou, C., Hansen, M.R.: Duration Calculus: A formal approach to real-time systems. EATCS Monographs in Theoretical Computer Science, Springer (2004)
24. Zimmermann, M.: Optimal bounds in parametric LTL games. *Theor. Comput. Sci.* 493, 30–45 (2013), <http://dx.doi.org/10.1016/j.tcs.2012.07.039>