

Verifying linear temporal specifications of constant-rate multi-mode systems

Michael Blondin
Université de Sherbrooke
 Sherbrooke, Canada
 michael.blondin@usherbrooke.ca

Philip Offtermatt
Université de Sherbrooke
 Sherbrooke, Canada
University of Warsaw
 Warsaw, Poland
 philip.offtermatt@usherbrooke.ca

Alex Sansfaçon-Buchanan
Université de Sherbrooke
 Sherbrooke, Canada
 alex.sansfacon-buchanan@usherbrooke.ca

Abstract—Constant-rate multi-mode systems (MMS) are hybrid systems with finitely many modes and real-valued variables that evolve over continuous time according to mode-specific constant rates. We introduce a variant of linear temporal logic (LTL) for MMS, and we investigate the complexity of the model-checking problem for syntactic fragments of LTL. We obtain a complexity landscape where each fragment is either P-complete, NP-complete or undecidable. These results generalize and unify several results on MMS and continuous counter systems.

I. INTRODUCTION

Constant-rate multi-mode systems (MMS) are hybrid systems with finitely many modes and a finite number of real-valued variables that evolve over continuous time according to mode-specific constant rates.

MMS were originally introduced by Alur *et al.* to model, *e.g.*, problems related to green scheduling and reducing energy peak consumption of systems [1]. There, they consider the problems of safe schedulability and safe reachability with respect to zones defined as bounded convex polytopes.

Safe schedulability asks whether a given MMS admits a non-Zeno¹ infinite execution that remains within a given safety zone. *Safe reachability* asks whether a given MMS has a finite execution that reaches a target point, while staying within a given safety zone along the way. Both problems were shown to be solvable in polynomial time [1].

A similar problem was studied by Krishna *et al.* in the context of motion planning [2]. There, the authors are interested in the *reach-avoid problem*. In the latter, the goal is to reach a given target point without ever entering any of the given obstacles. The authors of [2] consider obstacles specified as convex polytopes. They show that the reach-avoid problem is decidable if the obstacles are closed, and is undecidable in general. They further provide an implementation of their procedure which is benchmarked positively against the *Open Motion Planning Library*.

M. Blondin was supported by a Discovery Grant from the Natural Sciences and Engineering Research Council of Canada (NSERC), and by the Fonds de recherche du Québec – Nature et technologies (FRQNT). P. Offtermatt is now at Informal Systems, Munich, Germany.

¹Informally, this means that there cannot be infinitely many mode switches within a finite amount of time.

Contribution: The aforementioned problems were solved with ad hoc approaches. Moreover, many natural problems cannot be expressed in these existing frameworks. One such problem is *safe repeated reachability*, where the goal is to find a non-Zeno infinite execution that remains within a safety zone and visits a finite set of zones infinitely often.

We propose a framework that encompasses all of the above. More precisely, we introduce a linear temporal logic (LTL) for MMS. Our variant uses bounded convex polytopes as atomic propositions. We omit the next operator X which is ill-suited for the continuous behavior of MMS. Moreover, we use a strict-future interpretation of the until temporal operator U , inspired from metric temporal logic [3] (more precisely from MITL_{0,∞}). In particular, our logic can express

- Safe schedulability: $\mathbf{G} Z_{\text{safe}}$;
- Safe reachability: $Z_{\text{safe}} \mathbf{U} \{x_{\text{target}}\}$;
- Reach-avoid: $(\neg O_1 \wedge \dots \wedge \neg O_n) \mathbf{U} \{x_{\text{target}}\}$; and
- Safe repeated reachability: $(\mathbf{G} Z_{\text{safe}}) \wedge \bigwedge_{i=1}^n (\mathbf{GF} Z_i)$.

We investigate the computational complexity of LTL model checking, which asks, given an MMS M , a starting point x and an LTL formula φ , whether there is a non-Zeno infinite execution of M that satisfies φ from x , denoted $x \models_M \varphi$. We consider the syntactic fragments obtained by (dis)allowing operators from $\{U, F, G, \wedge, \vee, \neg\}$ and allowing at least one temporal operator². We establish the computational complexity of *all* of the $2^6 - 2^3 = 56$ fragments: Each one is either P-complete, NP-complete or undecidable.

Our work is also closely related to the study of counter systems like vector addition systems (VAS) and Petri nets. These models have countless applications ranging from program verification and synthesis, to the formal analysis of chemical, biological and business processes (*e.g.*, see [4]–[7]). Moreover, the continuous relaxation of counter systems has been successfully employed in practice to alleviate their tremendous computational complexity (*e.g.*, see [8], [9]).

The behavior of an MMS amounts to continuous pseudo-reachability of VAS and Petri nets, *i.e.* where the effect of transitions can be scaled by positive real values, and without

²Without any temporal operator, the logic has nothing to do with MMS; it becomes quantifier-free linear arithmetic.

the requirement that counters must remain non-negative. The latter requirement can be regained in our logic. While we do not investigate unbounded zones in their full generality, we consider semi-bounded linear formulas, which include formulas of the form $(GZ) \wedge \dots$ or $Z U \dots$, where Z is unbounded, and so can be set to $Z := \mathbb{R}_{\geq 0}^d$. In particular, our results imply the known fact that continuous reachability, *i.e.* checking $\mathbb{R}_{\geq 0}^d U \{x_{\text{target}}\}$, can be done in polynomial time [10]. Moreover, we establish the decidability of richer properties. Thus, our work can be seen as a unifying and more general framework for MMS and continuous VAS/Petri nets.

Results: Let us write $LTL_B(X)$ to denote the set of LTL formulas using only operators from X , and $LTL(X)$ for the same fragment but with zones possibly unbounded. We obtain the full complexity landscape depicted in Figure 1. Our contribution is summarized by the following three points.

I): We show that $LTL_B(\{F, G, \wedge\})$ is in NP, and hence that $LTL_B(\{F, \wedge, \vee\})$ is as well. More precisely, we prove that:

- 1) Formulas from this fragment can be put in a normal form, coined as *flat formulas*, where the nesting of temporal operators is restricted;
- 2) Flat formulas can be translated into generalized Büchi automata with transition-based acceptance, no cycles except for self-loops (“almost acyclic”) and linear width;
- 3) Testing whether an MMS M satisfies a specification given by such an automaton \mathcal{A} can be done in NP by guessing a so-called linear path scheme S of \mathcal{A} ; constructing a so-called linear formula ψ equivalent to S , and testing whether $x \models_M \psi$ in polynomial time.

Step 2 is inspired by the work of Křetínský and Esparza [11] on deterministic Muller automata for classical LTL restricted to $\{F, G, \wedge, \vee, \neg\}$. Our construction also deals with classical LTL, restricted to $\{F, G, \wedge\}$, and is thus an indirect contribution to logic and automata independent of MMS.

In particular, in Step 3 we establish a polynomial-time LTL fragment for MMS, namely *semi-bounded linear LTL formulas*. We do so by using a polynomial-time fragment of existential linear arithmetic, introduced in [12] for the purpose of characterizing reachability sets of continuous Petri nets. In particular, we show how to translate LTL formulas of the form $\psi = (GZ_0) \wedge \bigwedge_{i=1}^n GFZ_i$, with Z_0 unbounded, into the logic of [12]. This is challenging, in contrast to simply handling GZ_0 , with Z_0 bounded, as done in [1]. It involves a technical characterization of MMS and points that satisfy ψ , which, in particular, goes through a careful use of Farkas’ lemma.

As a corollary of Step 3, we show that $LTL_B(\{F, G, \neg\})$, $LTL(\{F, \vee\})$ and $LTL(\{G, \wedge\})$ are solvable in polynomial time. These fragments include safe schedulability and safe reachability, which generalizes their membership in P [1].

II): We show the NP-hardness of $LTL_B(\{F, \wedge\})$ by reducing from SUBSET-SUM. With the previous results, this shows that $LTL_B(\{F, G, \wedge\})$ and $LTL_B(\{F, \wedge, \vee\})$ are NP-complete.

III): We show that $LTL_B(\{U\})$ and $LTL_B(\{G, \vee\})$ are both undecidable, by reducing from the reachability problem for Petri nets with inhibitor arcs. This “generalizes” the undecidability of the reach-avoid problem established in [2]. Their

proof indirectly shows that the model checking problem is undecidable for formulas of the form $(Z_1 \vee \dots \vee Z_n) U \{x_{\text{target}}\}$ where each Z_i is a possibly unbounded zone. We strengthen this result by using bounded zones only.

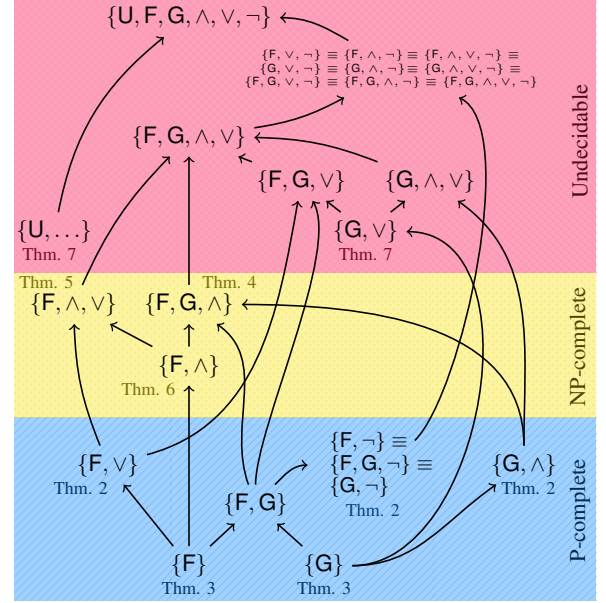


Fig. 1: Complexity landscape of LTL model checking for MMS. An edge from X to Y indicates that any formula from $LTL_B(X)$ is equivalent to some formula from $LTL_B(Y)$. Each expression “ $X \equiv Y$ ” stands for $X \leftrightarrow Y$, *i.e.*, an edge from X to Y and an edge from Y to X . Node $\{U, \dots\}$ stands for any LTL fragment that contains U .

Further related work: MMS are related to hybrid automata [13]. Contrary to MMS, however, the latter allow for a finite control structure, and modes of non-constant rates. Their immense modelling power leads to the undecidability of most problems, including reachability, *i.e.* formulas of the form $F \{x_{\text{target}}\}$. Yet, some researchers have investigated decision procedures for temporal specification languages such as signal temporal logic (*e.g.*, see [14]).

Timed automata [15] form another related type of hybrid system. In this model, all variables (known as clocks) increase at the *same* constant rate, as opposed to the case of MMS. On the other hand, timed automata are equipped with a finite control structure, which is not the case of MMS.

Bounded-rate multi-mode systems generalize MMS [16], [17]. In this model, the mode-dependent rates are given as bounded convex polytopes. The setting can be seen as a two-player game. Player 1 chooses a mode and a duration, and Player 2 chooses the rates from the set for that mode. The system evolves according to the rates chosen by Player 2. In this context, “schedulability” is a strategy for Player 1 that never leaves the safety zone, no matter the choices of Player 2.

Small fragments of classical LTL have been investigated in the literature, *e.g.* see [18, Table 1]. In particular, $LTL(\{F, G, \wedge, \vee, \neg\})$ has been studied in [11], [19] under the names

$L(F)$ and (F, G) , and $LTL(\{F, \wedge\})$ has been studied in [20] under the name $LTL_+(\diamond, \wedge)$. The authors of [20] show that a fragment, called LTL^{PODB} , and which is incomparable to $LTL(\{F, G, \wedge\})$, admits partially-ordered deterministic Büchi automata of exponential size and linear width. To the best of our knowledge, there is no work dedicated to $LTL(\{F, G, \wedge\})$, and in particular to its translation into automata of linear width.

Organization: In Section II, we introduce basic definitions, MMS and LTL. We further relate LTL over MMS with classical LTL over infinite words. In Section III, we show that any formula from classical $LTL(\{F, G, \wedge\})$ translates into a specific type of ω -automaton, which amounts to a disjunction of so-called linear LTL formulas. In Section IV, we show that linear LTL formulas over MMS can be model-checked in polynomial time. From this, we establish the P-completeness of some syntactic fragments. In Section V and Section VI, we respectively prove the NP-completeness and undecidability of the other fragments. We conclude in Section VII. Due to space limitation, many proofs are deferred to the full version which is freely available on arXiv.

II. PRELIMINARIES

We write \mathbb{N} to denote $\{0, 1, \dots\}$, \mathbb{Z} to denote the integers, and \mathbb{R} to denote the reals. We use subscripts to restrict these sets, e.g. $\mathbb{R}_{>0} := \{x \in \mathbb{R} : x > 0\}$. We write $[\alpha, \beta] := \{x \in \mathbb{R} : \alpha \leq x \leq \beta\}$ and $[a..b] := \{i \in \mathbb{N} : a \leq i \leq b\}$. We also use (semi-)open intervals, e.g. $(1, 2] = [1, 2] \setminus \{1\}$.

Let I be a set of indices and let $X \subseteq \mathbb{R}^I$. We write $e_i \in \mathbb{R}^I$ for the vector with $e_i(i) = 1$ and $e_i(j) = 0$ for all $j \neq i$, and $\mathbf{0}$ to denote the vector such that $\mathbf{0}(i) = 0$ for all $i \in I$. Let $\|\mathbf{x}\| := \max\{|\mathbf{x}(i)| : i \in I\}$ and $\|X\| := \sup\{\|\mathbf{x}\| : \mathbf{x} \in X\}$. We say that X is *convex* if $\lambda\mathbf{x} + (1-\lambda)\mathbf{y} \in X$ for all $\lambda \in [0, 1]$ and $\mathbf{x}, \mathbf{y} \in X$, and *bounded* if $\|X\| \leq b$ for some $b \in \mathbb{R}_{\geq 0}$.

We write 2^Σ for the powerset of Σ . Given a nonempty finite sequence w , let $w^\omega := ww \dots$. Let $\Sigma^\omega := \{w_0 w_1 \dots : w_i \in \Sigma\}$ be the set of infinite sequences with elements from Σ .

A. Constant-rate multi-mode systems

A d -dimensional constant-rate multi-mode system (MMS), with $d \in \mathbb{N}_{\geq 1}$, is a finite set $M \subseteq \mathbb{R}^d$ whose elements are called *modes*. A *schedule* is a (finite or infinite) sequence $\pi = (\alpha_1, \mathbf{m}_1)(\alpha_2, \mathbf{m}_2) \dots$, where each $(\alpha_i, \mathbf{m}_i) \in \mathbb{R}_{>0} \times M$. To ease the notation, we often write, e.g., $\mathbf{m} \frac{1}{2} \mathbf{m}'$ rather than $(1, \mathbf{m})(1/2, \mathbf{m}')$. Given $\lambda \in \mathbb{R}_{>0}$, we define the schedule $\lambda\pi$ as π with each α_i replaced by $\lambda\alpha_i$. The *size* of π , denoted $|\pi|$, is its number of pairs. The *effect* of π is $\Delta_\pi := \sum_i \alpha_i \mathbf{m}_i$. The *support* of π is $\text{supp}(\pi) := \{\mathbf{m}_1, \mathbf{m}_2, \dots\}$. Let $\text{time}_{\mathbf{m}}(\pi) := \sum_{i: \mathbf{m}_i = \mathbf{m}} \alpha_i$ and $\text{time}(\pi) := \sum_{\mathbf{m} \in M} \text{time}_{\mathbf{m}}(\pi)$. We say that an infinite schedule π is *non-Zeno* if $\text{time}(\pi) = \infty$. The *Parikh image* of a finite schedule π is denoted $\boldsymbol{\pi} \in \mathbb{R}_{\geq 0}^M$, i.e. $\boldsymbol{\pi}(\mathbf{m}) := \text{time}_{\mathbf{m}}(\pi)$. We say that two finite schedules are *equivalent*, denoted with \equiv , if they are equal after merging consecutive equal modes, i.e. using the rule $\pi(\alpha, \mathbf{m})(\beta, \mathbf{m})\pi' \equiv \pi(\alpha + \beta, \mathbf{m})\pi'$. Let $\pi[\tau.. \tau']$ be the schedule obtained from π starting where time τ has elapsed, and ending where time τ' has elapsed; e.g., for $\pi = (2, \mathbf{m}_1)(0.5, \mathbf{m}_2)(1, \mathbf{m}_3)^\omega$, we have




Fig. 2: An execution σ is depicted as a directed path. Three bounded zones X, Y, Z are depicted as filled colored polygons. A trace is obtained from σ from the marked points.

$\pi[0..1] = (1, \mathbf{m}_1)$, $\pi[0.5..2.25] = (1.5, \mathbf{m}_1)(0.25, \mathbf{m}_2)$ and $\pi[3..] = (0.5, \mathbf{m}_3)(1, \mathbf{m}_3)^\omega$.

An *execution* is a (finite or infinite) sequence $\sigma = x_0 I_0 x_1 I_1 x_2 \dots$ where $x_0, x_1, \dots \in \mathbb{R}^d$, $I_0, I_1, \dots \subseteq \mathbb{R}_{\geq 0}$ are closed intervals with distinct endpoints, $\min I_0 = 0$, and $\min I_j = \max I_{j-1}$ for all $j \in \mathbb{N}_{>0}$. Let $\text{dom } \sigma := I_0 \cup I_1 \cup \dots$. For every $\tau \in \text{dom } \sigma$, with $\tau \in I_j$, let

$$\sigma(\tau) := \mathbf{x}_j + \frac{\tau - \min I_j}{\max I_j - \min I_j} \cdot (\mathbf{x}_{j+1} - \mathbf{x}_j).$$

We define $\sigma[\tau.. \tau']$, where $[\tau, \tau'] \subseteq \text{dom } \sigma$, as the execution σ' that satisfies $\sigma'(\alpha) = \sigma(\tau + \alpha)$ for every $\alpha \in [0, \tau' - \tau]$.

A *schedule* $\pi = (\alpha_1, \mathbf{m}_1)(\alpha_2, \mathbf{m}_2) \dots$, together with a point \mathbf{x}_0 , gives rise to an execution $\text{exec}(\pi, \mathbf{x}_0) := \mathbf{x}_0 I_0 \mathbf{x}_1 \dots$ where $I_0 := [0, \alpha_1]$, $I_j := [\max I_{j-1}, \max I_{j-1} + \alpha_{j+1}]$ and $\mathbf{x}_j := \mathbf{x}_{j-1} + \alpha_j \mathbf{m}_j$. We use the notation $\mathbf{x} \rightarrow^\pi \mathbf{y}$ to denote the fact that π is a schedule that, from \mathbf{x} , gives rise to an execution ending in \mathbf{y} . If we only care about the existence of such a schedule, we may write $\mathbf{x} \rightarrow^* \mathbf{y}$, or write $\mathbf{x} \rightarrow^+ \mathbf{y}$ to denote that there is such a nonempty schedule. We sometimes omit either of the two endpoints if its value is irrelevant, e.g. $\mathbf{x} \rightarrow^\pi$ stands for $\mathbf{x} \rightarrow^\pi \mathbf{x} + \Delta_\pi$. Given a set $Z \subseteq \mathbb{R}^d$, we write $\mathbf{x} \rightarrow_Z^\pi$ to denote that the execution never leaves Z , i.e. $\sigma(\tau) \in Z$ for all $\tau \in \text{dom } \sigma$, where $\sigma := \text{exec}(\pi, \mathbf{x})$. We extend this notation to any set of sets \mathcal{X} , requiring that, for all $\tau \in \text{dom } \sigma$, there exists $Z \in \mathcal{X}$ such that $\sigma(\tau) \in Z$.

Example 1: Let $M := \{(0, 1), (1, 0), (1, 1), (-1, 1)\}$. Let $\pi := \frac{1}{2}(1, 1) \frac{1}{2}(1, 0) \frac{1}{2}(1, 1) \frac{1}{2}(1, 0) (-1, 1) (0, 1) \dots$ be a schedule. The execution $\sigma := \text{exec}(\pi, (1, 1))$ is depicted in Figure 2 as a directed path, with distinct styles to distinguish the modes (ignore the circular marks and colored polygons):

$$(1, 1) [0, 0.5] (1.5, 1.5) [0.5, 1] (2, 1.5) [1, 1.5] (2.5, 2) [1.5, 2] (3, 2) [2, 3] (2, 3) [3, 4] (2, 4) \dots \quad \square$$

B. A linear temporal logic for MMS

A *zone* $Z \subseteq \mathbb{R}^d$ is a convex polytope represented as the intersection of finitely many closed half-spaces, i.e. $Z = \{\mathbf{x} \in \mathbb{R}^d : \mathbf{A}\mathbf{x} \leq \mathbf{b}\}$ for some $\mathbf{A} \in \mathbb{Z}^{k \times d}$ and $\mathbf{b} \in \mathbb{Z}^k$.

Linear temporal logic (LTL), over a finite set of zones AP , has the following syntax:

$$\varphi ::= \text{true} \mid Z \mid \neg\varphi \mid \varphi \wedge \varphi' \mid \varphi \vee \varphi' \mid \mathbf{F}\varphi \mid \mathbf{G}\varphi \mid \varphi \mathbf{U} \varphi',$$

where $Z \in AP$. Let $\text{LTL}(X)$ denote the set of LTL formulas (syntactically) using only the operators from X . For example, $\text{LTL}(\{\mathbf{F}, \mathbf{G}, \wedge\})$ describes the LTL formulas using only operators \mathbf{F} , \mathbf{G} and \wedge . We write $\text{LTL}_{\mathbf{B}}$ to indicate that all zones from AP must be bounded. We say that an LTL formula is *negation-free* if it contains no occurrence of \neg .

We define the semantics over infinite executions:

$$\begin{aligned} \sigma, \tau \models \text{true} &\iff \text{true}, \\ \sigma, \tau \models Z &\iff \sigma(\tau) \in Z, \\ \sigma, \tau \models \neg\varphi &\iff \neg(\sigma, \tau \models \varphi), \\ \sigma, \tau \models \varphi \wedge \varphi' &\iff (\sigma, \tau \models \varphi) \wedge (\sigma, \tau \models \varphi'), \\ \sigma, \tau \models \varphi \vee \varphi' &\iff (\sigma, \tau \models \varphi) \vee (\sigma, \tau \models \varphi'), \\ \sigma, \tau \models \mathbf{F}\varphi &\iff \exists \tau' \geq \tau : \sigma, \tau' \models \varphi, \\ \sigma, \tau \models \mathbf{G}\varphi &\iff \forall \tau' \geq \tau : \sigma, \tau' \models \varphi, \\ \sigma, \tau \models \varphi \mathbf{U} \varphi' &\iff \exists \tau' \geq \tau : (\sigma, \tau' \models \varphi') \\ &\quad \wedge (\forall \tau'' \in [\tau, \tau') : \sigma, \tau'' \models \varphi). \end{aligned}$$

We write $\sigma \models \varphi$ iff $\sigma, 0 \models \varphi$. We say that two formulas are *equivalent*, denoted $\varphi \equiv \varphi'$, if they are satisfied by the same executions. In particular, $\mathbf{F}\psi \equiv \text{true} \mathbf{U} \psi$ and $\mathbf{G}\psi \equiv \neg\mathbf{F}\neg\psi$.

Let M be an MMS and let $\mathbf{x} \in \mathbb{R}^d$. We say that $\mathbf{x} \models_M \varphi$ iff M has a non-Zeno infinite schedule π such that $\text{exec}(\pi, \mathbf{x}) \models \varphi$. The *model-checking problem* of a fragment $\text{LTL}(X)$ asks, given M , \mathbf{x} and $\varphi \in \text{LTL}(X)$, whether $\mathbf{x} \models_M \varphi$.

Example 2: Recall the MMS M and the schedule π from Example 1. Let X , Y and Z be the bounded zones colored in Figure 2, e.g. $X := \{(x, y) \in \mathbb{R}^2 : 0.5 \leq x \leq 1.5, 0.5 \leq y \leq 1.5\}$. We have $(1, 1) \models_M X \wedge \mathbf{F}((Y \wedge \neg Z) \wedge \mathbf{F}Z)$. \square

C. Connection with classical LTL

Classical linear temporal logic (LTL) (without temporal operator \mathbf{X}) has the same syntax as the logic from Section II-B, but is interpreted over infinite words $w \in (2^{AP})^\omega$:

$$\begin{aligned} w, i \models \text{true}, &\iff \text{true}, \\ w, i \models a &\iff a \in w(i), \\ w, i \models \neg\varphi &\iff \neg(w, i \models \varphi), \\ w, i \models \varphi \wedge \varphi' &\iff (w, i \models \varphi) \wedge (w, i \models \varphi'), \\ w, i \models \varphi \vee \varphi' &\iff (w, i \models \varphi) \vee (w, i \models \varphi'), \\ w, i \models \mathbf{F}\varphi &\iff \exists j \geq i : w, j \models \varphi, \\ w, i \models \mathbf{G}\varphi &\iff \forall j \geq i : w, j \models \varphi, \\ w, i \models \varphi \mathbf{U} \varphi' &\iff \exists j \geq i : (w, j \models \varphi') \\ &\quad \wedge (\forall k \in [i..j-1] : w, k \models \varphi). \end{aligned}$$

We write $w \models \varphi$ iff $w, 0 \models \varphi$. Observe that $w, i \models \varphi$ holds iff $w[i..] \models \varphi$, where $w[i..] := w(i)w(i+1)\dots$. We write $\varphi \equiv \varphi'$ if φ and φ' are satisfied by the same infinite words. In order to relate LTL over executions with LTL over infinite

words, we introduce the notion of traces. Informally, a trace captures the zone changes within an execution.

Let $\chi_{AP} : \mathbb{R}^d \rightarrow AP$ be the function that yields the set of zones a given point lies in: $\chi_{AP}(\mathbf{x}) := \{Z \in AP : \mathbf{x} \in Z\}$. Let σ be an execution. We say that word w is a *trace* of σ if there exist $\tau_0 < \tau_1 < \dots \in \mathbb{R}_{\geq 0}$ such that

- $\text{dom } \sigma = [\tau_0, \tau_1] \cup [\tau_1, \tau_2] \cup \dots$,
- $w(i) = \chi_{AP}(\sigma(\tau_i))$ for every $i \in \mathbb{N}$, and
- for every $i \in \mathbb{N}$, there exists $j \in \{i, i+1\}$ such that: $\chi_{AP}(\sigma(\tau')) = \chi_{AP}(\sigma(\tau_j))$ for all $\tau' \in (\tau_i, \tau_{i+1})$.

Example 3: Recall execution σ from Example 1. The word $w := \{X\}\{X\}\emptyset\{Y\}\{Y\}\{Y, Z\}\{Y, Z\}\{Z\}\emptyset\emptyset\dots$ is a trace of σ . As depicted with circular marks in Figure 2, it is obtained from $\tau_0 := 0, \tau_1 := 0.5, \tau_2 := 1, \tau_3 := 1.5, \tau_4 := 2, \tau_5 := 2.5, \tau_6 := 2.75, \tau_7 := 3, \tau_8 := 3.5$ and so on. \square

With a bit of care, it is possible to prove that any execution admits a trace. Moreover, in the absence of negations, model-checking an execution amounts to model-checking any of its traces under the classical LTL semantics. Thus, equivalences of negation-free classical LTL also hold under the semantics over executions.

Proposition 1: Any execution σ has a trace.

Proposition 2: Let σ be an execution with $\text{dom } \sigma = \mathbb{R}_{\geq 0}$, let w be a trace of σ , and let φ be a negation-free LTL formula. It is the case that $\sigma \models \varphi$ iff $w \models \varphi$.

III. FROM $\text{LTL}(\{\mathbf{F}, \mathbf{G}, \wedge\})$ TO LINEAR LTL

This section deals with classical LTL formulas interpreted over infinite words. We show that any formula from $\text{LTL}(\{\mathbf{F}, \mathbf{G}, \wedge\})$ corresponds to an automaton of a certain shape, which amounts to what we call linear LTL formulas.

A. From $\text{LTL}(\{\mathbf{F}, \mathbf{G}, \wedge\})$ to flat formulas

We say that a formula is *pseudo-atomic* if it is a conjunction of atomic propositions. By convention, an empty conjunction amounts to *true*. An LTL formula φ is *flat* if it has this form:

$$\psi \wedge \mathbf{G}\psi' \wedge \bigwedge_{i \in I} \mathbf{G}\mathbf{F}\psi''_i \wedge \bigwedge_{j \in J} \mathbf{F}\varphi_j,$$

where ψ, ψ' and ψ''_i are pseudo-atomic; and φ_j is flat.

Given a formula $\varphi \in \text{LTL}(\{\mathbf{F}, \mathbf{G}, \wedge\})$ of this form:

$$\varphi = \psi \wedge \bigwedge_{i \in I} \mathbf{G}\varphi_i \wedge \bigwedge_{j \in J} \mathbf{F}\varphi_j,$$

we define these mappings:

$$\begin{aligned} \text{flat}_{\mathbf{G}}(\varphi) &:= \mathbf{G}\psi \wedge \bigwedge_{i \in I} \text{flat}_{\mathbf{G}}(\varphi_i) \wedge \bigwedge_{j \in J} \text{flat}_{\mathbf{GF}}(\varphi_j), \\ \text{flat}_{\mathbf{GF}}(\varphi) &:= \mathbf{GF}\psi \wedge \bigwedge_{i \in I} \text{flat}_{\mathbf{FG}}(\varphi_i) \wedge \bigwedge_{j \in J} \text{flat}_{\mathbf{GF}}(\varphi_j), \\ \text{flat}_{\mathbf{FG}}(\varphi) &:= \mathbf{FG}\psi \wedge \bigwedge_{i \in I} \text{flat}_{\mathbf{FG}}(\varphi_i) \wedge \bigwedge_{j \in J} \text{flat}_{\mathbf{GF}}(\varphi_j), \\ \text{flat}(\varphi) &:= \psi \wedge \bigwedge_{i \in I} \text{flat}_{\mathbf{G}}(\varphi_i) \wedge \bigwedge_{j \in J} \mathbf{F}\text{flat}(\varphi_j). \end{aligned}$$

As its name suggests, it follows by induction that formula $\text{flat}(\varphi)$ is flat. Moreover, the following holds.

Proposition 3: It is the case that $\text{flat}(\varphi) \equiv \varphi$.

B. From flat formulas to almost acyclic automata

We say that an automaton is *almost acyclic* if, for every pair of states $q \neq r$, it is the case that $q \rightarrow^* r$ implies $r \not\rightarrow^* q$, i.e. cycles must be self-loops. The *width* of an almost acyclic automaton is the maximal length among its simple paths.

We will prove that any formula $\varphi \in \text{LTL}(\{F, G, \wedge\})$ can be translated into an almost acyclic automaton \mathcal{A}_φ of linear width in the size of φ , and such that \mathcal{A}_φ accepts w iff $w \models \varphi$. We will formally define the acceptance condition later on, but for readers familiar with ω -automata: \mathcal{A}_φ will be a generalized Büchi automaton with accepting transitions.

In order to define \mathcal{A}_φ , we first provide intermediate definitions. Let $\mathfrak{U}: \text{LTL}(\{F, G, \wedge\}) \rightarrow 2^{\text{LTL}(\{F, G, \wedge\})}$ be defined by $\mathfrak{U}(\text{true}) := \{\text{true}\}$, $\mathfrak{U}(a) := \{a\}$, $\mathfrak{U}(G\varphi) := \{G\varphi\}$,

$$\begin{aligned} \mathfrak{U}(\varphi_1 \wedge \varphi_2) &:= \{\psi_1 \wedge \psi_2 : \psi_1 \in \mathfrak{U}(\varphi_1), \psi_2 \in \mathfrak{U}(\varphi_2)\}, \\ \mathfrak{U}(F\varphi) &:= \{F\varphi\} \cup \mathfrak{U}(\varphi). \end{aligned}$$

Example 4: The set $\mathfrak{U}(Ga \wedge F(b \wedge FGc))$ is equal to

$$\{Ga \wedge F(b \wedge FGc), Ga \wedge b \wedge FGc, Ga \wedge b \wedge Gc\}. \quad \square$$

Given $A \subseteq AP$, let $\text{prop}(\bigwedge_{a \in A} a) := A$. Given $A \subseteq AP$ and a flat formula $\varphi = \psi \wedge G\psi' \wedge \bigwedge_{i \in I} GF\psi''_i \wedge \bigwedge_{j \in J} F\varphi_j$, let

$$\varphi[A] := \begin{cases} G\psi' \wedge \bigwedge_{i \in I} GF\psi''_i \wedge \bigwedge_{j \in J} F\varphi_j & \text{if } \text{prop}(\psi \wedge \psi') \subseteq A, \\ \text{false} & \text{otherwise.} \end{cases}$$

Given $\varphi \in \text{LTL}(\{F, G, \wedge\})$, the automaton $\mathcal{A}_\varphi := (Q, \Sigma, \rightarrow, q_0)$ is defined respectively by the following states, alphabet, transitions and initial state:

$$\begin{aligned} Q &:= \{\psi \in \text{LTL}(\{F, G, \wedge\}) : \psi \text{ is flat}\}, \\ \Sigma &:= 2^{AP}, \\ \rightarrow &:= \{(\psi, A, \psi') : \exists \psi'' \in \mathfrak{U}(\psi) \text{ s.t. } \psi' = \psi''[A] \neq \text{false}\}, \\ q_0 &:= \text{flat}(\varphi). \end{aligned}$$

Example 5: Let $\varphi := a \wedge Fb$, which is flat. We have $\mathfrak{U}(\varphi) = \{a \wedge Fb, a \wedge b\}$. The automaton \mathcal{A}_φ is depicted at the top of Figure 3. Note that $w \models \varphi$ iff there is an infinite path from the initial state, labeled with w , that visits true .

Let $\varphi' := GF(a \wedge Gc) \wedge Fb$. We have $\text{flat}(\varphi') = GFa \wedge FGc \wedge Fb$. Hence, $\mathfrak{U}(\text{flat}(\varphi'))$ is equal to

$$\{GFa \wedge FGc \wedge Fb, GFa \wedge FGc \wedge b, GFa \wedge Gc \wedge Fb, GFa \wedge Gc \wedge b\}.$$

The automaton $\mathcal{A}_{\varphi'}$ is depicted at the bottom of Figure 3. Let $q := GFa \wedge Gc$. Note that $w \models \varphi'$ iff there is an infinite path from the initial state, labeled with w , that visits the set of transitions $\{(q, A, q) : A \supseteq \{a, c\}\}$ infinitely often. \square

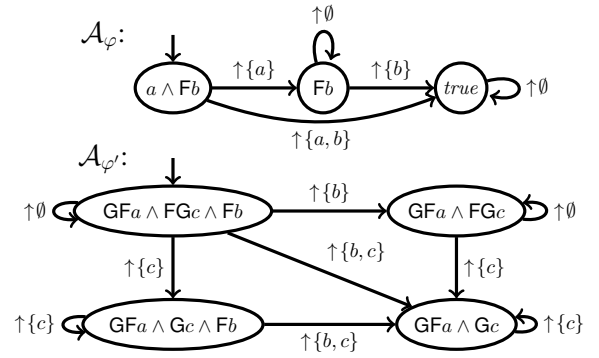


Fig. 3: Automata for $\varphi = a \wedge Fb$ (top) and $\varphi' = GF(a \wedge Gc) \wedge Fb$ (bottom). Each $\uparrow A$ stands for $\{A' \subseteq AP : A' \supseteq A\}$.

1) *Shape of automaton \mathcal{A}_φ :* We first seek to prove that \mathcal{A}_φ is almost acyclic. For every $\varphi \in \text{LTL}(\{F, G, \wedge\})$, let $|\text{true}| = |a| := 1$, $|\varphi_1 \wedge \varphi_2| := |\varphi_1| + 1 + |\varphi_2|$ and $|G\varphi| = |F\varphi| := 1 + |\varphi|$. Moreover, let $|\text{true}|_F = |a|_F = |G\varphi|_F := 0$, $|\varphi_1 \wedge \varphi_2|_F := |\varphi_1|_F + |\varphi_2|_F$ and $|F\varphi|_F := 1 + |\varphi|_F$. The properties below follow by induction.

Proposition 4: Let $\varphi \in \text{LTL}(\{F, G, \wedge\})$. This holds:

- 1) if φ is flat, then $|\varphi|_F > |\varphi'|_F$ for all $\varphi' \in \mathfrak{U}(\varphi) \setminus \{\varphi\}$,
- 2) if φ is flat, then $|\varphi|_F \geq |\varphi[A]|_F$ for all $A \subseteq AP$,
- 3) $|\varphi| \geq |\text{flat}(\varphi)|_F$.

This proposition follows from Proposition 4:

Proposition 5: Let $r_0 \xrightarrow{A_1} r_1 \xrightarrow{A_2} \dots \xrightarrow{A_n} r_n$ be a simple path of \mathcal{A}_φ . It is the case that $|r_1|_F > \dots > |r_n|_F$.

Proposition 6: Let $\varphi \in \text{LTL}(\{F, G, \wedge\})$. Automaton \mathcal{A}_φ is almost acyclic and its width belongs to $\mathcal{O}(|\varphi|)$.

Proof: Let us first prove almost acyclicity. For the sake of contradiction, suppose that \mathcal{A}_φ has a simple cycle $q \rightarrow^u r \rightarrow^v q$ where $q \neq r$. Since $q \rightarrow^* r$, it follows from Items 1 and 2 of Proposition 4 that $|q|_F \geq |r|_F$. Since $q \neq r$, we have $|u| > 1$ and $|v| > 1$. Thus, Proposition 5 yields $|r|_F > |q|_F$, which is a contradiction.

Let us now bound the width n of \mathcal{A}_φ . Let $q_0 \xrightarrow{A_1} q_1 \xrightarrow{A_2} \dots \xrightarrow{A_n} q_n$ be a simple path of \mathcal{A}_φ . We have

$$\begin{aligned} n &\leq |q_1|_F + 1 && \text{(by Proposition 5)} \\ &\leq |q_0|_F + 1 && \text{(by Items 1 and 2 of Proposition 4)} \\ &= |\text{flat}(\varphi)|_F + 1 && \text{(by def. of } q_0) \\ &\leq |\varphi| + 1 && \text{(by Item 3 of Proposition 4)}. \quad \square \end{aligned}$$

2) *Language of \mathcal{A}_φ :* Let us define the acceptance condition of automaton \mathcal{A}_φ . Let $F := \{q \in Q : q_0 \rightarrow^+ q \wedge |q|_F = 0\}$. By definition, each state $q \in F$ is of the form $G\psi \wedge \bigwedge_{j \in J} GF\psi'_j$. Given such a state q , we define

$$T_{q,j} := \{(q, A, q) \in \rightarrow : \text{prop}(\psi'_j) \subseteq A\}.$$

We say that word $w \in (2^{AP})^\omega$ is *accepted* by \mathcal{A}_φ , denoted $w \in L(\mathcal{A}_\varphi)$, iff there exist $q \in F$ and an infinite path from q_0 that visits q and, for each $j \in J$, the set $T_{q,j}$ infinitely often. In the remainder, we prove that $w \in L(\mathcal{A}_\varphi)$ iff $w \models \varphi$.

Lemma 1: Let $\varphi \in \text{LTL}(\{F, G, \wedge\})$ be a flat formula. These two properties are equivalent to $w \models \varphi$:

- 1) there exists φ' such that $\varphi \rightarrow^{w(0)} \varphi'$ and $w[1..] \models \varphi'$;
- 2) there exist $i \in \mathbb{N}$ and φ' such that $\varphi \rightarrow^{w(0)\dots w(i-1)} \varphi'$, $|\varphi'|_F = 0$ and $w[i..] \models \varphi'$.

Proposition 7: Let $\varphi \in \text{LTL}(\{F, G, \wedge\})$. It is the case that $w \models \varphi$ iff $w \in L(\mathcal{A}_\varphi)$.

Proof: \Rightarrow) By Lemma 1(2), there are $k \in \mathbb{N}$ and φ' with $\text{flat}(\varphi) \rightarrow^{w(0)\dots w(k-1)} \varphi'$, $|\varphi'|_F = 0$ and $w[k..] \models \varphi'$.

As $|\varphi'|_F = 0$, we have $\mathcal{U}(\varphi') = \{\varphi'\}$. So, Lemma 1(1) yields

$$\varphi' \rightarrow^{w(k)} \varphi' \rightarrow^{w(k+1)} \varphi' \rightarrow^{w(k+2)} \dots$$

As $\varphi' \in F$, it has the form $G\psi \wedge \bigwedge_{j \in J} GF\psi'_j$. In particular, this means that $w[k..] \models \bigwedge_{j \in J} GF\psi'_j$. Recall that $T_{\varphi',j} = \{(\varphi', A, \varphi') \in \rightarrow : \text{prop}(\psi'_j) \subseteq A\}$. So, for each $j \in J$, the set $T_{\varphi',j}$ is visited infinitely often.

\Leftarrow) By $w \in L(\mathcal{A}_\varphi)$, there exist $q \in F$ and $k \in \mathbb{N}$ such that

- $q = G\psi \wedge \bigwedge_{j \in J} GF\psi'_j$,
- $q_0 \rightarrow^{w(0)\dots w(k-1)} q$, and
- some infinite path $q \rightarrow^{w[k..]}$ visits, for each $j \in J$, the set $T_{q,j}$ infinitely often.

Recall that $T_{q,j} = \{(q, A, q) \in \rightarrow : \text{prop}(\psi'_j) \subseteq A\}$. Since $q \rightarrow^{w[k..]}$ visits each $T_{q,j}$ infinitely often, we have $w[k..] \models \bigwedge_{j \in J} GF\psi'_j$. By $\mathcal{U}(q) = \{q\}$ and by definition of \rightarrow , we have $w[k..] \models G\psi$. So, $w[k..] \models q$. By repeated applications of Lemma 1(1), this implies $w \models q_0 = \text{flat}(\varphi) \equiv \varphi$. \square

C. From almost acyclic automata to linear LTL

In this subsection, we show that almost acyclic automata are equivalent to finite sets of so-called linear LTL formulas, with the goal of showing that $\text{LTL}_B(\{F, G, \wedge\})$ belongs to NP in the forthcoming Section V-A.

For every $A \subseteq AP$, let $\uparrow A := \{A' \subseteq AP : A' \supseteq A\}$. We say that $X \subseteq 2^{AP}$ is *simple* if $X = \uparrow A$ for some $A \subseteq AP$.

Example 6: Consider the bottom automaton of Figure 3. Its infinite paths are captured by these three expressions:

- $\uparrow \emptyset^* \uparrow \{b\} \uparrow \emptyset^* \uparrow \{c\} \uparrow \{c\}^\omega$,
- $\uparrow \emptyset^* \uparrow \{b, c\} \uparrow \{c\}^\omega$,
- $\uparrow \emptyset^* \uparrow \{c\} \uparrow \{c\}^* \uparrow \{b, c\} \uparrow \{c\}^\omega$,

which respectively amount to these LTL formulas:

- $\text{true} \mathbf{U} (b \wedge (\text{true} \mathbf{U} (c \wedge Gc)))$,
- $\text{true} \mathbf{U} ((b \wedge c) \wedge Gc)$,
- $\text{true} \mathbf{U} (c \wedge (c \mathbf{U} ((b \wedge c) \wedge Gc)))$.

Taking into account the acceptance condition of the automaton on its bottom-right transition, we obtain these LTL formulas:

- $\text{true} \mathbf{U} (b \wedge (\text{true} \mathbf{U} (c \wedge (Gc \wedge GFa))))$,
- $\text{true} \mathbf{U} ((b \wedge c) \wedge (Gc \wedge GFa))$,
- $\text{true} \mathbf{U} (c \wedge (c \mathbf{U} ((b \wedge c) \wedge (Gc \wedge GFa))))$. \square

For every $q, r \in Q$, let $X_{q,r} := \{A \subseteq AP : q \rightarrow^A r\}$. In general, the paths of \mathcal{A}_φ can always be captured in the fashion of Example 6 due to the following structure of \mathcal{A}_φ .

Proposition 8: Let $q, r \in Q$. It is the case that

- 1) $X_{q,r}$ is either empty or simple,

2) if $X_{q,r} \neq \emptyset$, then $X_{r,r} \neq \emptyset$,

3) if $X_{q,q} \neq \emptyset$, then $X_{q,q} \supseteq X_{q,r}$.

Moreover, given $\theta \in \mathcal{U}(q)$ and $A \subseteq AP$ such that $r = \theta[A]$, the representation of $X_{q,r}$ can be obtained in polynomial time.

A *linear path scheme (LPS)* of \mathcal{A}_φ is a simple path $r_0 \rightarrow r_1 \rightarrow \dots \rightarrow r_n$ of \mathcal{A}_φ such that $r_n \in F$. A word $w \in (2^{AP})^\omega$ is *accepted* by such an LPS S , denoted $w \in L(S)$, iff \mathcal{A}_φ has an accepting path starting in r_0 and visiting r_1, \dots, r_n (possibly many times). For example, the LPS $(a \wedge Fb) \rightarrow Fb \rightarrow \text{true}$, from the top of Figure 3, accepts $w := \{a\} \emptyset \{a\} \{b\} \emptyset \{a\}^\omega$.

We say that an LTL formula is *linear* if it can be derived from ψ in this grammar:

$$\begin{aligned} \psi &::= A \wedge \psi' \mid \psi', \\ \psi' &::= B \mathbf{U} (B' \wedge \psi') \mid (GC_0) \wedge \bigwedge_{i=1}^n GFC_i, \end{aligned}$$

where $A, B, B', C_0, \dots, C_n \subseteq AP$, $\uparrow B \supseteq \uparrow B'$, and each subset D stands for formula $\bigwedge_{d \in D} d$. We prove that any LPS is equivalent to a linear formula.

Proposition 9: Given an LPS S of \mathcal{A}_φ , one can construct, in polynomial time, a linear formula ψ s.t. $w \in L(S)$ iff $w \models \psi$.

Proof: Let $r_0 \rightarrow r_1 \rightarrow \dots \rightarrow r_n \in F$ be the simple path given by S . We inductively construct a formula derived from ψ' if $X_{r_0, r_0} \neq \emptyset$, and from ψ otherwise.

If $n = 0$, then we have $r_0 = r_n$. As $r_0 \in F$, it is of the form $r_0 = G\psi \wedge \bigwedge_{j \in J} GF\psi'_j$. Recall that $T_{r_0, j} := \{(r_0, A, r_0) \in \rightarrow : A \in \uparrow \text{prop}(\psi'_j)\}$. Note that $X_{r_0, r_0} \neq \emptyset$ as $\text{prop}(\psi) \in X_{r_0, r_0}$. So, by Proposition 8, it is the case that $X_{r_0, r_0} = \uparrow C_0$ for some $C_0 \subseteq AP$. Let $\psi' := GC_0 \wedge \bigwedge_{j \in J} GFC_j$, where $C_j := \text{prop}(\psi'_j)$. By definition, $w \in L(S)$ iff $w \models \psi'$.

Assume $n > 0$. Let S' be the LPS $r_1 \rightarrow \dots \rightarrow r_n \in F$. By Proposition 8, since $X_{r_0, r_1} \neq \emptyset$, we have $X_{r_1, r_1} \neq \emptyset$. By induction hypothesis, there is a linear formula ψ' such that $w \in L(S')$ iff $w \models \psi'$. By Proposition 8, there exists $B \subseteq AP$ such that $X_{r_0, r_1} = \uparrow B$. If $X_{r_0, r_0} = \emptyset$, then we set $\psi := B \wedge \psi'$. Otherwise, by Proposition 8, there exists $A \subseteq AP$ such that $X_{r_0, r_0} = \uparrow A \supseteq \uparrow B$. Thus, we set $\psi := A \mathbf{U} (B \wedge \psi')$. By definition of acceptance, $w \in L(S)$ iff $w \models \psi$. \square

IV. P-COMPLETE FRAGMENTS

Let $\text{goals}(A \wedge \varphi) := \text{goals}(\varphi)$, $\text{goals}(B \mathbf{U} (B' \wedge \varphi)) := \text{goals}(\varphi)$, and $\text{goals}((GC_0) \wedge \bigwedge_{i=1}^n GFC_i) := \{C_1, \dots, C_n\}$. We say that a linear LTL formula φ , interpreted over executions, is *semi-bounded* if each zone of $\text{goals}(\varphi)$ is bounded. We will establish the following theorem by translating linear LTL formulas into a polynomial-time logic introduced in [12].

Theorem 1: The model-checking problem for semi-bounded linear LTL formulas is in P.

Before proving the above theorem, we use it to prove the previously announced P-completeness results.

Theorem 2: The model-checking problem is in P for these fragments: $\text{LTL}_B(\{F, G, \neg\})$, $\text{LTL}(\{F, \vee\})$ and $\text{LTL}(\{G, \wedge\})$.

Proof: Consider a formula from $\text{LTL}_B(\{F, G, \neg\})$. Note that $\neg F\varphi \equiv G\neg\varphi$ and $\neg\neg\varphi \equiv \varphi$. Thus, negations can be pushed inwards. Afterwards, we can simplify using $\text{FGF}\varphi \equiv$

$\text{GF}\varphi$ and $\text{GFG}\varphi \equiv \text{FG}\varphi$. If the resulting formula is negation-free, then it is of the form FZ , GZ , GFZ or FGZ . These are all linear as $\text{FZ} \equiv \mathbb{R}^d \cup Z$ and $\text{FGZ} \equiv \mathbb{R}^d \cup (\text{GZ})$. Thus, we are done by Theorem 1. If the resulting formula has a negation, then there are four forms to consider: (1) $\text{F}\neg Z$, (2) $\text{G}\neg Z$, (3) $\text{GF}\neg Z$ and (4) $\text{FG}\neg Z$. These are easy to handle:

- We have $x \models_M \text{F}\neg Z$ iff $x \models_M \text{GF}\neg Z$ iff $x \models_M \text{FG}\neg Z$ iff $x \notin Z$ or there is a mode $m \in M$ such that $m \neq 0$.
- We have $x \models_M \text{G}\neg Z$ iff $x \notin Z$ and there exists a mode $m \in M$ such that for all $\alpha \in \mathbb{R}_{>0}$: $x + \alpha m \notin Z$.

Since $\text{FF}\varphi \equiv \text{F}\varphi$ and $\text{F}(\varphi \vee \psi) \equiv (\text{F}\varphi) \vee (\text{F}\psi)$, any formula from $\text{LTL}(\{\text{F}, \vee\})$ can be turned into a disjunction of atomic propositions and formulas from $\text{LTL}(\{\text{G}, \wedge\})$. So, it suffices to check each disjunct in polynomial time.

As $\text{GG}\varphi \equiv \text{G}\varphi$ and $(\text{G}\varphi) \wedge (\text{G}\psi) \equiv \text{G}(\varphi \wedge \psi)$, we can transform formulas from $\text{LTL}(\{\text{G}, \wedge\})$ into the form $\psi \wedge \text{G}\psi'$, where ψ, ψ' are pseudo-atomic. The latter is linear, and hence can be model-checked in polynomial time. \square

Theorem 3: The model-checking problem is P-hard for both $\text{LTL}_B(\{\text{F}\})$ and $\text{LTL}_B(\{\text{G}\})$.

Proof: It follows by simple reductions from feasibility of linear programs and the monotone circuit-value problem. \square

A. A polynomial-time first-order logic

We recall a first-order logic over the reals introduced in [12]. It allows for conjunctions of *convex semi-linear Horn formulas*, i.e. formulas of this form:

$$\sum_{i=1}^d \mathbf{a}(i) \cdot \mathbf{x}(i) \sim c \vee \bigwedge_{i \in I} \bigwedge_{j \in J_i} \mathbf{x}(j) > 0,$$

where $\mathbf{a} \in \mathbb{Z}^d$, $c \in \mathbb{Z}$, $\sim \in \{<, \leq, =, \geq, >\}$, and I and each J_i is a finite set of indices. The problem of determining, given a formula φ from this logic, whether there exists $\mathbf{x} \in \mathbb{R}_{\geq 0}^d$ such that $\varphi(\mathbf{x})$ holds, can be solved in polynomial time [12].

This result extends easily to solutions where $\mathbf{x}(j) \in \mathbb{R}$ is allowed, provided that $\mathbf{x}(j)$ is never used in disjuncts. Indeed, it suffices to introduce two variables $y, z \in \mathbb{R}_{\geq 0}$ and replace each occurrence of $\mathbf{x}(j)$ with $y - z$.

Given $\mathbf{x}(i), \mathbf{x}(j) \in \mathbb{R}_{\geq 0}$, we will use $\mathbf{x}(i) > 0 \rightarrow \mathbf{x}(j) > 0$ as short for $\mathbf{x}(i) = 0 \vee \mathbf{x}(j) > 0$.

B. Expressing \rightarrow_Z^* in first-order logic

We first seek to build a formula φ_Z from the aforementioned logic such that $\varphi_Z(\mathbf{x}, \boldsymbol{\lambda}, \mathbf{y})$ holds iff there is a finite schedule π with $\mathbf{x} \rightarrow_Z^\pi \mathbf{y}$ and $\boldsymbol{\pi} = \boldsymbol{\lambda}$. Let us fix zone Z and modes $M = \{\mathbf{m}_1, \dots, \mathbf{m}_n\}$. We take inspiration from the characterization of continuous Petri nets reachability of [10, Thm. 20], which is equivalent to finding a Parikh image that (1) admits the right effect, (2) is forward fireable, and (3) is backward fireable. The forthcoming Proposition 13 similarly characterizes $\mathbf{x} \rightarrow_Z^\pi \mathbf{y}$.

Proposition 10: Let $\mathbf{x} \rightarrow_Z^\pi \mathbf{y}$. There exist π_x and π_y such that $\mathbf{x} \rightarrow_Z^{\pi_x} \mathbf{x}$, $\mathbf{x} \rightarrow_Z^{\pi_y} \mathbf{y}$, $\text{supp}(\pi_x) = \text{supp}(\pi_y) = \text{supp}(\pi)$ and $|\pi_x| = |\pi_y| = |\text{supp}(\pi)|$.

Lemma 2: Let $\rho(\alpha, \mathbf{m})\rho'$ be a schedule. This holds:

- If $\mathbf{x} \rightarrow_Z^{\rho(\alpha, \mathbf{m})\rho'}$, then $\mathbf{x} \rightarrow_Z^{\rho(\frac{\alpha}{2}, \mathbf{m})\frac{1}{2}\rho'(\frac{\alpha}{2}, \mathbf{m})}$,

- If $\mathbf{x} \rightarrow_Z^{\rho'(\alpha, \mathbf{m})\rho}$, then $\mathbf{x} \rightarrow_Z^{\rho'(\frac{\alpha}{2}, \mathbf{m})\frac{1}{2}\rho(\frac{\alpha}{2}, \mathbf{m})\rho}$.

Proposition 11: Let $\mathbf{x} \rightarrow_Z^\pi \mathbf{y}$. There exist $\beta \in \mathbb{N}_{\geq 1}$, $\mathbf{x} \rightarrow_Z^{\pi'} \mathbf{y}_Z$ and $\mathbf{x}_Z \rightarrow_Z^{\pi''} \mathbf{y}$ such that $|\pi| = |\pi'| = |\pi''|$, $\text{supp}(\pi) = \text{supp}(\pi') = \text{supp}(\pi'')$, and, for every $\mathbf{m} \in \text{supp}(\pi)$, it is the case that $\mathbf{x}_Z \rightarrow_Z^{(1/\beta)\mathbf{m}}$ and $\mathbf{x} \rightarrow_Z^{(1/\beta)\mathbf{m}} \mathbf{y}_Z$.

Proposition 12: Let $\mathbf{x} \rightarrow_Z^\pi \mathbf{y}$, $k := |\pi|$ and $\beta \in \mathbb{N}_{\geq 1}$ be such that $\mathbf{x} \rightarrow_Z^{(1/\beta)\pi(i)}$ and $\mathbf{x} \rightarrow_Z^{(1/\beta)\pi(i)} \mathbf{y}$ hold for all $i \in [1..k]$. It is the case that $\mathbf{x} \rightarrow_Z^{\pi'} \mathbf{y}$, where $\pi' := ((1/(\beta k))\pi)^{\beta k}$.

Proposition 13: It is the case that $\mathbf{x} \rightarrow_Z^\pi \mathbf{y}$ iff there exist $\pi', \pi_{\text{fwd}}, \pi_{\text{bwd}}$ with

- $\text{supp}(\pi') = \text{supp}(\pi_{\text{fwd}}) = \text{supp}(\pi_{\text{bwd}}) = \text{supp}(\pi)$,
- $\mathbf{x} \rightarrow_Z^{\pi'} \mathbf{y}$, $\mathbf{x} \rightarrow_Z^{\pi_{\text{fwd}}}$ and $\mathbf{x} \rightarrow_Z^{\pi_{\text{bwd}}}$.

Proof: \Rightarrow) It suffices to take $\pi' = \pi_{\text{fwd}} = \pi_{\text{bwd}} := \pi$.

\Leftarrow) Let β and the following be given by Proposition 11:

$$\mathbf{x} \rightarrow_Z^{\pi'} \mathbf{x}_Z \text{ and } \mathbf{y}_Z \rightarrow_Z^{\pi'_{\text{bwd}}} \mathbf{y}.$$

Let $\gamma \in \mathbb{N}_{\geq 1}$ be sufficiently large so that $\pi'' \geq \frac{1}{\gamma}(\pi'_{\text{fwd}} + \pi'_{\text{bwd}})$. Such a γ exists as $\text{supp}(\pi'_{\text{fwd}}) = \text{supp}(\pi'_{\text{bwd}}) = \text{supp}(\pi')$. Let π'' be any schedule with $\pi'' = \pi' - \frac{1}{\gamma}(\pi'_{\text{fwd}} + \pi'_{\text{bwd}})$. We have

$$\mathbf{x} \rightarrow_Z^{\frac{1}{\gamma}\pi'_{\text{fwd}}} \mathbf{x}'_Z \rightarrow_Z^{\pi''} \mathbf{y}'_Z \rightarrow_Z^{\frac{1}{\gamma}\pi'_{\text{bwd}}} \mathbf{y}.$$

So, by invoking Proposition 12 with $\beta \cdot \gamma \cdot \lceil \text{time}(\pi'') \rceil \in \mathbb{N}_{\geq 1}$, we obtain π with $\mathbf{x} \rightarrow_Z^\pi \mathbf{y}$. \square

We may now conclude this subsection by building a suitable first-order formula. First, we define

$$\begin{aligned} \theta_Z(\mathbf{x}, \mathbf{s}, \mathbf{y}) := & \exists \{z_{i,j} \in Z\}_{i \in [1..n], j \in [0..n]} \\ & \exists \{\lambda_{i,j} \in \mathbb{R}_{\geq 0}\}_{i,j \in [1..n]} \\ & \exists \boldsymbol{\alpha} \in \mathbb{R}_{\geq 0}^n : (\mathbf{x} = \mathbf{z}_{1,0}) \wedge (\mathbf{y} = \mathbf{z}_{n,n}) \wedge \\ & \bigwedge_{i \in [1..n]} \bigwedge_{j \in [1..n]} (z_{i,j} = z_{i,j-1} + \lambda_{i,j} \cdot \mathbf{m}_j) \wedge \\ & \bigwedge_{i \in [2..n]} (z_{i,0} = z_{i-1,n}) \wedge \\ & \bigwedge_{j \in [1..n]} (\boldsymbol{\alpha}(j) = \sum_{i \in [1..n]} \lambda_{i,j}) \wedge \\ & \bigwedge_{j \in [1..n]} (\mathbf{s}(j) > 0 \leftrightarrow \boldsymbol{\alpha}(j) > 0). \end{aligned}$$

Let $\varphi_Z(\mathbf{x}, \boldsymbol{\lambda}, \mathbf{y}) := \psi(\mathbf{x}, \boldsymbol{\lambda}, \mathbf{y}) \wedge \psi_{\text{fwd}}(\mathbf{x}, \boldsymbol{\lambda}) \wedge \psi_{\text{bwd}}(\boldsymbol{\lambda}, \mathbf{y})$ where

$$\psi(\mathbf{x}, \boldsymbol{\lambda}, \mathbf{y}) := (\mathbf{y} = \mathbf{x} + \lambda(1) \cdot \mathbf{m}_1 + \dots + \lambda(n) \cdot \mathbf{m}_n),$$

$$\psi_{\text{fwd}}(\mathbf{x}, \mathbf{s}) := \exists \mathbf{x}' \in \mathbb{R}^d : \theta_Z(\mathbf{x}, \mathbf{s}, \mathbf{x}'),$$

$$\psi_{\text{bwd}}(\mathbf{s}, \mathbf{y}) := \exists \mathbf{y}' \in \mathbb{R}^d : \theta_Z(\mathbf{y}', \mathbf{s}, \mathbf{y}).$$

Proposition 14: It is the case that $\varphi_Z(\mathbf{x}, \boldsymbol{\lambda}, \mathbf{y})$ holds iff $\mathbf{x} \rightarrow_Z^\pi \mathbf{y}$ for some finite schedule π such that $\boldsymbol{\pi} = \boldsymbol{\lambda}$.

Proof: First note that formula $\theta_Z(\mathbf{u}, \mathbf{s}, \mathbf{v})$ guesses a schedule of size at most n^2 from \mathbf{u} to \mathbf{v} that remains within Z using precisely the modes from $\{\mathbf{m}_j : j \in [1..n], \mathbf{s}(j) > 0\}$. The reason θ_Z uses a schedule of size n^2 , rather than n , is to guess the order in which modes are first used.

\Rightarrow) Suppose that $\psi(\mathbf{x}, \boldsymbol{\lambda}, \mathbf{y}) \wedge \psi_{\text{fwd}}(\mathbf{x}, \boldsymbol{\lambda}) \wedge \psi_{\text{bwd}}(\boldsymbol{\lambda}, \mathbf{y})$ holds. Let $\pi' := \prod_{i=1}^n \lambda(i) \mathbf{m}_i$,

$$\pi_{\text{fwd}} := \prod_{i=1}^n \prod_{j=1}^n \lambda_{i,j}^{\text{fwd}} \mathbf{m}_j \text{ and } \pi_{\text{bwd}} := \prod_{i=1}^n \prod_{j=1}^n \lambda_{i,j}^{\text{bwd}} \mathbf{m}_j,$$

with the convention that $0 \cdot \mathbf{m}_j$ stands for the empty schedule. We clearly have $\mathbf{x} \rightarrow^{\pi'} \mathbf{y}$. By the above observation on θ_Z , we further have $\mathbf{x} \rightarrow^{\pi_{\text{fwd}}}$ and $\rightarrow^{\pi_{\text{bwd}}} \mathbf{y}$. Moreover, $\text{supp}(\pi') = \text{supp}(\pi_{\text{fwd}}) = \text{supp}(\pi_{\text{bwd}}) = \boldsymbol{\lambda}$. By Proposition 13, we obtain $\mathbf{x} \rightarrow^{\pi} \mathbf{y}$ for some π with $\pi = \boldsymbol{\lambda}$.

\Leftarrow) Let $\mathbf{x} \rightarrow^{\pi} \mathbf{y}$ and $\boldsymbol{\lambda} := \pi$. By Proposition 10, there are π_x and π_y such that $\mathbf{x} \rightarrow^{\pi_x} \mathbf{x}$, $\rightarrow^{\pi_y} \mathbf{y}$, $\text{supp}(\pi_x) = \text{supp}(\pi_y) = \text{supp}(\pi)$ and $|\pi_x| = |\pi_y| = |\text{supp}(\pi)|$. Thus, we can use π , π_x and π_y to satisfy $\psi(\mathbf{x}, \boldsymbol{\lambda}, \mathbf{y})$, $\psi_{\text{fwd}}(\mathbf{x}, \boldsymbol{\lambda})$ and $\psi_{\text{bwd}}(\boldsymbol{\lambda}, \mathbf{y})$. \square

C. Expressing $\text{GZ} \wedge \text{GF}X \wedge \text{GF}Y$ in first-order logic

In this subsection, we build a formula $\varphi_{\text{GZ} \wedge \text{GF}X \wedge \text{GF}Y}$ from the logic of Section IV-A such that $\mathbf{x} \models_M \text{GZ} \wedge \text{GF}X \wedge \text{GF}Y$ iff $\varphi_{\text{GZ} \wedge \text{GF}X \wedge \text{GF}Y}(\mathbf{x})$ holds. Let us fix an MMS M .

Proposition 15: Let Z be a zone, let π be a schedule, let $\mathbf{x}, \mathbf{x}', \mathbf{y} \in Z$ and let $\beta \in (0, 1]$. Let $\mathbf{z} := \beta \mathbf{x} + (1 - \beta) \mathbf{y}$ and $\mathbf{z}' := \beta \mathbf{x}' + (1 - \beta) \mathbf{y}$. If $\mathbf{x} \rightarrow^{\pi} \mathbf{x}'$ holds, then $\mathbf{z} \rightarrow^{\beta \pi} \mathbf{z}'$.

Proposition 16: Let X, Y, Z be zones where at least one of the three zones is bounded. Let $\mathbf{z} \models_M \text{GZ} \wedge \text{GF}X \wedge \text{GF}Y$. There exist $\mathbf{x}_f \in X \cap Z$, $\mathbf{y}_f \in Y \cap Z$ and finite schedules π, π' such that $\mathbf{z} \rightarrow^* \mathbf{x}_f \rightarrow^{\pi} \mathbf{y}_f \rightarrow^{\pi'} \mathbf{x}_f$ and $\|\pi + \pi'\| \geq 1$.

Proof: Let $X' := X \cap Z$ and $Y' := Y \cap Z$. By assumption, there exist $\mathbf{x}_0, \mathbf{x}_1, \dots \in X'$ and $\mathbf{y}_0, \mathbf{y}_1, \dots \in Y'$ such that

$$\mathbf{z} \rightarrow_Z^* \mathbf{x}_0 \rightarrow_Z^{\pi_0} \mathbf{y}_0 \rightarrow_Z^{\pi'_0} \mathbf{x}_1 \rightarrow_Z^{\pi_1} \mathbf{y}_1 \rightarrow_Z^{\pi'_1} \dots,$$

and $\|\pi_i + \pi'_i\| \geq 1$ for all $i \in \mathbb{N}$. Note that the latter follows from non-Zenoness.

Let $\mathbf{A}_1 \ell \leq \mathbf{b}_1$ and $\mathbf{A}_2 \ell \leq \mathbf{b}_2$ be the systems of inequalities that respectively represent zones X' and Y' . We define \mathbf{M} as the matrix such that each column is a mode from M . Let \mathcal{S} denote the following system:

$$\begin{aligned} & \exists \mathbf{u}_1, \mathbf{u}_2, \mathbf{u}_3 \geq \mathbf{0} : \\ & \begin{bmatrix} \mathbf{A}_1 \mathbf{M} & \mathbf{0} & \mathbf{0} \\ \mathbf{A}_2 \mathbf{M} & \mathbf{A}_2 \mathbf{M} & \mathbf{0} \\ \mathbf{0} & \mathbf{M} & \mathbf{M} \\ \mathbf{0} & -\mathbf{M} & -\mathbf{M} \\ \mathbf{0}^T & -\mathbf{1}^T & -\mathbf{1}^T \end{bmatrix} \begin{bmatrix} \mathbf{u}_1 \\ \mathbf{u}_2 \\ \mathbf{u}_3 \end{bmatrix} \leq \begin{bmatrix} \mathbf{b}_1 - \mathbf{A}_1 \mathbf{z} \\ \mathbf{b}_2 - \mathbf{A}_2 \mathbf{z} \\ \mathbf{0} \\ \mathbf{0} \\ -1 \end{bmatrix}. \end{aligned}$$

Observe that \mathcal{S} is equivalent to the existence of $\mathbf{x}_f \in X'$, $\mathbf{y}_f \in Y'$ and π, π' such that

$$\mathbf{z} \rightarrow^* \mathbf{x}_f \rightarrow^{\pi} \mathbf{y}_f \rightarrow^{\pi'} \mathbf{x}_f \text{ and } \|\pi + \pi'\| \geq 1.$$

For the sake of contradiction, suppose that \mathcal{S} has no solution. By Farkas' lemma, the following system \mathcal{S}' has a solution:

$$\begin{aligned} & \exists \mathbf{v}_1, \mathbf{v}_2 \in \mathbb{R}_{\geq 0}^d, \mathbf{v}_3 \in \mathbb{R}^d, \mathbf{v}_4 \in \mathbb{R}_{\geq 0} : \\ & \begin{bmatrix} \mathbf{M}^T \mathbf{A}_1^T & \mathbf{M}^T \mathbf{A}_2^T & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{M}^T \mathbf{A}_2^T & \mathbf{M}^T & -\mathbf{1} \\ \mathbf{0} & \mathbf{0} & \mathbf{M}^T & -\mathbf{1} \end{bmatrix} \begin{bmatrix} \mathbf{v}_1 \\ \mathbf{v}_2 \\ \mathbf{v}_3 \\ \mathbf{v}_4 \end{bmatrix} \geq \mathbf{0}, \\ & [(\mathbf{b}_1 - \mathbf{A}_1 \mathbf{z})^T \quad (\mathbf{b}_2 - \mathbf{A}_2 \mathbf{z})^T \quad -\mathbf{1}] \begin{bmatrix} \mathbf{v}_1 \\ \mathbf{v}_2 \\ \mathbf{v}_4 \end{bmatrix} < 0. \end{aligned}$$

Using the above, we will construct linear functions g and h such that $\lim_{i \rightarrow \infty} g(\mathbf{x}_i) = \lim_{i \rightarrow \infty} h(\mathbf{y}_i) = \infty$. Since either X' or Y' is bounded, this yields a contradiction. We make a case distinction on the value of v_4 .

Case $v_4 > 0$. We have $\mathbf{v}_3^T \mathbf{M} \geq \mathbf{1}^T v_4 > 0$. Let $g(\mathbf{x}_i) := \mathbf{v}_3^T (\mathbf{x}_i - \mathbf{x}_0)$ and $h(\mathbf{y}_i) := \mathbf{v}_3^T (\mathbf{y}_i - \mathbf{y}_0)$. For every $i \geq 1$,

$$\begin{aligned} g(\mathbf{x}_i) &= \mathbf{v}_3^T (\mathbf{x}_i - \mathbf{x}_0) \\ &= \mathbf{v}_3^T (\mathbf{x}_{i-1} - \mathbf{x}_0) + \mathbf{v}_3^T (\mathbf{x}_i - \mathbf{x}_{i-1}) \\ &= g(\mathbf{x}_{i-1}) + \mathbf{v}_3^T \mathbf{M} (\pi_{i-1} + \pi'_{i-1}) \\ &\geq g(\mathbf{x}_{i-1}) + v_4 \mathbf{1}^T (\pi_{i-1} + \pi'_{i-1}) \\ &\geq g(\mathbf{x}_{i-1}) + v_4 \quad (\text{by } \|\pi_{i-1} + \pi'_{i-1}\| \geq 1). \end{aligned}$$

By the above, we conclude that $\lim_{i \rightarrow \infty} g(\mathbf{x}_i) = \infty$. The proof for function h is symmetric.

Case $v_4 = 0$. In this case, \mathcal{S}' amounts to:

$$\mathbf{v}_1^T \mathbf{A}_1 \mathbf{M} + \mathbf{v}_2^T \mathbf{A}_2 \mathbf{M} \geq \mathbf{0}^T \quad (1)$$

$$\mathbf{v}_2^T \mathbf{A}_2 \mathbf{M} + \mathbf{v}_3^T \mathbf{M} \geq \mathbf{0}^T \quad (2)$$

$$\mathbf{v}_3^T \mathbf{M} \geq \mathbf{0}^T \quad (3)$$

$$\mathbf{v}_1^T (\mathbf{b}_1 - \mathbf{A}_1 \mathbf{z}) + \mathbf{v}_2^T (\mathbf{b}_2 - \mathbf{A}_2 \mathbf{z}) < 0. \quad (4)$$

Let $\lambda := -(\mathbf{v}_1^T (\mathbf{b}_1 - \mathbf{A}_1 \mathbf{z}) + \mathbf{v}_2^T (\mathbf{b}_2 - \mathbf{A}_2 \mathbf{z}))$. By (4), we have $\lambda > 0$. Let

$$\begin{aligned} f'(\mathbf{x}, \mathbf{y}) &:= -(\mathbf{v}_1^T \mathbf{A}_1 (\mathbf{x} - \mathbf{z}) + \mathbf{v}_2^T \mathbf{A}_2 (\mathbf{y} - \mathbf{z})) + \\ &\quad \mathbf{v}_2^T \mathbf{A}_2 (\mathbf{y} - \mathbf{x}) + \mathbf{v}_3^T (\mathbf{y} - \mathbf{x}), \\ f(\mathbf{x}, \mathbf{y}) &:= \mathbf{v}_3^T (\mathbf{y} - \mathbf{x}). \end{aligned}$$

Using (1) and (2), it can be shown that $f(\mathbf{x}, \mathbf{y}) \geq \lambda$. From this, we can conclude. Let $g(\mathbf{x}_i) := f(\mathbf{x}_0, \mathbf{x}_i)$ and $h(\mathbf{y}_i) := f(\mathbf{x}_0, \mathbf{y}_i)$ for all $i \in \mathbb{N}$. For all $i \geq 1$, we have

$$\begin{aligned} & f(\mathbf{x}_0, \mathbf{y}_i) \\ &= \mathbf{v}_3^T (\mathbf{y}_i - \mathbf{x}_0) \\ &= \mathbf{v}_3^T ((\mathbf{y}_i - \mathbf{x}_i) + (\mathbf{x}_i - \mathbf{y}_{i-1}) + (\mathbf{y}_{i-1} - \mathbf{x}_0)) \\ &= \mathbf{v}_3^T (\mathbf{y}_i - \mathbf{x}_i) + \mathbf{v}_3^T (\mathbf{x}_i - \mathbf{y}_{i-1}) + \mathbf{v}_3^T (\mathbf{y}_{i-1} - \mathbf{x}_0) \\ &= f(\mathbf{x}_i, \mathbf{y}_i) + \mathbf{v}_3^T \mathbf{M} \pi'_{i-1} + f(\mathbf{x}_0, \mathbf{y}_{i-1}) \\ &\geq f(\mathbf{x}_0, \mathbf{y}_{i-1}) + f(\mathbf{x}_i, \mathbf{y}_i) \quad (\text{by (3)}) \\ &\geq f(\mathbf{x}_0, \mathbf{y}_{i-1}) + \lambda. \end{aligned}$$

By the above, we have $\lim_{i \rightarrow \infty} h(\mathbf{y}_i) = \infty$.

Similarly, for every $i \geq 1$, we have

$$f(\mathbf{x}_0, \mathbf{x}_i) = f(\mathbf{x}_0, \mathbf{y}_{i-1}) + \mathbf{v}_3^T(\mathbf{x}_i - \mathbf{y}_{i-1}) \geq f(\mathbf{x}_0, \mathbf{y}_{i-1}),$$

which implies $g(\mathbf{x}_i) \geq h(\mathbf{y}_{i-1})$. So, $\lim_{i \rightarrow \infty} g(\mathbf{x}_i) = \infty$. \square

We now seek to show the following proposition.

Proposition 17: Let X, Y and Z be zones. Let $\mathbf{z}, \mathbf{z}' \in Z$, $\mathbf{x}_0, \mathbf{x}', \mathbf{x}_f \in X \cap Z$ and $\mathbf{y}_0, \mathbf{y}_f \in Y \cap Z$ be such that

- $\mathbf{z} \rightarrow_Z^* \mathbf{z}' \rightarrow_Z^{\pi''} \mathbf{x}_0 \rightarrow_Z^{\pi} \mathbf{y}_0 \rightarrow_Z^{\pi'} \mathbf{x}'$,
- $\mathbf{x}' \rightarrow_Z^{\rho} \mathbf{x}_f \rightarrow_Z^{\rho'} \mathbf{y}_f \rightarrow_Z^{\rho''} \mathbf{x}_f$,
- $\text{supp}(\rho) = \text{supp}(\pi) = \text{supp}(\pi') = \text{supp}(\pi'')$,
- $\text{supp}(\rho') \cup \text{supp}(\rho'') \subseteq \text{supp}(\rho)$ and $\|\rho' + \rho''\| \geq 1$.

It is the case that $\mathbf{z} \models \mathbf{GZ} \wedge \mathbf{GF}X \wedge \mathbf{GF}Y$.

To prove the above proposition, we build a schedule from the given schedules. By assumption, there exists a sufficiently small $\epsilon \in \mathbb{R}_{>0}$ such that $\rho \geq \epsilon \cdot (\pi + \pi')$. Let $\lambda := 1 - (1/(1 + \epsilon))$. For every $n \in \mathbb{N}_{\geq 1}$, let

$$\begin{aligned} \mathbf{x}_n &:= \mathbf{y}_{n-1} + \lambda^{n-1} \left(\Delta_{\pi'} + \frac{\Delta_{\rho} - \epsilon \Delta_{\pi \pi'}}{1 + \epsilon} \right) \\ &\quad + (1 - \lambda^{n-1}) \Delta_{\rho''}, \\ \mathbf{y}_n &:= \mathbf{x}_n + \lambda^n \Delta_{\pi} + (1 - \lambda^n) \Delta_{\rho'}. \end{aligned}$$

Proposition 18: For every $n \in \mathbb{N}$, it is the case that $\mathbf{x}_n = \lambda^n \mathbf{x}_0 + (1 - \lambda^n) \mathbf{x}_f$ and $\mathbf{y}_n = \lambda^n \mathbf{y}_0 + (1 - \lambda^n) \mathbf{y}_f$.

The following proposition proves Proposition 17.

Proposition 19: It is the case that (1) $\mathbf{x}_n \rightarrow_Z^* \mathbf{y}_n$ and (2) $\mathbf{y}_n \rightarrow_Z^* \mathbf{x}_{n+1}$ for all $n \in \mathbb{N}$.

Proof: (1) Recall that $\mathbf{x}_0 \rightarrow_Z^{\pi} \mathbf{y}_0$ and $\mathbf{x}_f \in Z$. Therefore, by Propositions 15 and 18, we have

$$\begin{aligned} \mathbf{x}_n &= (\lambda^n \mathbf{x}_0 + (1 - \lambda^n) \mathbf{x}_f) \\ &\quad \rightarrow_Z^{\lambda^n \pi} (\lambda^n \mathbf{y}_0 + (1 - \lambda^n) \mathbf{x}_f). \end{aligned} \quad (5)$$

Similarly, by Propositions 15 and 18, we have

$$\begin{aligned} (\lambda^n \mathbf{x}_0 + (1 - \lambda^n) \mathbf{x}_f) \\ \rightarrow_Z^{\lambda^n \pi} (\lambda^n \mathbf{y}_0 + (1 - \lambda^n) \mathbf{y}_f) = \mathbf{y}_n. \end{aligned} \quad (6)$$

By definition, we have

$$\mathbf{y}_n = \mathbf{x}_n + \lambda^n \Delta_{\pi} + (1 - \lambda^n) \Delta_{\rho'}. \quad (7)$$

Altogether, (5)–(7) yield

$$\mathbf{x}_n \rightarrow_Z^{\lambda^n \pi (1 - \lambda^n) \rho'} \mathbf{y}_n, \mathbf{x}_n \rightarrow_Z^{\lambda^n \pi} \text{ and } \rightarrow_Z^{\lambda^n \pi} \mathbf{y}_n.$$

As $\text{supp}(\rho') \subseteq \text{supp}(\pi)$, we have $\text{supp}(\lambda^n \pi) = \text{supp}(\lambda^n \pi (1 - \lambda^n) \rho')$. So, by Proposition 13, we conclude that $\mathbf{x}_n \rightarrow_Z^* \mathbf{y}_n$.

(2) By Propositions 15 and 18, we have

$$\mathbf{y}_n = \lambda^n \mathbf{y}_0 + (1 - \lambda^n) \mathbf{y}_f \rightarrow_Z^{\lambda^n \pi'} (\lambda^n \mathbf{x}' + (1 - \lambda^n) \mathbf{y}_f). \quad (8)$$

Similarly, by Propositions 15 and 18, we have

$$(\lambda^{n+1} \mathbf{z}' + (1 - \lambda^{n+1}) \mathbf{x}_f) \rightarrow_Z^{\lambda^{n+1} \pi''} \mathbf{x}_{n+1}. \quad (9)$$

Let π_n be any finite schedule such that $\pi_n = \rho - \epsilon(\pi + \pi')$. We have

$$\mathbf{y}_n \rightarrow_Z^{\lambda^n \pi' (\lambda^n / (1 + \epsilon)) \pi_n (1 - \lambda^n) \rho''} \mathbf{x}_{n+1}. \quad (10)$$

By (8)–(10), $\text{supp}(\rho'') \subseteq \text{supp}(\rho) = \text{supp}(\pi) = \text{supp}(\pi') = \text{supp}(\pi'')$ and Proposition 13, we obtain $\mathbf{y}_n \rightarrow_Z^* \mathbf{x}_{n+1}$. \square

We may now conclude this subsection by building a suitable first-order formula. Let $\mathbf{x} \rightarrow_Z^{\lambda} \mathbf{y}$ be a shorthand for formula $\varphi_Z(\mathbf{x}, \lambda, \mathbf{y})$ from Section IV-B, and let $\mathbf{x} \rightarrow_Z^* \mathbf{y}$ stand for $\exists \lambda \geq \mathbf{0} : \mathbf{x} \rightarrow_Z^{\lambda} \mathbf{y}$. Let $M = \{m_1, \dots, m_n\}$. We define $\varphi_{\mathbf{GZ} \wedge \mathbf{GF}X \wedge \mathbf{GF}Y}(\mathbf{z})$ by

$$\begin{aligned} \exists \mathbf{z}' \in Z; \mathbf{x}_0, \mathbf{x}', \mathbf{x}_f \in X \cap Z; \mathbf{y}_0, \mathbf{y}_f \in Y \cap Z; \\ \pi, \pi', \pi'', \rho, \rho', \rho'' \geq \mathbf{0} : \\ \mathbf{z} \rightarrow_Z^* \mathbf{z}' \rightarrow_Z^{\pi''} \mathbf{x}_0 \rightarrow_Z^{\pi} \mathbf{y}_0 \rightarrow_Z^{\pi'} \mathbf{x}' \wedge \end{aligned} \quad (11)$$

$$\mathbf{x}' \rightarrow_Z^{\rho} \mathbf{x}_f \rightarrow_Z^{\rho'} \mathbf{y}_f \rightarrow_Z^{\rho''} \mathbf{x}_f \wedge \quad (12)$$

$$\bigwedge_{j \in [1..n]} \theta_j \wedge \sum_{j \in [1..n]} (\rho'(j) + \rho''(j)) \geq 1,$$

where

$$\begin{aligned} \theta_j &= (\pi(j) > 0 \leftrightarrow \pi'(j) > 0 \leftrightarrow \pi''(j) > 0 \leftrightarrow \rho(j) > 0) \\ &\quad \wedge (\rho'(j) > 0 \rightarrow \rho(j) > 0) \wedge (\rho''(j) > 0 \rightarrow \rho(j) > 0). \end{aligned}$$

Proposition 20: It is the case that $\mathbf{z} \models_M \mathbf{GZ} \wedge \mathbf{GF}X \wedge \mathbf{GF}Y$ iff $\varphi_{\mathbf{GZ} \wedge \mathbf{GF}X \wedge \mathbf{GF}Y}(\mathbf{z})$ holds.

Proof: \Leftarrow) It follows directly from Proposition 17, since φ_Z is the same statement written in logic.

\Rightarrow) Let π be a non-Zeno infinite schedule such that $\sigma := \text{exec}(\pi, \mathbf{z}) \models \mathbf{GZ} \wedge \mathbf{GF}X \wedge \mathbf{GF}Y$. Let M' be the set of modes used infinitely often in π . From \mathbf{z} , we can move along σ to a point \mathbf{z}' from which only modes of M' are used. We can further go to a point $\mathbf{x}_0 \in X \cap Z$ where all modes of M' are used from \mathbf{z}' to \mathbf{x}_0 . Similarly, we can use all modes of M' from \mathbf{x}_0 to some $\mathbf{y}_0 \in Y \cap Z$, and likewise from \mathbf{y}_0 to some $\mathbf{x}' \in X \cap Z$. The resulting sequence satisfies (11).

Note that $\mathbf{GZ} \wedge \mathbf{GF}X \wedge \mathbf{GF}Y$ is satisfied from \mathbf{x}' in M' . So, we can invoke Proposition 16 to satisfy (12). We are done since $\text{supp}(\pi'') = \text{supp}(\pi) = \text{supp}(\pi') = \text{supp}(\rho) = M'$, $\text{supp}(\rho') \subseteq M'$ and $\text{supp}(\rho'') \subseteq M'$. \square

D. Expressing \mathbf{GZ} in first-order logic

Let us now handle the special case $n = 0$ of the previous subsection. We claim that $\mathbf{z} \models_M \mathbf{GZ}$ iff $\varphi_{\mathbf{GZ}}(\mathbf{z})$ holds, where

$$\begin{aligned} \varphi_{\mathbf{GZ}}(\mathbf{z}) &:= \exists \mathbf{z}' \in Z, \pi, \pi' \geq \mathbf{0} : \mathbf{z} \rightarrow_Z^{\pi} \wedge \mathbf{z} \rightarrow_Z^{\pi'} \mathbf{z}' \wedge \\ &\quad \mathbf{Az}' \leq \mathbf{Az} \wedge \bigwedge_{j \in [1..n]} (\pi'(j) > 0 \rightarrow \pi(j) > 0) \wedge \\ &\quad \sum_{j \in [1..n]} \pi'(j) \geq 1. \end{aligned}$$

A simpler proof to the one of Proposition 16 yields:

Proposition 21: If $\mathbf{z} \models_M \mathbf{GZ}$, then there exist π and \mathbf{z}' such that $\mathbf{z} \rightarrow_Z^{\pi} \mathbf{z}'$, $\mathbf{Az}' \leq \mathbf{Az}$ and $\|\pi\| \geq 1$.

Proposition 22: Let $\mathbf{z}, \mathbf{z}' \in Z$ and ρ be a finite schedule. If $\mathbf{z} \rightarrow_Z^{\rho}$ and $\mathbf{Az}' \leq \mathbf{Az}$ then $\mathbf{z}' \rightarrow_Z^{\rho}$.

Proposition 23: It is the case that $\mathbf{z} \models_M \mathbf{GZ}$ iff $\varphi_{\mathbf{GZ}}(\mathbf{z})$.

Proof: \Leftarrow) Let π be the schedule such that $\mathbf{z} \rightarrow_Z^{\pi}$. By Proposition 11, we obtain some $\beta \in \mathbb{N}_{\geq 1}$ and $\mathbf{z} \rightarrow_Z^{\rho_{\beta}} \mathbf{z}_0$ with $\mathbf{z}_0 \rightarrow_Z^{(1/\beta)m}$ for every $m \in \text{supp}(\pi)$. Let π' be the schedule such that $\mathbf{z} \rightarrow_Z^{\pi'} \mathbf{z}'$.

Let $\rho' := (1/\beta \cdot \text{time}(\pi'))\pi'$. For all $i \in \mathbb{N}$, let $z_{i+1} := z_i + \Delta_{\rho'}$. Since $\text{supp}(\pi') \subseteq \text{supp}(\pi)$, we have $z_0 \rightarrow_Z^{\rho'} z_1$. Moreover, as $\mathbf{A}\Delta_{\pi'} = \mathbf{A}(z' - z) \leq \mathbf{0}$, we have $\mathbf{A}\Delta_{\rho'} \leq \mathbf{0}$ and hence $\mathbf{A}z_1 \leq \mathbf{A}z_0$. By Proposition 22, we obtain $z_0 \rightarrow_Z^{\rho'} z_1$. By the same reasoning, we conclude that

$$z \rightarrow_Z^{\rho'} z_0 \rightarrow_Z^{\rho'} z_1 \rightarrow_Z^{\rho'} z_2 \rightarrow_Z^{\rho'} \dots$$

\Rightarrow) Let π'' be a non-Zeno infinite schedule such that $\sigma := \text{exec}(\pi'', z) \models \mathbf{G}Z$. Let M' be the set of modes used in π'' . From z , we move along π'' to some point where all modes from M' have been used. We take π as such a prefix. The other constraints hold by Proposition 21. \square

E. From $\mathbf{G}Z_0 \wedge \bigwedge_{i=1}^n \mathbf{G}FZ_i$ to $\mathbf{G}Z \wedge \mathbf{G}F\mathbf{X} \wedge \mathbf{G}F\mathbf{Y}$

Lemma 3: Given a d -dimensional MMS M , point $\mathbf{x} \in \mathbb{R}^d$ and zones $Z_0, \dots, Z_n \subseteq \mathbb{R}^d$, it is possible to construct, in polynomial time, an nd -dimensional MMS M' and zones $X, Y, Z \subseteq \mathbb{R}^{nd}$ such that $\mathbf{x} \models_M \mathbf{G}Z_0 \wedge \mathbf{G}FZ_1 \wedge \dots \wedge \mathbf{G}FZ_n$ iff $(\mathbf{x}, \dots, \mathbf{x}) \models_{M'} \mathbf{G}Z \wedge \mathbf{G}F\mathbf{X} \wedge \mathbf{G}F\mathbf{Y}$. Furthermore, zone Z is bounded iff zone Z_0 is bounded, and zones $\{X, Y\}$ are all bounded iff zones $\{Z_1, \dots, Z_n\}$ are all bounded.

Proof: We consider each $\mathbf{s} \in \mathbb{R}^{nd}$ as a sequence of n points from \mathbb{R}^d , i.e. $\mathbf{s} = (\mathbf{s}[1], \dots, \mathbf{s}[n])$. Formally, for all $\mathbf{s} \in \mathbb{R}^{nd}$ and $i \in [1..n]$, let $\mathbf{s}[i] := (\mathbf{s}((i-1) \cdot d + 1), \dots, \mathbf{s}(i \cdot d))$.

For each $\mathbf{m} \in M$, let $\mathbf{m}_i \in \mathbb{R}^{nd}$ be such that $\mathbf{m}_i[i] = \mathbf{m}$ and $\mathbf{m}_i[j] = \mathbf{0}$ for all $j \neq i$. Let,

$$\begin{aligned} M' &:= \{\mathbf{m}_i : \mathbf{m} \in M, i \in [1..n]\}, \\ Z &:= Z_0 \times Z_0 \times \dots \times Z_0, \\ X &:= Z_1 \times Z_2 \times \dots \times Z_n, \text{ and} \\ Y &:= \{\mathbf{y} \in \mathbb{R}^{nd} : \mathbf{y}[1] = \dots = \mathbf{y}[n] \in Z_1\}. \end{aligned}$$

Let $\varphi := \mathbf{G}Z_0 \wedge \mathbf{G}FZ_1 \wedge \dots \wedge \mathbf{G}FZ_n$ and $\varphi' := \mathbf{G}Z \wedge \mathbf{G}F\mathbf{X} \wedge \mathbf{G}F\mathbf{Y}$. It is the case that $\mathbf{x} \models_M \varphi$ iff $(\mathbf{x}, \dots, \mathbf{x}) \models_{M'} \varphi'$. \square

F. Model checking linear formulas

We may now prove Theorem 1, i.e. show that the model-checking problem for linear LTL formulas is in P.

Proof of Theorem 1: Let M be a d -dimensional MMS, let $\mathbf{x} \in \mathbb{R}^d$, and let ψ be a semi-bounded linear LTL formula. We recursively build a formula φ_ψ from the polynomial-time first-order logic of Section IV-A such that $\mathbf{x} \models_M \psi$ iff $\varphi_\psi(\mathbf{x})$.

For every $A \subseteq AP$, let $\text{zone}(A)$ denote the zone obtained by taking the intersection of the zones from A .

Case $\psi = A \wedge \psi'$. We take $\varphi_\psi(\mathbf{x}) := \mathbf{x} \in \text{zone}(A) \wedge \varphi_{\psi'}(\mathbf{x})$, which can be expressed as $\text{zone}(A)$ is represented by a system of inequalities.

Case $\psi = B \mathbf{U} (B' \wedge \psi')$. Note that $\mathbf{x} \models_M B \mathbf{U} B'$ almost amounts to $\mathbf{x} \rightarrow_{\text{zone}(B)}^* \mathbf{y} \in \text{zone}(B')$, except that, contrary to the former, the latter requires \mathbf{y} to be part of $\text{zone}(B)$.

In our case, we show that $\text{zone}(B') \subseteq \text{zone}(B)$. Recall that $\uparrow B \supseteq \uparrow B'$ by definition of linear LTL formulas. Let $\mathbf{z} \in \text{zone}(B')$. We have $\chi_{AP}(\mathbf{z}) \supseteq B'$ and $\chi_{AP}(\mathbf{z}) \in \uparrow B' \subseteq \uparrow B$. Thus, $\chi_{AP}(\mathbf{z}) \supseteq B$ and so $\mathbf{z} \in \text{zone}(B)$. Thus, we take

$$\varphi_\psi(\mathbf{x}) := \exists \lambda \geq \mathbf{0}, \mathbf{y} \in \text{zone}(B') : \varphi_{\text{zone}(B)}(\mathbf{x}, \lambda, \mathbf{y}) \wedge \varphi_{\psi'}(\mathbf{y}),$$

where φ_Z is the formula of Section IV-B with $Z := \text{zone}(B)$.

Case $\psi = (\mathbf{G}C_0) \wedge \bigwedge_{i=1}^n \mathbf{G}F C_i$. Let $Z_i := \text{zone}(C_i)$ for all $i \in [0..n]$. If $n = 0$, then we use formula $\varphi_{\mathbf{G}Z_0}$ from Section IV-D. If $n = 1$, then we artificially define $Z_2 := Z_1$. So, assume that $n \geq 2$. By Lemma 3, we can construct, in polynomial time, an nd -dimensional MMS M' and zones $X, Y, Z \subseteq \mathbb{R}^{nd}$ such that $\mathbf{x} \models_M \mathbf{G}Z_0 \wedge \mathbf{G}FZ_1 \wedge \dots \wedge \mathbf{G}FZ_n$ iff $(\mathbf{x}, \dots, \mathbf{x}) \models_{M'} \mathbf{G}Z \wedge \mathbf{G}F\mathbf{X} \wedge \mathbf{G}F\mathbf{Y}$. Furthermore, zones X and Y are bounded since $\{Z_1, \dots, Z_n\}$ are all bounded. Thus, we take $\varphi_\psi(\mathbf{x}) := \varphi_{\mathbf{G}Z \wedge \mathbf{G}F\mathbf{X} \wedge \mathbf{G}F\mathbf{Y}}((\mathbf{x}, \dots, \mathbf{x}))$ from Section IV-C for M' . \square

V. NP-COMPLETE FRAGMENTS

In this section, we establish the NP-completeness of fragments $\text{LTL}_B(\{\mathbf{F}, \mathbf{G}, \wedge\})$, $\text{LTL}_B(\{\mathbf{F}, \wedge, \vee\})$ and $\text{LTL}_B(\{\mathbf{F}, \wedge\})$.

A. Membership

Theorem 4: The fragment $\text{LTL}_B(\{\mathbf{F}, \mathbf{G}, \wedge\})$ belongs to NP.

Proof: Let M be a d -dimensional MMS, $\mathbf{x} \in \mathbb{R}^d$ and $\varphi \in \text{LTL}_B(\{\mathbf{F}, \mathbf{G}, \wedge\})$. Let $\text{sch}(M)$ denote the set of non-Zeno infinite schedules of M . Let $\text{lps}(\mathcal{A}_\varphi)$ denote the (finite) set of LPS of \mathcal{A}_φ starting from q_0 . Let $\text{lin}(S)$ be the linear LTL formula obtained from LPS S by Proposition 9. We have

$$\mathbf{x} \models \varphi$$

$$\iff \exists \pi \in \text{sch}(M) : \text{exec}(\pi, \mathbf{x}) \models \varphi$$

$$\iff \exists \pi \in \text{sch}(M), w \in \text{tr}(\text{exec}(\pi, \mathbf{x})) : w \models \varphi \quad (13)$$

$$\iff \exists \pi \in \text{sch}(M), w \in \text{tr}(\text{exec}(\pi, \mathbf{x})) : w \in L(\mathcal{A}_\varphi) \quad (14)$$

$$\iff \exists \pi \in \text{sch}(M), w \in \text{tr}(\text{exec}(\pi, \mathbf{x})), S \in \text{lps}(\mathcal{A}_\varphi) : w \models \text{lin}(S) \quad (15)$$

$$\iff \exists \pi \in \text{sch}(M), S \in \text{lps}(\mathcal{A}_\varphi) : \text{exec}(\pi, \mathbf{x}) \models \text{lin}(S) \quad (16)$$

$$\iff \exists S \in \text{lps}(\mathcal{A}_\varphi) : \mathbf{x} \models \text{lin}(S),$$

where

- (13) follows from Propositions 1 and 2,
- (14) follows from Proposition 7,
- (15) follows from Proposition 9 and from the fact that $L(\mathcal{A}_\varphi) = \bigcup_{S \in \text{lps}(\mathcal{A}_\varphi)} L(S)$,
- (16) follows from Propositions 1 and 2.

Thus, to check $\mathbf{x} \models \varphi$, we can nondeterministically construct a linear path scheme S of \mathcal{A}_φ , which is of linear size by Proposition 6, and the linear formula $\psi = \text{lin}(S)$, and check whether $\mathbf{x} \models \psi$ in polynomial time by Theorem 1. \square

Theorem 5: The fragment $\text{LTL}_B(\{\mathbf{F}, \wedge, \vee\})$ belongs to NP.

Proof: We give a nondeterministic polynomial-time procedure which, given an MMS M , \mathbf{x} and $\varphi \in \text{LTL}_B(\{\mathbf{F}, \wedge, \vee\})$, decides $\mathbf{x} \models_M \varphi$. For each disjunction $\varphi_1 \vee \varphi_2$ of φ , we nondeterministically replace $\varphi_1 \vee \varphi_2$ with either φ_1 or φ_2 . Let φ' denote the resulting formula. A simple induction shows that $\mathbf{x} \models_M \varphi$ iff $\mathbf{x} \models_M \varphi'$ for some such formula φ' . Moreover, φ' has no disjunctions and so $\varphi' \in \text{LTL}_B(\{\mathbf{F}, \wedge\}) \subseteq \text{LTL}_B(\{\mathbf{F}, \mathbf{G}, \wedge\})$. Therefore, deciding whether $\mathbf{x} \models_M \varphi'$ can be done in NP by Theorem 4. \square

B. NP-hardness

Theorem 6: Fragment $\text{LTL}_B(\{F, \wedge\})$ is strongly NP-hard.

Proof: We reduce from the rational variant of SUBSET-SUM [21], which asks, given $S \subseteq \mathbb{Q}$ and $t \in \mathbb{Q}$, whether some subset $V \subseteq S$ satisfies $\sum_{v \in V} v = t$. Given an instance where $S = \{s_1, \dots, s_n\}$, we give a $(4n + 1)$ -dimensional MMS M and a formula $\varphi \in \text{LTL}_B(\{F, \wedge\})$ such that $\mathbf{0} \models_M \varphi$ holds iff there is a solution for (S, t) .

A simple faulty approach is as follows. For each $s_i \in S$, we could associate the modes $\mathbf{y}_i = (0, \dots, 0, 1, 0, \dots, 0, s_i)$ and $\mathbf{n}_i = (0, \dots, 0, 1, 0, \dots, 0, 0)$, where “1” appears in dimension i . The goal would be to sum the modes in order to obtain $(1, \dots, 1, t)$. However, this is too naive. For example, consider $S = \{8, 9\}$ and $t = 4$. By taking $(0.5, \mathbf{y}_1)$ $(0.5, \mathbf{n}_1)$ $(1, \mathbf{n}_2)$, we obtain $(1, 1, 4)$ even though 4 cannot be obtained from S . We need a mechanism to ensure that, for each i , either \mathbf{y}_i or \mathbf{n}_i is used by exactly one unit. For this reason, we will introduce the additional modes $\bar{\mathbf{y}}_i$ and $\bar{\mathbf{n}}_i$, and zones Y_i , N_i and C_i . We will require that Y_i , N_i and C_i are all reached. As (partially) depicted in Figure 4, the only way to do so will be to either use schedule $(1, \mathbf{y}_i)$ $(1, \bar{\mathbf{y}}_i)$ or schedule $(1, \mathbf{n}_i)$ $(1, \bar{\mathbf{n}}_i)$. Moreover, the first zone reached, among Y_i and N_i , will determine whether $s_i \in S$ has been used.

Definition of M and φ : Let us now proceed. We will refer to the first $4n$ dimensions as $c_{i,1}, c_{i,2}, c_{i,3}, c_{i,4}$ for every $i \in [1..n]$, and to the remaining dimension as c^* . Intuitively, at the end of a satisfying execution, c^* will store the number t , which was derived by summing up elements from S . The other dimensions ensure that each element of S is added to c^* by a factor of 1 or 0, *i.e.* neither partially nor more than once.

For all $i \in [1..n]$, modes $\{\mathbf{y}_i, \mathbf{n}_i, \bar{\mathbf{y}}_i, \bar{\mathbf{n}}_i\}$ are defined by:

j	$\mathbf{y}_i(j)$	$\mathbf{n}_i(j)$	$\bar{\mathbf{y}}_i(j)$	$\bar{\mathbf{n}}_i(j)$
$c_{i,1}$	0.5	-0.5	-1	1
$c_{i,2}$	1	1	1	1
$c_{i,3}$	1	1	0	0
$c_{i,4}$	0	0	1	1
c^*	s_i	0	0	0
else	0	0	0	0

Let $\gamma := \max(2, |t|, n \cdot \max(|s_1|, \dots, |s_n|))$. For every $i \in [1..n]$, we define zones Y_i , N_i and C_i by the constraints:

	Y_i	N_i	C_i
$c_{i,1}$	$= 0.5$	$= -0.5$	$\in [-0.5, 0.5]$
$c_{i,2}$	$\in [1, 2]$	$\in [1, 2]$	$= 2$
$c_{i,3}$	$= 1$	$= 1$	$= 1$
$c_{i,4}$	$\in [0, 1]$	$\in [0, 1]$	$= 1$
else	$\in [-\gamma, \gamma]$	$\in [-\gamma, \gamma]$	$\in [-\gamma, \gamma]$

Let T be the zone $T := \{\mathbf{x} \in C_1 \cap \dots \cap C_n : \mathbf{x}(c^*) = t\}$. We define $\varphi := FT \wedge \bigwedge_{i=1}^n (FY_i \wedge FN_i)$.

Intuitively, the first zone that is reached among Y_i and N_i indicates whether number s_i is used in the solution to the SUBSET-SUM instance. Mode \mathbf{y}_i can be used to reach Y_i




Fig. 4: Schedules $(1, \mathbf{y}_i)$ $(1, \bar{\mathbf{y}}_i)$ and $(1, \mathbf{n}_i)$ $(1, \bar{\mathbf{n}}_i)$, where the x and y axes respectively correspond to $c_{i,1}$ and $c_{i,2}$.

first, and likewise with \mathbf{n}_i for N_i . Mode $\bar{\mathbf{y}}_i$ can be used to go from Y_i to N_i ; and likewise for $\bar{\mathbf{n}}_i$ for N_i to Y_i . See Figure 4.

Correctness: The proof appears in the full version. \square

VI. UNDECIDABLE FRAGMENTS

In this section, we show that $\text{LTL}_B(\{U\})$ and $\text{LTL}_B(\{G, \vee\})$ are undecidable, by reducing from the reachability problem for Petri nets with inhibitor arcs (*i.e.* zero-tests). A *Petri net with inhibitor arcs* is a tuple $\mathcal{N} = (P, T, \sim, \Delta)$ where

- P is a finite set of elements called *places*,
- T is a disjoint finite set of elements called *transitions*,
- $\sim : T \rightarrow \{\geq, =\}^P$, and
- $\Delta : T \rightarrow \mathbb{Z}^P$.

A transition t is *enabled* in $\mathbf{x} \in \mathbb{N}^P$ if $\mathbf{x} \sim_t \mathbf{0}$ and $\mathbf{x} + \Delta_t \geq \mathbf{0}$. If it is enabled, then its *firing* leads to $\mathbf{x}' := \mathbf{x} + \Delta_t$, denoted $\mathbf{x} \rightarrow^t \mathbf{x}'$. We write $\mathbf{x} \rightarrow \mathbf{x}'$ if $\mathbf{x} \rightarrow^t \mathbf{x}'$ for some t . We define \rightarrow^+ as the transitive closure of \rightarrow , and \rightarrow^* as the reflexive closure of \rightarrow^+ . The *reachability problem* asks, given a Petri nets with inhibitor arcs \mathcal{N} , and $\mathbf{x}_{\text{src}}, \mathbf{x}_{\text{tgt}}$, whether $\mathbf{x}_{\text{src}} \rightarrow^+ \mathbf{x}_{\text{tgt}}$ in \mathcal{N} . This problem is undecidable, *e.g.* see [22].

A. From Petri nets with inhibitor arcs to MMS

In this subsection, we will prove the following proposition through a series of intermediate propositions:

Proposition 24: Given a Petri net with inhibitor arcs \mathcal{N} and $\mathbf{x}_{\text{src}}, \mathbf{x}_{\text{tgt}}$, it is possible to compute an MMS M , two points \mathbf{x}, \mathbf{x}' , and a finite set of bounded zones AP such that

- 1) $\mathbf{x}_{\text{src}} \rightarrow^+ \mathbf{x}_{\text{tgt}}$ in \mathcal{N} iff $\mathbf{x} \rightarrow_{AP}^* \mathbf{x}'$ in M , and
- 2) no infinite non-Zeno schedule π satisfies $\mathbf{x} \rightarrow_{AP}^\pi$ in M .

Let $\mathcal{N} = (P, T, \sim, \Delta)$ be a Petri net with inhibitor arcs. We define a $(|P| + 3|T|)$ -dimensional MMS M together with zones $AP := \bigcup_{t \in T} \{A_t, A'_t, B_t, B'_t, C_t, C'_t\}$. We associate $|P|$ dimensions to P , which we collectively denote \mathbf{p} . Each transition $t \in T$ is associated to dimensions $\{t_A, t_B, t_C\}$.

Each transition $t \in T$ gives rise to modes $\{\mathbf{a}_t, \mathbf{b}_t, \mathbf{c}_t\}$. Informally, these three modes are respectively used to “request the firing of t ”, “fire t ” and “release the control on t ”. For

every $s \neq t$ and $I \in \{A, B, C\}$, we have $\mathbf{a}_t(s_I) = \mathbf{b}_t(s_I) = \mathbf{c}_t(s_I) := 0$. The rest of the values are defined as follows:

j	$\mathbf{a}_t(j)$	$\mathbf{b}_t(j)$	$\mathbf{c}_t(j)$
\mathbf{p}	$\mathbf{0}$	Δ_t	$\mathbf{0}$
t_A	-1	0	1
t_B	1	-1	0
t_C	0	1	-1

The six zones associated to $t \in T$ are defined by these constraints, where s stands for ‘‘any transition $s \neq t$ ’’:

	A_t	A'_t	B_t	B'_t	C_t	C'_t
\mathbf{p}	$\geq \mathbf{0}$	$\geq \mathbf{0}$	$\sim_t \mathbf{0}$	$\geq \mathbf{0}$	$\geq \mathbf{0}$	$\geq \mathbf{0}$
t_A	≥ 0	≥ 0	$= 0$	$= 0$	$= 0$	≥ 0
t_B	$= 0$	≥ 0	≥ 0	≥ 0	$= 0$	$= 0$
t_C	$= 0$	$= 0$	$= 0$	≥ 0	≥ 0	≥ 0
s_A	≥ 0	≥ 0	≥ 0	≥ 0	≥ 0	≥ 0
s_B	$= 0$	$= 0$	$= 0$	$= 0$	$= 0$	$= 0$
s_C	$= 0$	$= 0$	$= 0$	$= 0$	$= 0$	$= 0$

Since $A_t = A_s$ for all $s, t \in T$, we simply call this zone A . Informally, the MMS operates as follows:

- from A , we can take a mode \mathbf{a}_t , must go through A'_t , and end up in B_t after maximally taking \mathbf{a}_t ;
- in B_t , we test whether $\mathbf{p} \sim_t \mathbf{0}$ through the constraints;
- from B_t , we must take mode \mathbf{b}_t , go through B'_t and end up in C_t after maximally taking \mathbf{b}_t (adding Δ_t to \mathbf{p});
- from C_t , we must take mode \mathbf{c}_t , go through C'_t , and end up in A after maximally taking \mathbf{c}_t .

More formally, the following holds.

Lemma 4: Let $\mathbf{x}_A, \mathbf{x}'_A \in A$ and let π be a finite schedule such that $\mathbf{x}_A(t_A) = 1$, $|\pi| > 0$, and $\mathbf{x}_A \xrightarrow{\pi}_{AP} \mathbf{x}'_A$ holds with no intermediate points in A , i.e. $\text{exec}(\pi, \mathbf{x}_A)(\tau) \in A$ iff $\tau \in \{0, \text{time}(\pi)\}$. It is the case that $\pi \equiv \mathbf{a}_t \mathbf{b}_t \mathbf{c}_t$ and there exist $\mathbf{x}_B \in B_t$, $\mathbf{x}_C \in C_t$ such that

$$\mathbf{x}_A \xrightarrow{\mathbf{a}_t}_{A'_t} \mathbf{x}_B \xrightarrow{\mathbf{b}_t}_{B'_t} \mathbf{x}_C \xrightarrow{\mathbf{c}_t}_{C'_t} \mathbf{x}'_A.$$

A zone Z is *closed under scaling* if $\lambda \mathbf{z} \in Z$ holds for all $\lambda \in \mathbb{R}_{>0}$ and $\mathbf{z} \in Z$. A set of zones AP is closed under scaling if each $Z \in AP$ is closed under scaling. From Lemma 4, the following can be shown:

Proposition 25: Given a Petri net with inhibitor arcs \mathcal{N} and $\mathbf{x}_{\text{src}}, \mathbf{x}_{\text{tgt}}$, it is possible to compute an MMS M , points \mathbf{x}, \mathbf{x}' , and a finite set of zones AP closed under scaling, such that

- 1) $\mathbf{x}_{\text{src}} \xrightarrow{+} \mathbf{x}_{\text{tgt}}$ in \mathcal{N} iff $\mathbf{x} \xrightarrow{\pi}_{AP} \mathbf{x}'$ in M for some finite schedule π with $\text{time}(\pi) \geq 1$,
- 2) $\mathbf{0} \not\xrightarrow{+}_{AP} \mathbf{0}$ in M .

Proposition 26: Given a Petri net with inhibitor arcs \mathcal{N} and $\mathbf{x}_{\text{src}}, \mathbf{x}_{\text{tgt}}$, one can compute an MMS M , points \mathbf{x}, \mathbf{x}' , and a finite set of bounded zones AP such that $\mathbf{x}_{\text{src}} \xrightarrow{+} \mathbf{x}_{\text{tgt}}$ in \mathcal{N} iff $\lambda \mathbf{x} \xrightarrow{\pi}_{AP} \lambda \mathbf{x}'$ in M for some $\lambda \in [0, 1]$ and π with $\text{time}(\pi) = 1$.

Proof: Let $M, \mathbf{x}, \mathbf{x}'$ and AP be given by Proposition 25. Let $\gamma := \|\mathbf{x}\| + \|M\|$ and let d denote the dimension of M . We show the proposition with $AP' := \{Z \cap [-\gamma, \gamma]^d : Z \in AP\}$.

\Rightarrow) As $\mathbf{x}_{\text{src}} \xrightarrow{+} \mathbf{x}_{\text{tgt}}$ in \mathcal{N} , Proposition 25 yields $\mathbf{x} \xrightarrow{\pi}_{AP} \mathbf{x}'$ in M and $\text{time}(\pi) \geq 1$ for some finite schedule π . As AP is closed under scaling, we have $\lambda \mathbf{x} \xrightarrow{\lambda \pi}_{AP} \lambda \mathbf{x}'$ in M for any $\lambda \in \mathbb{R}_{>0}$. By picking $\lambda := 1/\text{time}(\pi)$, each point \mathbf{y} along the resulting execution satisfies $\|\mathbf{y}\| \leq \gamma$, and hence $\mathbf{y} \in Z'$ for some $Z' \in AP'$. Further, $\text{time}(\lambda \pi) = 1$.

\Leftarrow) Let $\lambda \mathbf{x} \xrightarrow{\pi}_{AP'} \lambda \mathbf{x}'$ with $\text{time}(\pi) = 1$. As π is nonempty, and $\mathbf{0} \not\xrightarrow{+}_{AP} \mathbf{0}$ by Proposition 25, we have $\lambda > 0$.

Let $Z \in AP$. If $\mathbf{y} \in Z \cap [-\gamma, \gamma]^d$, then in particular $\mathbf{y} \in Z$. Since Z is closed under scaling, we have $\lambda \mathbf{y} \in Z$. Thus,

$$\mathbf{x} = (1/\lambda) \cdot \lambda \mathbf{x} \xrightarrow{\frac{1}{\lambda} \pi}_{AP} (1/\lambda) \cdot \lambda \mathbf{x}' = \mathbf{x}' \text{ in } M.$$

By Proposition 25, this implies that $\mathbf{x}_{\text{src}} \xrightarrow{+} \mathbf{x}_{\text{tgt}}$ in \mathcal{N} . \square

We are now ready to prove Proposition 24.

Proof of Proposition 24: Let $M, \mathbf{x}, \mathbf{x}'$ and AP be given by Proposition 26. Let us define an MMS M' and a set of zones AP' . MMS M' has the same dimensions as M , plus four more: $\{\top, \vdash, \perp, \star\}$. The modes of M' are $\{\mathbf{a}_\top, \bar{\mathbf{a}}_\top\} \cup \{\mathbf{m}_\vdash : \mathbf{m} \in M\} \cup \{\mathbf{a}_\perp, \bar{\mathbf{a}}_\perp\}$. They are defined as follows:

j	$\mathbf{a}_\top(j)$	$\bar{\mathbf{a}}_\top(j)$	$\mathbf{m}_\vdash(j)$	$\mathbf{a}_\perp(j)$	$\bar{\mathbf{a}}_\perp(j)$
\top	-1	-1	0	0	0
\vdash	0	0	1	0	0
\perp	0	0	0	1	1
\star	1	0	0	-1	0
rest	\mathbf{x}	$\mathbf{0}$	\mathbf{m}	$-\mathbf{x}'$	$\mathbf{0}$

The set AP' contains two new zones, plus each zone from AP extended with the constraint $\top = \perp = 0$ and $\vdash, \star \in [0, 1]$. Since AP consists of bounded zones, we can extract $\gamma \in \mathbb{N}$ such that each dimension must remain within $[-\gamma, \gamma]$. We add zones $\{A_\top, A_\perp\}$ defined by these constraints:

	A_\top	A_\perp
\top	$\in [0, 1]$	$= 0$
\vdash	$= 0$	$= 1$
\perp	$= 0$	$\in [0, 1]$
\star	$\in [0, 1]$	$\in [0, 1]$
else	$\in [-\gamma, \gamma]$	$\in [-\gamma, \gamma]$

Informally, the MMS operates as follows:

- from A_\top , modes \mathbf{a}_\top and $\bar{\mathbf{a}}_\top$ empty \top to generate $\lambda \mathbf{x}$, and keep a copy of $\lambda \in [0, 1]$ in \star ;
- modes of M are used until the time reaches $\vdash = 1$;
- from A_\perp , modes \mathbf{a}_\perp and $\bar{\mathbf{a}}_\perp$ increase \perp to 1, in order to consume $\lambda \mathbf{x}'$, using \star to infer λ .

Formally, by definition of modes and zones, there exist $\lambda \in [0, 1]$ and $\lambda \mathbf{x} \xrightarrow{\pi}_{AP} \lambda \mathbf{x}'$ in M such that $\text{time}(\pi) = 1$ iff

there exist finite schedules ρ_{\top} and ρ_{\perp} , respectively using only modes $\{\mathbf{a}_{\top}, \overline{\mathbf{a}}_{\top}\}$ and $\{\mathbf{a}_{\perp}, \overline{\mathbf{a}}_{\perp}\}$, such that

$$\begin{pmatrix} 1 \\ 0 \\ 0 \\ 0 \\ 0 \end{pmatrix} \xrightarrow{\rho_{\top}} \begin{pmatrix} 0 \\ 0 \\ 0 \\ \lambda \\ \lambda \mathbf{x} \end{pmatrix} \xrightarrow{\pi_{\top}} \begin{pmatrix} 0 \\ 1 \\ 0 \\ \lambda \\ \lambda \mathbf{x}' \end{pmatrix} \xrightarrow{\rho_{\perp}} \begin{pmatrix} 0 \\ 1 \\ 1 \\ 0 \\ 0 \end{pmatrix}$$

in M' . The above holds iff $(1, 0, 0, 0, \mathbf{0}) \xrightarrow{*}_{AP'} (0, 1, 1, 0, \mathbf{0})$ in M' since zones enforce this ordering.

It remains to show Item 2. For the sake of contradiction, suppose there exists an infinite non-Zeno schedule π such that $\mathbf{x} \xrightarrow{\pi}_{AP}$. All zones of AP enforce $\top, \perp, \vdash, \star \in [0, 1]$. Thus, we obtain a contradiction since:

- If $\text{time}_{\mathbf{a}_{\top}}(\pi) + \text{time}_{\overline{\mathbf{a}}_{\top}}(\pi) = \infty$, then \top drops below 0;
- If $\sum_{\mathbf{m} \in M} \text{time}_{\mathbf{m}_{\vdash}}(\pi) = \infty$, then \vdash exceeds 1;
- If $\text{time}_{\mathbf{a}_{\perp}}(\pi) + \text{time}_{\overline{\mathbf{a}}_{\perp}}(\pi) = \infty$, then \perp exceeds 1. \square

B. Undecidability

We prove the undecidability of the fragments $\text{LTL}_{\text{B}}(\{\mathbf{U}\})$ and $\text{LTL}_{\text{B}}(\{\mathbf{G}, \vee\})$ using Proposition 24.

Lemma 5: Given $\psi_1, \dots, \psi_n, \varphi \in \text{LTL}_{\text{B}}(\{\mathbf{U}\})$, it is possible to compute a formula from $\text{LTL}_{\text{B}}(\{\mathbf{U}\})$ that is equivalent to formula $(\psi_1 \vee \dots \vee \psi_n) \mathbf{U} \varphi$.

Theorem 7: $\text{LTL}_{\text{B}}(\{\mathbf{U}\})$ and $\text{LTL}_{\text{B}}(\{\mathbf{G}, \vee\})$ are undecidable.

Proof: Let \mathcal{N} be a Petri net with inhibitor arcs and let $\mathbf{x}_{\text{src}}, \mathbf{x}_{\text{tgt}}$. Let $M, \mathbf{x}, \mathbf{x}'$ and AP be given by Proposition 24. Let $X' := \{\mathbf{x}'\}$ and $\psi := (\bigvee_{Z \in AP} Z) \mathbf{U} X'$. By Lemma 5, we can compute a formula $\varphi \in \text{LTL}_{\text{B}}(\{\mathbf{U}\})$ with $\varphi \equiv \psi$. By Proposition 24, we have $\mathbf{x}_{\text{src}} \xrightarrow{+} \mathbf{x}_{\text{tgt}}$ in \mathcal{N} iff $\mathbf{x} \xrightarrow{*}_{AP} \mathbf{x}'$ in M iff $\mathbf{x} \models_M \psi$ iff $\mathbf{x} \models_M \varphi$.

The proof for $\text{LTL}_{\text{B}}(\{\mathbf{G}, \vee\})$ is essentially the same, but requires an extra “dummy dimension” that can be increased and decreased once (and only once) \mathbf{x}' is reached. \square

VII. CONCLUSION

We have introduced a linear temporal logic for MMS and established the complexity of model checking for each syntactic fragments: Each one is either P-complete, NP-complete or undecidable. This generalizes and unifies existing work on MMS and continuous vector addition systems/Petri nets.

Future work includes fully dealing with unbounded zones; allowing for time constraints on temporal operators; and algorithmically optimizing objective functions on schedules satisfying a given LTL specification. It would also be interesting to go from theory to practice by providing a solver for linear LTL formulas, and more generally $\text{LTL}_{\text{B}}(\{\mathbf{F}, \mathbf{G}, \wedge\})$.

REFERENCES

- [1] R. Alur, A. Trivedi, and D. Wojtczak, “Optimal scheduling for constant-rate multi-mode systems,” in *Proc. 15th ACM International Conference on Hybrid Systems: Computation and Control (HSCC)*, 2012, pp. 75–84.
- [2] S. N. Krishna, A. Kumar, F. Somenzi, B. Touri, and A. Trivedi, “The reach-avoid problem for constant-rate multi-mode systems,” in *Proc. 15th International Symposium Automated Technology for Verification and Analysis (ATVA)*. Springer, 2017, pp. 463–479.
- [3] R. Koymans, “Specifying real-time properties with metric temporal logic,” *Real-time systems*, vol. 2, no. 4, pp. 255–299, 1990.
- [4] S. M. German and A. P. Sistla, “Reasoning about systems with many processes,” *Journal of the ACM*, vol. 39, no. 3, pp. 675–735, 1992.
- [5] J. Esparza, P. Ganty, J. Leroux, and R. Majumdar, “Verification of population protocols,” *Acta Informatica*, vol. 54, no. 2, pp. 191–215, 2017.
- [6] M. Heiner, D. R. Gilbert, and R. Donaldson, “Petri nets for systems and synthetic biology,” in *Formal Methods for Computational Systems Biology*, 2008, pp. 215–264.
- [7] W. van der Aalst, “The application of Petri nets to workflow management,” *Journal of Circuits, Systems, and Computers*, vol. 8, no. 1, pp. 21–66, 1998.
- [8] J. Esparza, R. Ledesma-Garza, R. Majumdar, P. J. Meyer, and F. Nikšić, “An SMT-based approach to coverability analysis,” in *Proc. 26th International Conference on Computer Aided Verification (CAV)*, vol. 8559, 2014, pp. 603–619.
- [9] M. Blondin, C. Haase, and P. Offtermatt, “Directed reachability for infinite-state systems,” in *Proc. 27th International Conference on Tools and Algorithms for the Construction and Analysis of Systems (TACAS)*, 2021, pp. 3–23.
- [10] E. Fraca and S. Haddad, “Complexity analysis of continuous Petri nets,” *Fundamenta Informaticae*, vol. 137, no. 1, pp. 1–28, 2015.
- [11] J. Křetínský and J. Esparza, “Deterministic automata for the (F, G)-fragment of LTL,” in *Proc. 24th International Conference on Computer Aided Verification (CAV)*, 2012, pp. 7–22.
- [12] M. Blondin and C. Haase, “Logics for continuous reachability in Petri nets and vector addition systems with states,” in *Proc. 32nd Annual ACM/IEEE Symposium on Logic in Computer Science (LICS)*, 2017, pp. 1–12.
- [13] T. A. Henzinger, “The theory of hybrid automata,” in *Proc. 11th Annual IEEE Symposium on Logic in Computer Science (LICS)*. IEEE Computer Society, 1996, pp. 278–292.
- [14] K. Bae and J. Lee, “Bounded model checking of signal temporal logic properties using syntactic separation,” *Proc. ACM on Programming Languages (POPL)*, vol. 3, no. POPL, pp. 51:1–51:30, 2019.
- [15] R. Alur, “Timed automata,” in *Proc. International Conference on Computer Aided Verification (CAV)*, 1999, pp. 8–22.
- [16] D. Bhave, S. Jha, S. N. Krishna, S. Schewe, and A. Trivedi, “Bounded-rate multi-mode systems based motion planning,” in *Proc. 18th International Conference on Hybrid Systems: Computation and Control (HSCC)*. ACM, 2015, pp. 41–50.
- [17] R. Alur, V. Forejt, S. Moarref, and A. Trivedi, “Schedulability of bounded-rate multimode systems,” *ACM Transactions on Embedded Computing Systems (TECS)*, vol. 16, no. 3, 2017.
- [18] S. Demri and Philippe Schnoebelen, “The complexity of propositional linear temporal logics in simple cases,” *Information and Computation*, vol. 174, no. 1, pp. 84–103, 2002.
- [19] A. P. Sistla and E. M. Clarke, “The complexity of propositional linear temporal logics,” *Journal of the ACM*, vol. 32, no. 3, pp. 733–749, 1985.
- [20] R. Alur and S. L. Torre, “Deterministic generators and games for LTL fragments,” *ACM Transactions on Computational Logic (TOCL)*, vol. 5, no. 1, pp. 1–25, 2004.
- [21] D. Wojtczak, “On strong NP-completeness of rational problems,” in *Proc. 13th International Computer Science Symposium in Russia on Computer Science (CSR)*, 2018, pp. 308–320.
- [22] K. Reinhardt, “Reachability in Petri nets with inhibitor arcs,” *Electronic Notes in Theoretical Computer Science (ENTCS)*, vol. 223, pp. 239–264, 2008.
- [23] R. Greenlaw, H. J. Hoover, and W. L. Ruzzo, *Limits to parallel computation: P-completeness theory*. Oxford University Press, 1995.

APPENDIX

Before proving Proposition 1, which we recall shortly, let us prove the following technical lemma.

Lemma 6: Let σ be an execution and let $\tau, \tau' \in \text{dom } \sigma$ belong to a common interval I_j of σ . Let $X \subseteq \mathbb{R}^d$ be a convex set. If $\tau'' \in [\tau, \tau']$ and $\sigma(\tau), \sigma(\tau') \in X$, then $\sigma(\tau'') \in X$.

Proof: Let $\lambda \in [0, 1]$ be such that $\tau'' = \lambda \cdot \tau + (1 - \lambda) \cdot \tau'$. Let $b := \min I_j$, $c := \max I_j - \min I_j$ and $\mathbf{y} := \mathbf{x}_{j+1} - \mathbf{x}_j$. We have:

$$\begin{aligned} & \sigma(\tau'') \\ &= \mathbf{x}_j + (\tau'' - b)/c \cdot \mathbf{y} \\ &= \mathbf{x}_j + (\lambda\tau + (1 - \lambda)\tau' - b)/c \cdot \mathbf{y} \\ &= \lambda(\mathbf{x}_j + (\tau - b)/c \cdot \mathbf{y}) + (1 - \lambda)(\mathbf{x}_j + (\tau' - b)/c \cdot \mathbf{y}) \\ &= \lambda\sigma(\tau) + (1 - \lambda)\sigma(\tau') \\ &\in X. \end{aligned}$$

Proposition 1: Any execution σ has a trace. □

Proof: Let $\sigma = \mathbf{x}_0 I_0 \mathbf{x}_1 \dots$. Let $\tau_0 := 0$. We construct the rest of the sequence inductively. Let $X_i := \{\tau' \in \text{dom } \sigma : \tau' > \tau_{i-1} \text{ and } \chi_{AP}(\tau') \neq \chi_{AP}(\tau_{i-1})\}$. Let us make a case distinction.

X_i is empty. If $\tau_{i-1} = \text{sup dom } \sigma$, then the process ends. Otherwise, we add $\tau_i := \max I_j$, where $j \in \mathbb{N}$ is the maximal index such that $\tau_{i-1} \in I_j$.

$\inf X_i > \tau_{i-1}$. We add $\tau_i := \inf X_i$, as it satisfies $\chi_{AP}(\sigma(\tau')) = \chi_{AP}(\sigma(\tau_{i-1}))$ for all $\tau' \in [\tau_{i-1}, \tau_i]$.

$\inf X_i = \tau_{i-1}$. Let j be the maximal index such that $\tau_{i-1} \in I_j$. There is a sequence $\alpha_0 > \alpha_1 > \dots \in X_i \cap I_j$ that converges to τ_{i-1} . By Lemma 6 and convexity of zones, there exists $k \in \mathbb{N}$ such that $\chi_{AP}(\sigma(\tau')) = \chi_{AP}(\sigma(\alpha_k))$ for all $\tau' \in (\tau_{i-1}, \alpha_k]$. Thus, we add $\tau_i := \alpha_k$ to the sequence.

It remains to show that $\text{dom } \sigma = [\tau_0, \tau_1] \cup [\tau_1, \tau_2] \cup \dots$. For the sake of contradiction, suppose that τ_0, τ_1, \dots is infinite and converges to some $\alpha \leq \text{sup dom } \sigma$. Let $j \in \mathbb{N}$ be the minimal index such that $\alpha \in I_j$. There exists $k \in \mathbb{N}$ such that $\tau_k, \tau_{k+1}, \dots \in I_j$. By Lemma 6 and convexity of zones, there exists $\ell \geq k$ such that $\chi_{AP}(\sigma(\tau')) = \chi_{AP}(\sigma(\alpha))$ holds for all $\tau' \in [\tau_\ell, \alpha]$. This means that either $X_{\ell+1} = \emptyset$ or $\inf X_{\ell+1} \geq \alpha$. In both cases, this means that $\tau_{\ell+1} \geq \alpha$. If $\tau_{\ell+1} > \alpha$, then this contradicts $\lim_{i \rightarrow \infty} \tau_i = \alpha$. If $\tau_{\ell+1} = \alpha$, then we must have $\tau_{\ell+1} = \text{sup dom } \sigma$, which contradicts the fact that the sequence is infinite. □

Before proving Proposition 2, which we recall shortly, let us prove the following technical lemma.

Lemma 7: Let φ be a negation-free LTL formula, let σ be an execution with $\text{dom } \sigma = \mathbb{R}_{\geq 0}$, let $\tau, \tau' \in \mathbb{R}_{\geq 0}$ and let $T_\varphi := \{\tau'' \in [\tau, \tau'] : \sigma, \tau'' \models \varphi\}$. If the set T_φ is nonempty, then it has a minimum.

Proof: We proceed by induction on the structure of φ .

Case $\varphi = \text{true}$. We trivially have $\min T_\varphi = \tau$.

Case $\varphi = Z \in AP$. Informally, the claim follows from the fact that Z is defined as the intersection of *closed* half-spaces, and hence σ must intersect with a face of Z .

Formally, let $\sigma = \mathbf{x}_0 I_0 \mathbf{x}_1 \dots$ and let Z be defined by the system $\mathbf{A}\mathbf{x} \leq \mathbf{b}$. Let $\alpha := \inf T_Z$. If $\sigma(\alpha) \in Z$, then we are done. For the sake of contradiction, assume that this is not the case. We have $(\mathbf{A} \cdot \sigma(\alpha))(\ell) > \mathbf{b}(\ell)$ for some ℓ .

Let $\alpha_0 > \alpha_1 > \dots$ be a sequence from T_Z that converges to α . For every $i \in \mathbb{N}$, let $j_i \in \mathbb{N}$ be the last index such that $\alpha_i \in I_{j_i}$. Let $j \in \mathbb{N}$ be the last index such that $\alpha \in I_j$. There exists $k \in \mathbb{N}$ such that $j_k = j_{k+1} = \dots = j$. Let $b := \min I_j$, $c := \max I_j - \min I_j$ and $\mathbf{y} := \mathbf{x}_{j+1} - \mathbf{x}_j$. By definition of an execution, we have

$$\mathbf{A} \cdot \sigma(\alpha) = \mathbf{A}\mathbf{x}_j + \frac{\alpha - b}{c} \cdot \mathbf{A}\mathbf{y}.$$

Moreover, for every $k' \geq k$, we have

$$\mathbf{A} \cdot \sigma(\alpha_{k'}) = \mathbf{A} \cdot \sigma(\alpha) + \frac{\alpha_{k'} - \alpha}{c} \cdot \mathbf{A}\mathbf{y}.$$

For every $k' \geq k$, we have $\alpha_{k'} > \alpha$ and $(\mathbf{A} \cdot \sigma(\alpha_{k'}))(\ell) \leq \mathbf{b}(\ell)$. Hence, $(\mathbf{A}\mathbf{y})(\ell) < 0$. Let $k' \geq k$ be sufficiently large so that $\alpha_{k'} - \alpha$ is small enough for the following to hold:

$$(\mathbf{A} \cdot \sigma(\alpha))(\ell) + \frac{\alpha_{k'} - \alpha}{c} \cdot (\mathbf{A}\mathbf{y})(\ell) > \mathbf{b}(\ell).$$

We obtain $\sigma(\alpha_{k'}) \notin Z$, which is a contradiction.

Case $\varphi = \psi \wedge \psi'$. Let $\alpha := \inf T_\varphi$, and let

$$\begin{aligned} B &:= \{\alpha' \in [\alpha, \tau'] : \sigma, \alpha' \models \psi\}, \\ B' &:= \{\alpha' \in [\alpha, \tau'] : \sigma, \alpha' \models \psi'\}. \end{aligned}$$

As $T_\varphi \neq \emptyset$, both B and B' are nonempty. Thus, by induction, $\beta := \min B$ and $\beta' := \min B'$ are well-defined. We must have $\inf B = \inf B' = \inf T_\varphi$, since $\varphi = \psi \wedge \psi'$. Thus, $\min B = \min B' = \alpha$, which means that $\min T_\varphi = \alpha$.

Case $\varphi = \psi \vee \psi'$. Follows from $T_\varphi = T_\psi \cup T_{\psi'}$ and induction.

Case $\varphi = \psi \mathbf{U} \psi'$. Let $\alpha := \inf T_\varphi$ and let

$$T' := \{\alpha' \in [\alpha, \tau'] : \sigma, \alpha' \models \psi'\},$$

Since $T_\varphi \neq \emptyset$, we must have $T' \neq \emptyset$. By induction hypothesis, $\beta := \min T'$ is well-defined. If $\beta = \alpha$, then we are done as $\sigma, \alpha \models \psi \mathbf{U} \psi'$ and hence $\min T_\varphi = \alpha$.

Otherwise, $\alpha < \beta$. We claim that $\sigma, \gamma \models \psi$ for all $\gamma \in (\alpha, \beta)$. By induction on $\{\alpha' \in [\alpha, \beta) : \sigma, \alpha' \models \psi\}$, the claim implies $\min T_\varphi = \alpha$. It remains to show the claim. Let $\gamma \in (\alpha, \beta)$. As $\alpha = \inf T_\varphi$, there is $\gamma' \in (\alpha, \gamma)$ such that $\sigma, \gamma' \models \psi \mathbf{U} \psi'$. By minimality of β , we must have $\sigma, \delta \models \psi$ for all $\delta \in [\gamma', \beta)$. Thus, $\sigma, \gamma' \models \psi \mathbf{U} \psi'$, and so $\sigma, \gamma \models \psi \mathbf{U} \psi'$.

Case $\varphi = F\psi$. Follows from $F\psi \equiv \text{true} \mathbf{U} \psi$.

Case $\varphi = G\psi$. Let $\alpha := \inf T_\varphi$. It is the case that $\sigma, \beta \models G\psi$ for infinitely many $\beta \in (\alpha, \tau')$ arbitrarily closer to α . Thus, we have $\sigma, \beta \models \psi$ for all $\beta > \alpha$. By induction hypothesis, the set $\{\beta \in [\alpha, \tau'] : \sigma, \beta \models \psi\}$ has a minimum, which must be α . Hence, $\sigma, \alpha \models G\psi$, which means that $\min T_\varphi = \alpha$. □

Proposition 2: Let σ be an execution with $\text{dom } \sigma = \mathbb{R}_{\geq 0}$, let w be a trace of σ , and let φ be a negation-free LTL formula. It is the case that $\sigma \models \varphi$ iff $w \models \varphi$.

Proof: Let $\tau_0, \tau_1, \dots \in \mathbb{R}_{\geq 0}$ yield w . Let $f: \mathbb{R}_{\geq 0} \rightarrow \mathbb{N}$ be the function that satisfies the following, for all $\tau \in [\tau_i, \tau_{i+1})$:

$$f(\tau) := \begin{cases} i & \text{if } \chi_{AP}(\sigma(\tau)) = \chi_{AP}(\sigma(\tau_i)), \\ i+1 & \text{otherwise.} \end{cases}$$

By definition of traces, f is non-decreasing. We show that $\sigma, \tau \models \varphi$ iff $w, f(\tau) \models \varphi$ by induction on the structure of φ .

Case $\varphi = Z \in AP$. We have $w(i) = \chi_{AP}(\sigma(\tau_i))$ for all $i \in \mathbb{N}$. Thus, $\sigma, \tau \models \varphi$ iff $\sigma(\tau) \in Z$ iff $Z \in w(f(\tau))$ iff $w, f(\tau) \models \varphi$.

Case $\varphi = \psi \wedge \psi'$. We have $\sigma, \tau \models \varphi$ iff $\sigma, \tau \models \psi \wedge \sigma, \tau \models \psi'$ iff $w, f(\tau) \models \psi \wedge w, f(\tau) \models \psi'$ iff $w, f(\tau) \models \varphi$.

Case $\varphi = \psi \vee \psi'$. Symmetric to \wedge .

Case $\varphi = \psi \mathbf{U} \psi'$. \Rightarrow) Since $\sigma, \tau \models \varphi$, there exists $\tau' \geq \tau$ such that $\sigma, \tau' \models \psi'$, and $\sigma, \tau'' \models \psi$ for all $\tau'' \in [\tau, \tau')$. By induction hypothesis, we have $w, f(\tau') \models \psi'$, and $w, f(\tau'') \models \psi$ for all $\tau'' \in [\tau, \tau')$. By definition of f , we have

$$[f(\tau)..f(\tau') - 1] \subseteq \{f(\tau'') : \tau'' \in [\tau, \tau')\}.$$

Thus, $w, j \models \psi$ holds for all $j \in [f(\tau)..f(\tau') - 1]$. This means that $w, f(\tau) \models \psi \mathbf{U} \psi'$.

\Leftarrow) Since $w, f(\tau) \models \varphi$, there is a minimal $i \geq f(\tau)$ such that $w, i \models \psi'$, and $w, j \models \psi$ for all $j \in [f(\tau)..i - 1]$.

If $f(\tau) = i$, then we are done by the induction hypothesis. Thus, we assume $f(\tau) \neq i$. We claim that $\tau < \tau_i$. Indeed, if $\tau_i \leq \tau$ would hold, then, as f is non-decreasing, we would have $i = f(\tau_i) \leq f(\tau) \leq i$, which contradicts $f(\tau) \neq i$.

By induction hypothesis, it is the case that $\sigma, \tau_i \models \psi'$. Let $\tau' \in [\tau, \tau_i)$. It remains to show that $\sigma, \tau' \models \psi$. We have $f(\tau) \leq f(\tau') \leq f(\tau_i) = i$ as f is non-decreasing. So, $f(\tau') \in [f(\tau)..i]$. If $f(\tau') \leq i - 1$, then we are done by the induction hypothesis. Therefore, suppose that $f(\tau') = i$. Let

$$A := \{\alpha \in [\tau_{i-1}, \tau_i) : \sigma, \alpha \models \psi'\}.$$

As $\tau' < \tau_i$, the definition of f yields $f(\alpha) = i$ for all $\alpha \in (\tau_{i-1}, \tau_i)$. So, by induction hypothesis, we have $\inf A = \tau_{i-1}$, and hence $\min A = \tau_{i-1}$ by Lemma 7. Thus, $\sigma, \tau_{i-1} \models \psi'$, and so $w, i - 1 \models \psi'$, which contradicts the minimality of i .

Case $\varphi = F\psi$. Follows from $F\psi \equiv \text{true} \mathbf{U} \psi$.

Case $\varphi = G\psi$. We have

$$\begin{aligned} & \sigma, \tau \models G\psi \\ \iff & \sigma, \tau' \models \psi \text{ for all } \tau' \geq \tau \\ \iff & w, f(\tau') \models \psi \text{ for all } \tau' \geq \tau \quad (\text{by ind. hyp.}) \\ \iff & w, i \models \psi \text{ for all } i \geq f(\tau) \quad (\text{by def. of } f) \\ \iff & w, f(\tau) \models G\psi. \end{aligned}$$

Recall that in the forthcoming proposition, we have

$$\varphi = \psi \wedge \bigwedge_{i \in I} G\varphi_i \wedge \bigwedge_{j \in J} F\varphi_j.$$

Proposition 3: It is the case that $\text{flat}(\varphi) \equiv \varphi$.

Proof: We claim that the following equivalences hold:

- 1) $G\varphi \equiv G\psi \wedge \bigwedge_{i \in I} G\varphi_i \wedge \bigwedge_{j \in J} GF\varphi_j$,
- 2) $GF\varphi \equiv GF\psi \wedge \bigwedge_{i \in I} FG\varphi_i \wedge \bigwedge_{j \in J} GF\varphi_j$,
- 3) $FG\varphi \equiv FG\psi \wedge \bigwedge_{i \in I} FG\varphi_i \wedge \bigwedge_{j \in J} GF\varphi_j$.

By a routine induction, these equivalences yield $\text{flat}_G(\varphi) \equiv G\varphi$, $\text{flat}_{GF}(\varphi) \equiv GF\varphi$, $\text{flat}_{FG}(\varphi) \equiv FG\varphi$ and $\text{flat}(\varphi) \equiv \varphi$. It remains to prove the equivalences.

- 1) It follows from distributivity of G over \wedge , and idempotence of G .
- 2) \Rightarrow) Let $w \models GF\varphi$. Let $i_0 < i_1 < \dots \in \mathbb{N}$ be such that $w, i_k \models \varphi$ for every $k \in \mathbb{N}$. Recall that $\varphi = \psi \wedge \bigwedge_{i \in I} G\varphi_i \wedge \bigwedge_{j \in J} F\varphi_j$. Clearly, we have $w \models GF\psi$ and $w \models \bigwedge_{j \in J} GF\varphi_j$. Moreover, since $w, i_0 \models \bigwedge_{i \in I} G\varphi_i$, we have $w \models \bigwedge_{i \in I} FG\varphi_i$.
 \Leftarrow) Let $w \models GF\psi \wedge \bigwedge_{i \in I} FG\varphi_i \wedge \bigwedge_{j \in J} GF\varphi_j$. Let $i_0 \in \mathbb{N}$ satisfy $w, i_0 \models GF\psi \wedge \bigwedge_{i \in I} G\varphi_i \wedge \bigwedge_{j \in J} GF\varphi_j$. There exist $i_0 < i_1 < \dots \in \mathbb{N}$ such that $w, i_k \models \psi \wedge \bigwedge_{i \in I} G\varphi_i \wedge \bigwedge_{j \in J} F\varphi_j$ for every $k \geq 1$. Thus, $w \models GF\varphi$.
- 3) It follows from distributivity of G and FG over \wedge , and idempotence of G . \square

Proposition 4: Let $\varphi \in \text{LTL}(\{F, G, \wedge\})$. This holds:

- 1) if φ is flat, then $|\varphi|_F > |\varphi'|_F$ for all $\varphi' \in \mathfrak{L}(\varphi) \setminus \{\varphi\}$,
- 2) if φ is flat, then $|\varphi|_F \geq |\varphi[A]|_F$ for all $A \subseteq AP$,
- 3) $|\varphi| \geq |\text{flat}(\varphi)|_F$.

We prove Items 1 and 3. Item 2 follows easily by definition.

Proof of Proposition 4(1): We proceed by induction on the structure of the flat formula.

True, atomic propositions and operator G. It follows trivially from $\mathfrak{L}(\varphi) \setminus \{\varphi\} = \emptyset$.

Conjunction. Let $\varphi' \in \mathfrak{L}(\varphi_1 \wedge \varphi_2) \setminus \{\varphi_1 \wedge \varphi_2\}$. By definition,

$$\varphi' \in \{\psi_1 \wedge \psi_2 : \psi_1 \in \mathfrak{L}(\varphi_1), \psi_2 \in \mathfrak{L}(\varphi_2)\} \setminus \{\varphi_1 \wedge \varphi_2\}.$$

Thus, we have $\varphi' = \psi_1 \wedge \psi_2$ where $\psi_1 \in \mathfrak{L}(\varphi_1)$, $\psi_2 \in \mathfrak{L}(\varphi_2)$, and either $\psi_1 \neq \varphi_1$ or $\psi_2 \neq \varphi_2$. Let us consider the first case. The second one is symmetric. We have:

$$\begin{aligned} |\varphi'|_F &= |\psi_1 \wedge \psi_2|_F \\ &= |\psi_1|_F + |\psi_2|_F \\ &< |\varphi_1|_F + |\psi_2|_F \quad (\text{by ind. hyp. as } \psi_1 \neq \varphi_1) \\ &\leq |\varphi_1|_F + |\varphi_2|_F \quad (\text{by ind. hyp. or } \psi_2 = \varphi_2) \\ &= |\varphi_1 \wedge \varphi_2|_F. \end{aligned}$$

Operator F. Let $\varphi' \in \mathfrak{L}(F\varphi) \setminus \{F\varphi\}$. By definition, it is the case that $\varphi' \in \mathfrak{L}(\varphi)$. We have $|F\varphi|_F > |\varphi'|_F$ since

$$\begin{aligned} |F\varphi|_F &= 1 + |\varphi|_F \\ &\geq 1 + |\varphi'|_F \quad (\text{by ind. hyp. or } \varphi' = \varphi). \end{aligned} \quad \square$$

Proof of Proposition 4(3): We prove this by structural induction: $|\text{flat}(\varphi)|_F \leq |\varphi|$, $|\text{flat}_G(\varphi)|_F \leq |\varphi|$, $|\text{flat}_{GF}(\varphi)|_F \leq |\varphi|$ and $|\text{flat}_{FG}(\varphi)|_F \leq |\varphi|$.

Let φ be of the form $\psi \wedge \bigwedge_{i \in I} G\varphi_i \wedge \bigwedge_{j \in J} F\varphi_j$.

Case **G**. We have

$$\begin{aligned} |\text{flat}_G(\varphi)|_F &= \sum_{i \in I} |\text{flat}_G(\varphi_i)|_F + \sum_{j \in J} |\text{flat}_{GF}(\varphi_j)|_F \\ &\leq \sum_{i \in I} |\varphi_i| + \sum_{j \in J} |\varphi_j| && \text{(by ind.)} \\ &\leq |\varphi|. \end{aligned}$$

Case **GF**. We have

$$\begin{aligned} |\text{flat}_{GF}(\varphi)|_F &= \sum_{i \in I} |\text{flat}_{FG}(\varphi_i)|_F + \sum_{j \in J} |\text{flat}_{GF}(\varphi_j)|_F \\ &\leq \sum_{i \in I} |\varphi_i| + \sum_{j \in J} |\varphi_j| && \text{(by ind.)} \\ &\leq |\varphi|. \end{aligned}$$

Case **FG**. We have

$$\begin{aligned} |\text{flat}_{FG}(\varphi)|_F &= 1 + \sum_{i \in I} |\text{flat}_{FG}(\varphi_i)|_F + \sum_{j \in J} |\text{flat}_{GF}(\varphi_j)|_F \\ &\leq |\psi| + \sum_{i \in I} |\varphi_i| + \sum_{j \in J} |\varphi_j| \\ &\leq |\varphi|. \end{aligned}$$

General case. We have

$$\begin{aligned} |\text{flat}(\varphi)|_F &= |\psi|_F + \sum_{i \in I} |\text{flat}_G(\varphi_i)|_F + \sum_{j \in J} |\text{flat}(\varphi_j)|_F \\ &\leq |\psi| + \sum_{i \in I} |\varphi_i| + \sum_{j \in J} |\varphi_j| && \text{(by ind.)} \\ &\leq |\varphi|. \quad \square \end{aligned}$$

Proposition 27: Let $\varphi \in \text{LTL}(\{F, G, \wedge\})$ be flat, and let $A, A' \subseteq AP$. It is the case that $\varphi[A][A'] \in \{\varphi[A], \text{false}\}$.

Proof: If $\varphi[A] = \text{false}$, then $\varphi[A][A'] = \text{false}$. Otherwise, by definition, $\varphi[A][A']$ is either $\varphi[A]$ or false . \square

Proposition 5: Let $r_0 \xrightarrow{A_1} r_1 \xrightarrow{A_2} \dots \xrightarrow{A_n} r_n$ be a simple path of \mathcal{A}_φ . It is the case that $|r_1|_F > \dots > |r_n|_F$.

Proof: Let $i \in [1..n-1]$. By definition of \mathcal{A}_φ , there exist $\psi_i \in \mathfrak{U}(r_{i-1})$ and $\psi_{i+1} \in \mathfrak{U}(r_i)$ such that $r_i = \psi_i[A_i] \neq \text{false}$ and $r_{i+1} = \psi_{i+1}[A_{i+1}] \neq \text{false}$. There are two cases.

Case $\psi_{i+1} = r_i$. We have

$$r_{i+1} = \psi_{i+1}[A_{i+1}] = r_i[A_{i+1}] = \psi_i[A_i][A_{i+1}].$$

As $r_{i+1} \neq \text{false}$, it is the case that $\psi_i[A_i][A_{i+1}] = \psi_i[A_i] = r_i$ by Proposition 27. Therefore, $r_i = r_{i+1}$, which contradicts the fact that the path is simple.

Case $\psi_{i+1} \neq r_i$. We have $|r_{i+1}|_F = |\psi_{i+1}[A_{i+1}]|_F \leq |\psi_{i+1}|_F$ by Item 2 of Proposition 4. Moreover, we have $|\psi_{i+1}|_F < |r_i|_F$ by Item 1 of Proposition 4, since $\psi_{i+1} \in \mathfrak{U}(r_i) \setminus \{r_i\}$. \square

Lemma 1: Let $\varphi \in \text{LTL}(\{F, G, \wedge\})$ be a flat formula. These two properties are equivalent to $w \models \varphi$:

- 1) there exists φ' such that $\varphi \xrightarrow{w(0)} \varphi'$ and $w[1..] \models \varphi'$;
- 2) there exist $i \in \mathbb{N}$ and φ' such that $\varphi \xrightarrow{w(0)\dots w(i-1)} \varphi'$, $|\varphi'|_F = 0$ and $w[i..] \models \varphi'$.

Proof of Item 1: We proceed by induction on $|\varphi|_F$.

If $|\varphi|_F = 0$, then φ is of the form $\psi \wedge G\psi' \wedge \bigwedge_{i \in I} GF\psi''_i$. Let $\varphi' := \varphi[w(0)]$. Note that $\mathfrak{U}(\varphi) = \{\varphi\}$. We have

$$w \models \varphi$$

$$\iff \text{prop}(\psi \wedge \psi') \subseteq w(0) \wedge w[1..] \models (G\psi' \wedge \bigwedge_{i \in I} GF\psi''_i)$$

$$\iff \varphi[w(0)] \neq \text{false} \wedge w[1..] \models \varphi[w(0)]$$

$$\iff \varphi \xrightarrow{w(0)} \varphi' \wedge w[1..] \models \varphi'.$$

Now, assume that $\varphi = \theta \wedge F\psi$ where θ and ψ are flat.

\Rightarrow) Let $w \models \varphi$. We have $w \models \theta$ and $w \models F\psi$. By induction hypothesis, there exists θ' such that $\theta \xrightarrow{w(0)} \theta'$ and $w[1..] \models \theta'$. By definition of $\mathfrak{U}(\cdot)$, we have $\theta \wedge F\psi \xrightarrow{w(0)} \theta' \wedge F\psi$. So, if $w[1..] \models F\psi$, then we are done by taking $\varphi' := \theta' \wedge F\psi$. Otherwise, we must have $w \models \psi$. By induction hypothesis, there exists ψ' such that $\psi \xrightarrow{w(0)} \psi'$ and $w[1..] \models \psi'$. Since $\mathfrak{U}(\psi) \subseteq \mathfrak{U}(F\psi)$, we have $F\psi \xrightarrow{w(0)} \psi'$, and hence $\theta \wedge F\psi \xrightarrow{w(0)} \theta' \wedge \psi'$. So, we are done by taking $\varphi' := \theta' \wedge \psi'$.

\Leftarrow) Let φ' satisfy $\varphi \xrightarrow{w(0)} \varphi'$ and $w[1..] \models \varphi'$. By definition of \rightarrow , there exist θ', ψ' such that $\theta \xrightarrow{w(0)} \theta'$, $F\psi \xrightarrow{w(0)} \psi'$ and $\varphi' = \theta' \wedge \psi'$. Thus, $w[1..] \models \theta'$ and $w[1..] \models \psi'$. By induction hypothesis, we have $w \models \theta$. It remains to show that $w \models F\psi$.

If $\psi' \neq F\psi$, then $|\psi'|_F < |F\psi|_F$ by Proposition 4, and hence $w \models F\psi$ by induction hypothesis. Otherwise, we have $\psi' = F\psi$ and so $w[1..] \models F\psi$, and in particular $w \models F\psi$. \square

Proof of Item 2: We proceed by induction on $|\varphi|_F$. If $|\varphi|_F = 0$, then the claim trivially holds. Let $\varphi = \theta \wedge F\psi$ where θ and ψ are flat.

\Leftarrow) Since $\varphi \xrightarrow{w(0)\dots w(i-1)} \varphi'$ and $w[i..] \models \varphi'$, repeated applications of Item 1 yields $w \models \varphi$.

\Rightarrow) Let $w \models \varphi$. It is the case that $w \models \theta$ and $w \models F\psi$. Let $j \in \mathbb{N}$ be such that $w[j..] \models \psi$. By Item 1, there exists ψ' such that $\psi \xrightarrow{w(j)} \psi'$ and $w[j+1..] \models \psi'$. By definition of \rightarrow , there exists $\psi'' \in \mathfrak{U}(\psi)$ such that $\psi' = \psi''[w(j)]$. As $\mathfrak{U}(\psi) \subseteq \mathfrak{U}(F\psi)$, we have $\psi'' \in \mathfrak{U}(F\psi)$, and so $F\psi \xrightarrow{w(j)} \psi'$. Moreover, we have $F\psi \xrightarrow{w(0)\dots w(j-1)} F\psi$. Thus, $F\psi \xrightarrow{w(0)\dots w(j)} \psi'$.

By repeated applications of Item 1, there exists θ' such that $\theta \xrightarrow{w(0)\dots w(j)} \theta'$ and $w[j+1..] \models \theta'$. Hence,

$$(\theta \wedge F\psi) \xrightarrow{w(0)\dots w(j)} (\theta' \wedge \psi') \text{ and } w[j+1..] \models \theta' \wedge \psi'.$$

By Proposition 4, we have $|\theta|_F \geq |\theta'|_F$ and $|F\psi|_F > |\psi'|_F$. Thus, $|\theta \wedge F\psi|_F > |\theta' \wedge \psi'|_F$. So, by induction hypothesis, there exist $k \geq j+1$ and φ' such that $(\theta' \wedge \psi') \xrightarrow{w(j+1)\dots w(k-1)} \varphi'$, $|\varphi'|_F = 0$ and $w[k..] \models \varphi'$. We are done since

$$\varphi = (\theta \wedge F\psi) \xrightarrow{w(0)\dots w(j)} (\theta' \wedge \psi') \xrightarrow{w(j+1)\dots w(k-1)} \varphi'. \quad \square$$

Proposition 8: Let $q, r \in Q$. It is the case that

- 1) $X_{q,r}$ is either empty or simple,
- 2) if $X_{q,r} \neq \emptyset$, then $X_{r,r} \neq \emptyset$,
- 3) if $X_{q,q} \neq \emptyset$, then $X_{q,q} \supseteq X_{q,r}$.

Moreover, given $\theta \in \mathfrak{U}(q)$ and $A \subseteq AP$ such that $r = \theta[A]$, the representation of $X_{q,r}$ can be obtained in polynomial time.

Proof:

- 1) We assume that $X_{q,r} \neq \emptyset$, as we are otherwise trivially done. Let us first show that $X_{q,r}$ is closed under intersection. Let $A, B \in X_{q,r}$. By definition of \rightarrow , there must exist $\theta \in \mathfrak{U}(q)$ such that $r = \theta[A] = \theta[B]$. Moreover, θ is of the form

$$\theta = \psi \wedge G\psi' \wedge \bigwedge_{i \in I} GF\psi''_i \wedge \bigwedge_{j \in J} F\varphi_j,$$

and it is the case that $\text{prop}(\psi \wedge \psi') \subseteq A$ and $\text{prop}(\psi \wedge \psi') \subseteq B$. This means that $\text{prop}(\psi \wedge \psi') \subseteq A \cap B$, and hence that $r = \theta[A \cap B]$, which implies $q \rightarrow^{A \cap B} r$.

Since $X_{q,r}$ is closed under intersection, it has a minimal element A , which is in fact $A := \text{prop}(\psi \wedge \psi')$ and can thus be obtained in polynomial time from θ . We have $X_{q,r} = \uparrow A$, as for any $A' \supseteq A$, we have $\text{prop}(\psi \wedge \psi') = A \subseteq A'$, and so $q \rightarrow^{A'} \theta[A'] = r$.

- 2) Since $X_{q,r} \neq \emptyset$, there exist $\theta \in \mathcal{U}(q)$ and $A \subseteq AP$ such that $r = \theta[A]$. Thus, r is of the form $\mathbf{G}\psi' \wedge \bigwedge_{i \in I} \mathbf{GF}\psi_i'' \wedge \bigwedge_{j \in J} \mathbf{F}\varphi_j$. Note that $r \in \mathcal{U}(r)$ and $r[B] = r$ where $B := \text{prop}(\psi')$. Thus, $r \rightarrow^B r$, and so $B \in X_{r,r}$.
- 3) Since $X_{q,q} \neq \emptyset$, q is of the form $q = \mathbf{G}\psi' \wedge \bigwedge_{i \in I} \mathbf{GF}\psi_i'' \wedge \bigwedge_{j \in J} \mathbf{F}\varphi_j$. Let $\theta \in \mathcal{U}(q)$. The latter is of the form

$$\theta = \psi_{\mathcal{U}} \wedge \mathbf{G}(\psi' \wedge \psi'_{\mathcal{U}}) \wedge \bigwedge_{i \in I_{\mathcal{U}}} \mathbf{GF}\psi_i'' \wedge \bigwedge_{j \in J_{\mathcal{U}}} \mathbf{F}\varphi_j.$$

Testing $\text{prop}(\psi') \subseteq A$ is less restrictive than $\text{prop}(\psi_{\mathcal{U}} \wedge \psi' \wedge \psi'_{\mathcal{U}}) \subseteq A$. So, if $\theta[A] \neq \text{false}$, then $q[A] \neq \text{false}$. As $q \in \mathcal{U}(q)$, this means that $q \rightarrow^A r$ implies $q \rightarrow^A q$. \square

Theorem 2: The model-checking problem is in P for these fragments: $\text{LTL}_B(\{\mathbf{F}, \mathbf{G}, \neg\})$, $\text{LTL}(\{\mathbf{F}, \vee\})$ and $\text{LTL}(\{\mathbf{G}, \wedge\})$.

Proof: Let us give the details missing from the main text.

- 1) We show that $x \models_M \mathbf{F}\neg Z$ iff $x \notin Z$ or there exists a mode $m \in M$ such that $m \neq \mathbf{0}$.

\Rightarrow) For the sake of contradiction, suppose that $x \in Z$ and that there is no $m \in M$ with $m \neq \mathbf{0}$. Clearly, we can never reach a point other than x .

\Leftarrow) If $x \notin Z$, then clearly $x \models_M \mathbf{F}\neg Z$. Otherwise, we have $m \in Z$ with $m \neq \mathbf{0}$. Since we assume Z to be a bounded zone, it must hold that by simply scheduling m forever, we eventually leave Z .

- 2) We show that $x \models_M \mathbf{G}\neg Z$ iff $x \notin Z$ and there exists a mode $m \in M$ such that for all $\alpha \in \mathbb{R}_{>0}$ it is the case that $x + \alpha m \notin Z$.

\Leftarrow) It is easy to see that the execution obtained by scheduling m for an infinite duration satisfies the formula since $x + \alpha m \notin Z$ for all $\alpha \in \mathbb{R}_{>0}$.

\Rightarrow) Clearly, it must hold that $x \notin Z$. Let the modes of M be $\{m_1, \dots, m_n\}$. To prove the other half, assume for contradiction that for all $i \in [1..n]$, it holds that $x + \alpha_i m_i \in Z$ for some $\alpha_i \in \mathbb{R}_{>0}$. Let us demonstrate that for any non-Zeno infinite schedule ρ , it is the case that $\text{exec}(\rho, x) \not\models \mathbf{G}\neg Z$. We do so by showing the existence of $\tau \in \mathbb{R}_{\geq 0}$ and $z \in Z$ such that $x \rightarrow^{\rho[1..\tau]} z$.

Let $m'_i := \alpha_i m_i$, and let us consider the MMS $M' := \{\alpha_i m_i : i \in [1..n]\}$. Clearly, schedules of M and M' can be related: A schedule π of M amounts to the schedule π' of M' where we replace we replace each occurrence (β, m_i) with $(\beta/\alpha_i, m'_i)$. Thus, for every y, y' , we have $y \rightarrow^\pi y'$ iff $y \rightarrow^{\pi'} y'$.

Let us consider a schedule π of M' , and let $\rho = \pi[1..]$. It is the case that $x \rightarrow^\rho z$ with

$$z = x + \sum_{i=1}^n \lambda_i m'_i,$$

for some $\lambda_1, \dots, \lambda_n$ such that $\sum_{i=1}^n \lambda_i = 1$. Thus, z is a convex combination of the points $x + m'_i$ for any $i \in [1..n]$. But note that $x + m'_i = x + \alpha_i m_i$ by definition of M' , and $x + \alpha_i m_i \in Z$ by definition of α_i . So, z is a convex combination of points from Z , and is thus included in Z , as it is convex.

- 3) The proof for $\mathbf{GF}\neg Z$ is as in Item 1, except that it remains to note that once an execution has left Z by scheduling m , if it keeps scheduling m , then it will not re-enter Z .
- 4) The proof for $\mathbf{FG}\neg Z$ is as in Item 3. \square

Theorem 3: The model-checking problem is P-hard for both $\text{LTL}_B(\{\mathbf{F}\})$ and $\text{LTL}_B(\{\mathbf{G}\})$.

Proof: Fragment $\text{LTL}_B(\{\mathbf{F}\})$. We reduce from linear programming feasibility. This problem asks whether a given zone $Z \subseteq \mathbb{R}^d$, described by a system of inequalities $\mathbf{A}x \leq \mathbf{b}$, has a non-negative solution. The problem is P-complete even if zone Z is bounded [23, Prob. A.4.1]. Let $M := \{e_i : i \in [1..d]\}$. It is readily seen that $Z \neq \emptyset$ iff $\mathbf{0} \models_M \mathbf{F}Z$.

Fragment $\text{LTL}_B(\{\mathbf{G}\})$. We reduce from the monotone circuit-value problem (CVP) [23, Prob. A.1.3]. Let $C(x_1, \dots, x_n)$ be a boolean circuit with gates from $\{\wedge, \vee\}$, and let $w_1, \dots, w_n \in \{0, 1\}$. We construct a d -dimensional MMS M , $x \in \mathbb{R}^d$ and a bounded zone $Z \subseteq \mathbb{R}^d$ such that $C(w) = 1$ iff $x \models_M \mathbf{G}Z$.

Each gate g is associated to a dimension g . We add two dimensions $\{\heartsuit, \heartsuit\}$. Let $x := e_\heartsuit + \sum_{i=1}^n w_i \cdot e_{x_i}$. Let Z be the zone defined by $c \in [0, 1]$ for every dimension c .

We denote as g_{out} the output gate of C . We associate a mode m_g , to each gate $g = u \wedge v$, defined by $m_g := -e_u - e_v + e_g$. We associate the modes m_g and m'_g to each gate $g = u \vee v$, defined by $m_g := -e_u + e_g$ and $m'_g := -e_v + e_g$. We add two modes: $m_\heartsuit := -e_{g_{\text{out}}} - e_\heartsuit + e_{\heartsuit}$ and $m_\heartsuit := e_{g_{\text{out}}} + e_\heartsuit - e_{\heartsuit}$.

Let g_1, \dots, g_m be a topological ordering of the gates of C . Let C_j be the subcircuit obtained by setting g_j as the output gate. A routine induction shows that $C_j(w) = 1$ iff there exists y such that $x \rightarrow_Z^* y$ and $y(g_j) > 0$. The claim implies that $C(w) = 1$ iff $x \models_M \mathbf{G}Z$. Indeed, any non-Zeno infinite schedule must use m_\heartsuit . Moreover, upon reaching some y such that $y(g_{\text{out}}) = \alpha > 0$, it is possible to alternate between αm_\heartsuit and αm_\heartsuit indefinitely. \square

Proposition 10 follows inductively from this lemma:

Lemma 8: Let $\rho(\alpha, m)\rho'(\beta, m)\rho''$ be a schedule. This holds:

- If $x \rightarrow_Z^{\rho(\alpha, m)\rho'(\beta, m)\rho''} y$, then $x \rightarrow_Z^{\rho(\alpha, m)\frac{\alpha}{\alpha+\beta}[\rho'\rho'']}$,
- If $x \rightarrow_Z^{\rho''(\beta, m)\rho'(\alpha, m)\rho} y$, then $\rightarrow_Z^{\frac{\alpha}{\alpha+\beta}[\rho'\rho''](\alpha, m)\rho} y$.

Proof: We only prove the first item. The second follows symmetrically. We have $x \rightarrow_Z^\rho u$ for some u . By convexity, there exist v, v', w such that:

$$u \rightarrow_Z^{\frac{\alpha}{\alpha+\beta}[(\alpha, m)\rho'(\beta, m)]} v \rightarrow_Z^{\frac{\alpha}{\alpha+\beta}\rho''} w$$

and

$$\mathbf{u} \rightarrow_Z^{(\alpha, \mathbf{m}) \frac{\alpha}{\alpha+\beta} \rho'} \mathbf{v}'.$$

Let us show that $\mathbf{v} = \mathbf{v}'$:

$$\begin{aligned} \mathbf{v} &= \mathbf{u} + \frac{\alpha}{\alpha + \beta} (\alpha \mathbf{m} + \Delta_{\rho'} + \beta \mathbf{m}) \\ &= \mathbf{u} + \frac{\alpha}{\alpha + \beta} ((\alpha + \beta) \mathbf{m} + \Delta_{\rho'}) \\ &= \mathbf{u} + \alpha \mathbf{m} + \frac{\alpha}{\alpha + \beta} \Delta_{\rho'} \\ &= \mathbf{v}'. \end{aligned}$$

From this, we get $\mathbf{x} \rightarrow_Z^\rho \mathbf{u} \rightarrow_Z^{(\alpha, \mathbf{m}) \frac{\alpha}{\alpha+\beta} \rho'} \mathbf{v} \rightarrow_Z^{\frac{\alpha}{\alpha+\beta} \rho'} \mathbf{w}$. \square

Lemma 2: Let $\rho(\alpha, \mathbf{m})\rho'$ be a schedule. This holds:

- If $\mathbf{x} \rightarrow_Z^{\rho(\alpha, \mathbf{m})\rho'}$, then $\mathbf{x} \rightarrow_Z^{\rho(\frac{\alpha}{2}, \mathbf{m})\frac{1}{2}\rho'(\frac{\alpha}{2}, \mathbf{m})}$,
- If $\mathbf{y} \rightarrow_Z^{\rho'(\alpha, \mathbf{m})\rho}$, then $\mathbf{y} \rightarrow_Z^{(\frac{\alpha}{2}, \mathbf{m})\frac{1}{2}\rho'(\frac{\alpha}{2}, \mathbf{m})\rho}$.

Proof: We only prove the first item. The second follows symmetrically. By convexity, there exist \mathbf{y}' and \mathbf{w} such that

$$\mathbf{x} \rightarrow_Z^{\rho(\frac{\alpha}{2}, \mathbf{m})\frac{1}{2}\rho'} \mathbf{y}' \text{ and } \mathbf{x} \rightarrow_Z^{\rho(\alpha, \mathbf{m})\frac{1}{2}\rho'} \mathbf{w}.$$

Moreover, we have $\mathbf{y}' \rightarrow_Z^{(\frac{\alpha}{2}, \mathbf{m})} \mathbf{w}$. Since $\mathbf{y}', \mathbf{w} \in Z$, by convexity, we conclude that

$$\mathbf{x} \rightarrow_Z^{\rho(\frac{\alpha}{2}, \mathbf{m})\frac{1}{2}\rho'} \mathbf{y}' \rightarrow_Z^{(\frac{\alpha}{2}, \mathbf{m})} \mathbf{w}. \quad \square$$

Proposition 11: Let $\mathbf{x} \rightarrow_Z^\pi \mathbf{y}$. There exist $\beta \in \mathbb{N}_{\geq 1}$, $\mathbf{x} \rightarrow_Z^{\pi'} \mathbf{y}_Z$ and $\mathbf{x}_Z \rightarrow_Z^{\pi''} \mathbf{y}$ such that $|\pi| = |\pi'| = |\pi''|$, $\text{supp}(\pi) = \text{supp}(\pi') = \text{supp}(\pi'')$, and, for every $\mathbf{m} \in \text{supp}(\pi)$, it is the case that $\mathbf{x}_Z \rightarrow_Z^{(1/\beta)\mathbf{m}}$ and $\mathbf{y}_Z \rightarrow_Z^{(1/\beta)\mathbf{m}}$.

Proof: By Lemma 2, we can pick

$$\pi' := \prod_{i=1}^{|\pi|} \left(\frac{1}{2^i} \cdot \pi(i) \right), \quad \pi'' := \prod_{i=1}^{|\pi|} \left(\frac{1}{2^{(|\pi|-i+1)}} \cdot \pi(i) \right),$$

and $\beta := 2^{|\pi|} \cdot \lceil \text{time}(\pi) \rceil$. \square

Proposition 12: Let $\mathbf{x} \rightarrow_Z^\pi \mathbf{y}$, $k := |\pi|$ and $\beta \in \mathbb{N}_{\geq 1}$ be such that $\mathbf{x} \rightarrow_Z^{(1/\beta)\pi(i)}$ and $\mathbf{y} \rightarrow_Z^{(1/\beta)\pi(i)}$ hold for all $i \in [1..k]$. It is the case that $\mathbf{x} \rightarrow_Z^{\pi'} \mathbf{y}$, where $\pi' := ((1/(\beta k))\pi)^{\beta k}$.

Proof: Let $\pi := (\alpha_1, \mathbf{m}_1)(\alpha_2, \mathbf{m}_2) \cdots (\alpha_k, \mathbf{m}_k)$. Let $\mathbf{x}_0^0 := \mathbf{x}$, $\mathbf{x}_{\beta k-1}^k := \mathbf{y}$ and $\mathbf{x}_j^k := \mathbf{x}_{j+1}^0$. Let

$$\begin{aligned} \mathbf{x}_j^i &:= \mathbf{x}_0^0 + j \sum_{h=1}^k \frac{\alpha_h \mathbf{m}_h}{\beta k} + \sum_{h=1}^i \frac{\alpha_h \mathbf{m}_h}{\beta k} \\ &= \mathbf{x}_0^i + \sum_{h=i+1}^k \frac{\alpha_h \mathbf{m}_h}{\beta k} + (j-1) \sum_{h=1}^k \frac{\alpha_h \mathbf{m}_h}{\beta k} + \sum_{h=1}^i \frac{\alpha_h \mathbf{m}_h}{\beta k} \\ &= \mathbf{x}_0^i + \sum_{h=1}^k \frac{\alpha_h \mathbf{m}_h}{\beta k} + (j-1) \sum_{h=1}^k \frac{\alpha_h \mathbf{m}_h}{\beta k} \\ &= \mathbf{x}_0^i + j \sum_{h=1}^k \frac{\alpha_h \mathbf{m}_h}{\beta k}. \end{aligned}$$

Let $\mathbf{A} \ell \leq \mathbf{b}$ be the system of inequalities that represents zone Z . By assumption, the following holds for all $i \in [1..k]$:

$$\mathbf{A} \left(\mathbf{x}_0^0 + \frac{\alpha_i \mathbf{m}_i}{\beta} \right) \leq \mathbf{b} \text{ and } \mathbf{A} \left(\mathbf{x}_{\beta k-1}^k - \frac{\alpha_i \mathbf{m}_i}{\beta} \right) \leq \mathbf{b}.$$

The following holds:

$$\mathbf{x}_0^0 \rightarrow_Z^{\frac{1}{\beta k} \pi(1)} \mathbf{x}_0^1 \rightarrow_Z^{\frac{1}{\beta k} \pi(2)} \mathbf{x}_0^2 \rightarrow_Z^{\frac{1}{\beta k} \pi(3)} \cdots \rightarrow_Z^{\frac{1}{\beta k} \pi(k)} \mathbf{x}_0^k,$$

since

$$\begin{aligned} \mathbf{A} \mathbf{x}_0^i &= \mathbf{A} \mathbf{x}_0^0 + \sum_{j=1}^i \mathbf{A} \frac{\alpha_j \mathbf{m}_j}{\beta k} \\ &= \frac{k-i}{k} \mathbf{A} \mathbf{x}_0^0 + \frac{i}{k} \mathbf{A} \mathbf{x}_0^0 + \frac{1}{k} \sum_{j=1}^i \mathbf{A} \frac{\alpha_j \mathbf{m}_j}{\beta} \\ &= \frac{k-i}{k} \mathbf{A} \mathbf{x}_0^0 + \frac{1}{k} \sum_{j=1}^i \mathbf{A} \left(\mathbf{x}_0^0 + \frac{\alpha_j \mathbf{m}_j}{\beta} \right) \\ &\leq \frac{k-i}{k} \mathbf{b} + \frac{1}{k} \mathbf{b} \\ &= \mathbf{b}. \end{aligned}$$

Similarly, the following holds:

$$\mathbf{x}_{\beta k-1}^0 \rightarrow_Z^{\frac{1}{\beta k} \pi(1)} \mathbf{x}_{\beta k-1}^1 \rightarrow_Z^{\frac{1}{\beta k} \pi(2)} \mathbf{x}_{\beta k-1}^2 \cdots \rightarrow_Z^{\frac{1}{\beta k} \pi(k)} \mathbf{x}_{\beta k-1}^k.$$

It remains to show that, for all $i \in [1..k]$, $j \in [1..(\beta k - 1)]$, we have $\mathbf{x}_j^i = (1 - \lambda) \mathbf{x}_0^i + \lambda \mathbf{x}_{\beta k-1}^i$ where $\lambda = \frac{j}{\beta k - 1}$.

We have:

$$\begin{aligned} \mathbf{x}_j^i &= \mathbf{x}_0^i + j \sum_{h=1}^k \frac{\alpha_h \mathbf{m}_h}{\beta k} \\ &= \mathbf{x}_0^i + \frac{\beta k - 1}{\beta k - 1} \cdot j \sum_{h=1}^k \frac{\alpha_h \mathbf{m}_h}{\beta k} \\ &= \mathbf{x}_0^i + \lambda (\beta k - 1) \sum_{h=1}^k \frac{\alpha_h \mathbf{m}_h}{\beta k} \\ &= (1 - \lambda) \mathbf{x}_0^i + \lambda \left(\mathbf{x}_0^i + (\beta k - 1) \sum_{h=1}^k \frac{\alpha_h \mathbf{m}_h}{\beta k} \right) \\ &= (1 - \lambda) \mathbf{x}_0^i + \lambda \mathbf{x}_{\beta k-1}^i. \end{aligned}$$

Since $\lambda \in [0, 1]$, each \mathbf{x}_j^i belongs to Z by convexity. \square

Proposition 15: Let Z be a zone, let π be a schedule, let $\mathbf{x}, \mathbf{x}', \mathbf{y} \in Z$ and let $\beta \in (0, 1]$. Let $\mathbf{z} := \beta \mathbf{x} + (1 - \beta) \mathbf{y}$ and $\mathbf{z}' := \beta \mathbf{x}' + (1 - \beta) \mathbf{y}$. If $\mathbf{x} \rightarrow_Z^\pi \mathbf{x}'$ holds, then $\mathbf{z} \rightarrow_Z^{\beta \pi} \mathbf{z}'$.

Proof: Let us prove the case where $\pi = \alpha \mathbf{m}$ for some $\alpha \in \mathbb{R}_{>0}$ and mode $\mathbf{m} \in M$. The general case follows by induction. Note that $\mathbf{x}' = \mathbf{x} + \alpha \mathbf{m}$ and $\mathbf{x}' \in Z$. Let $\mathbf{z}' := \mathbf{z} + \beta \alpha \mathbf{m}$. We clearly have $\mathbf{z} \rightarrow_Z^{\beta \pi} \mathbf{z}'$. It remains to show that $\mathbf{z} \rightarrow_Z^{\beta \pi} \mathbf{z}'$. By convexity of Z , it suffices to show that \mathbf{z}' is on the line passing through $\mathbf{x}' \in Z$ and $\mathbf{y} \in Z$:

$$\begin{aligned} \mathbf{z}' &= (\beta \mathbf{x} + (1 - \beta) \mathbf{y}) + \beta \alpha \mathbf{m} \\ &= \beta (\mathbf{x} + \alpha \mathbf{m}) + (1 - \beta) \mathbf{y} \\ &= \beta \mathbf{x}' + (1 - \beta) \mathbf{y}. \end{aligned} \quad \square$$

Within the proof of Proposition 16, we have claimed that $f(\mathbf{x}, \mathbf{y}) \geq \lambda$. Let us prove this.

Proof of Proposition 16 (missing claim): We first make two additional claims. Given $\mathbf{x} \in X'$ and $\mathbf{y} \in Y'$ such that $\mathbf{z} \rightarrow_Z^\pi \mathbf{x} \rightarrow_Z^{\pi'} \mathbf{y}$, it is the case that:

- 1) $f'(\mathbf{x}, \mathbf{y}) \geq \lambda$, and
- 2) $f'(\mathbf{x}, \mathbf{y}) \leq \mathbf{v}_3^T(\mathbf{y} - \mathbf{x})$.

To prove Claim 1, note that we have

$$\begin{aligned}
f'(\mathbf{x}, \mathbf{y}) &= -(\mathbf{v}_1^T \mathbf{A}_1(\mathbf{x} - \mathbf{z}) + \mathbf{v}_2^T \mathbf{A}_2(\mathbf{y} - \mathbf{z})) + \\
&\quad \mathbf{v}_2^T \mathbf{A}_2(\mathbf{y} - \mathbf{x}) + \mathbf{v}_3^T(\mathbf{y} - \mathbf{x}) \\
&= -(\mathbf{v}_1^T \mathbf{A}_1(\mathbf{x} - \mathbf{z}) + \mathbf{v}_2^T \mathbf{A}_2(\mathbf{y} - \mathbf{z})) + \\
&\quad (\mathbf{v}_2^T \mathbf{A}_2 \mathbf{M} + \mathbf{v}_3^T \mathbf{M}) \boldsymbol{\pi}' \\
&\geq -(\mathbf{v}_1^T \mathbf{A}_1(\mathbf{x} - \mathbf{z}) + \mathbf{v}_2^T \mathbf{A}_2(\mathbf{y} - \mathbf{z})) \quad (\text{by (2)}) \\
&\geq -(\mathbf{v}_1^T(\mathbf{b}_1 - \mathbf{A}_1 \mathbf{z}) + \mathbf{v}_2^T(\mathbf{b}_2 - \mathbf{A}_2 \mathbf{z})) \\
&\quad (\text{by } \mathbf{x} \in X', \mathbf{y} \in Y') \\
&= \lambda.
\end{aligned}$$

Next, to prove Claim 2, we have

$$\begin{aligned}
f'(\mathbf{x}, \mathbf{y}) &= -(\mathbf{v}_1^T \mathbf{A}_1(\mathbf{x} - \mathbf{z}) + \mathbf{v}_2^T \mathbf{A}_2(\mathbf{y} - \mathbf{z})) + \\
&\quad \mathbf{v}_2^T \mathbf{A}_2(\mathbf{y} - \mathbf{x}) + \mathbf{v}_3^T(\mathbf{y} - \mathbf{x}) \\
&= -(\mathbf{v}_1^T \mathbf{A}_1(\mathbf{x} - \mathbf{z}) + \mathbf{v}_2^T \mathbf{A}_2(\mathbf{x} - \mathbf{z})) + \mathbf{v}_3^T(\mathbf{y} - \mathbf{x}) \\
&= -(\mathbf{v}_1^T \mathbf{A}_1 \mathbf{M} + \mathbf{v}_2^T \mathbf{A}_2 \mathbf{M}) \boldsymbol{\pi} + \mathbf{v}_3^T(\mathbf{y} - \mathbf{x}) \\
&\leq \mathbf{v}_3^T(\mathbf{y} - \mathbf{x}) \quad (\text{by (1)}).
\end{aligned}$$

Finally, by using Claims 2 and 1, we can prove the original claim: $f(\mathbf{x}, \mathbf{y}) = \mathbf{v}_3^T(\mathbf{y} - \mathbf{x}) \geq f'(\mathbf{x}, \mathbf{y}) \geq \lambda$. \square

Proposition 18: For every $n \in \mathbb{N}$, it is the case that $\mathbf{x}_n = \lambda^n \mathbf{x}_0 + (1 - \lambda^n) \mathbf{x}_f$ and $\mathbf{y}_n = \lambda^n \mathbf{y}_0 + (1 - \lambda^n) \mathbf{y}_f$.

Proof:

By unrolling the definition of \mathbf{x}_n and \mathbf{y}_n , we have

$$\begin{aligned}
&\mathbf{x}_n \quad (17) \\
&= \mathbf{x}_0 + \sum_{i=0}^{n-1} \left(\lambda^i \left(\boldsymbol{\Delta}_{\pi \pi'} + \frac{\boldsymbol{\Delta}_\rho - \epsilon \boldsymbol{\Delta}_{\pi \pi'}}{1 + \epsilon} \right) + (1 - \lambda^i) \boldsymbol{\Delta}_{\rho' \rho''} \right) \\
&= \mathbf{x}_0 + \sum_{i=0}^{n-1} \left(\lambda^i \left(\boldsymbol{\Delta}_{\pi \pi'} + \frac{\boldsymbol{\Delta}_\rho - \epsilon \boldsymbol{\Delta}_{\pi \pi'}}{1 + \epsilon} \right) + (1 - \lambda^i) \cdot \mathbf{0} \right) \\
&= \mathbf{x}_0 + \sum_{i=0}^{n-1} \lambda^i \left(\boldsymbol{\Delta}_{\pi \pi'} + \frac{\boldsymbol{\Delta}_\rho - \epsilon \boldsymbol{\Delta}_{\pi \pi'}}{1 + \epsilon} \right) \\
&= \mathbf{x}_0 + \frac{\lambda^n - 1}{\lambda - 1} \left(\boldsymbol{\Delta}_{\pi \pi'} + \frac{\boldsymbol{\Delta}_\rho - \epsilon \boldsymbol{\Delta}_{\pi \pi'}}{1 + \epsilon} \right) \quad (18) \\
&= \mathbf{x}_0 + \frac{\lambda^n - 1}{-1/(1 + \epsilon)} \left(\boldsymbol{\Delta}_{\pi \pi'} + \frac{\boldsymbol{\Delta}_\rho - \epsilon \boldsymbol{\Delta}_{\pi \pi'}}{1 + \epsilon} \right) \\
&= \mathbf{x}_0 - (\lambda^n - 1)(1 + \epsilon) \left(\boldsymbol{\Delta}_{\pi \pi'} + \frac{\boldsymbol{\Delta}_\rho - \epsilon \boldsymbol{\Delta}_{\pi \pi'}}{1 + \epsilon} \right) \\
&= \mathbf{x}_0 + (1 - \lambda^n)((1 + \epsilon) \boldsymbol{\Delta}_{\pi \pi'} + \boldsymbol{\Delta}_\rho - \epsilon \boldsymbol{\Delta}_{\pi \pi'}) \\
&= \mathbf{x}_0 + (1 - \lambda^n) \boldsymbol{\Delta}_{\pi \pi' \rho} \\
&= \lambda^n \mathbf{x}_0 + (1 - \lambda^n)(\mathbf{x}_0 + \boldsymbol{\Delta}_{\pi \pi' \rho}) \\
&= \lambda^n \mathbf{x}_0 + (1 - \lambda^n) \mathbf{x}_f,
\end{aligned}$$

where (18) follows from $\sum_{i=0}^{n-1} \lambda^i = (1 - \lambda^n)/(1 - \lambda)$. Furthermore, we have

$$\begin{aligned}
\mathbf{y}_n &= \mathbf{x}_n + \lambda^n \boldsymbol{\Delta}_\pi + (1 - \lambda^n) \boldsymbol{\Delta}_{\rho'} \\
&= \lambda^n \mathbf{x}_0 + (1 - \lambda^n) \mathbf{x}_f + \lambda^n \boldsymbol{\Delta}_\pi + (1 - \lambda^n) \boldsymbol{\Delta}_{\rho'} \quad (19) \\
&= \lambda^n (\mathbf{x}_0 + \boldsymbol{\Delta}_\pi) + (1 - \lambda^n) (\mathbf{x}_f + \boldsymbol{\Delta}_{\rho'}) \\
&= \lambda^n \mathbf{y}_0 + (1 - \lambda^n) \mathbf{y}_f,
\end{aligned}$$

where (19) follows from the fact that $\mathbf{x}_n = \lambda^n \mathbf{x}_0 + (1 - \lambda^n) \mathbf{x}_f$ for all $n \in \mathbb{N}$, as proven above. \square

Proposition 21: If $\mathbf{z} \models_M \mathbf{GZ}$, then there exist $\boldsymbol{\pi}$ and \mathbf{z}' such that $\mathbf{z} \rightarrow^\pi \mathbf{z}'$, $\mathbf{A} \mathbf{z}' \leq \mathbf{A} \mathbf{z}$ and $\|\boldsymbol{\pi}\| \geq 1$.

Proof: We define \mathbf{M} as the matrix such that each column is a mode from M . Observe that the following set of constraints \mathcal{S} is equivalent to $\exists \boldsymbol{\pi}, \mathbf{z}' : \mathbf{z} \rightarrow^\pi \mathbf{z}' \wedge \mathbf{A} \mathbf{z}' \leq \mathbf{A} \mathbf{z} \wedge \|\boldsymbol{\pi}\| \geq 1$:

$$\exists \mathbf{u} \geq \mathbf{0} : \begin{bmatrix} \mathbf{A} \mathbf{M} \\ -\mathbf{1}^T \end{bmatrix} \mathbf{u} \leq \begin{bmatrix} \mathbf{0} \\ -1 \end{bmatrix}.$$

For the sake of contradiction, suppose that $\mathbf{z} \models \mathbf{GZ}$ and that \mathcal{S} has no solution. By Farkas' lemma, the following system \mathcal{S}' has a solution:

$$\begin{aligned}
&\exists \mathbf{v}_1 \geq \mathbf{0}, \mathbf{v}_2 \geq \mathbf{0} : \\
&\quad \begin{bmatrix} \mathbf{M}^T \mathbf{A}^T & -\mathbf{1} \end{bmatrix} \begin{bmatrix} \mathbf{v}_1 \\ \mathbf{v}_2 \end{bmatrix} \geq \mathbf{0}, \quad [\mathbf{0}^T \ -1] \begin{bmatrix} \mathbf{v}_1 \\ \mathbf{v}_2 \end{bmatrix} < 0.
\end{aligned}$$

The latter can be rewritten equivalently as follows:

$$\begin{aligned}
&\exists \mathbf{v}_1 \geq \mathbf{0}, \mathbf{v}_2 \geq \mathbf{0} : \mathbf{M}^T \mathbf{A}^T \mathbf{v}_1 - \mathbf{1} \mathbf{v}_2 \geq \mathbf{0}, \quad -\mathbf{v}_2 < 0 \\
&\iff \exists \mathbf{v}_1 \geq \mathbf{0}, \mathbf{v}_2 \geq \mathbf{0} : \mathbf{v}_1^T \mathbf{A} \mathbf{M} \geq \mathbf{1}^T \mathbf{v}_2, \quad \mathbf{v}_2 > 0 \\
&\iff \exists \mathbf{v}_1 \geq \mathbf{0}, \mathbf{v}_2 \geq \mathbf{0} : \mathbf{v}_1^T \mathbf{A} \mathbf{M} \geq \mathbf{1}^T \mathbf{v}_2 > \mathbf{0}^T.
\end{aligned}$$

Since $\mathbf{z} \models \mathbf{GZ}$, there exists a non-Zeno infinite schedule π such that $\mathbf{z} \rightarrow_{\pi}^{\infty}$. Let $n \in \mathbb{N}$, $\pi_n := \pi[1..n]$ and $\mathbf{z}_n := \mathbf{z} + \boldsymbol{\Delta}_{\pi_n}$. We have $\mathbf{z}_n \in Z$ and hence $\mathbf{A} \mathbf{z}_n \leq \mathbf{b}$. Thus,

$$\begin{aligned}
\mathbf{v}_2 \mathbf{1}^T \boldsymbol{\pi}_n &\leq \mathbf{v}_1^T \mathbf{A} \mathbf{M} \boldsymbol{\pi}_n \\
&= \mathbf{v}_1^T \mathbf{A} (\mathbf{z}_n - \mathbf{z}) \\
&\leq \mathbf{v}_1^T (\mathbf{b} - \mathbf{A} \mathbf{z}). \quad (20)
\end{aligned}$$

Since π is non-Zeno, we further have $\lim_{n \rightarrow \infty} \mathbf{v}_2 \mathbf{1}^T \boldsymbol{\pi}_n = \infty$, which contradicts (20). \square

Proposition 22: Let $\mathbf{z}, \mathbf{z}' \in Z$ and ρ be a finite schedule. If $\mathbf{z} \rightarrow_Z^\rho$ and $\mathbf{A} \mathbf{z}' \leq \mathbf{A} \mathbf{z}$ then $\mathbf{z}' \rightarrow_Z^\rho$.

Proof: We consider the case where $\rho = \alpha m$. The general case follows inductively. Since $\mathbf{z} \rightarrow_Z^\rho$, we have $\mathbf{A} \mathbf{z} + \mathbf{A} \alpha m \leq \mathbf{b}$. Thus, $\mathbf{A} \mathbf{z}' + \mathbf{A} \alpha m \leq \mathbf{A} \mathbf{z} + \mathbf{A} \alpha m \leq \mathbf{b}$. \square

Proof of the construction in Lemma 3:

\Rightarrow Let π be a non-Zeno infinite schedule that satisfies $\text{exec}(\pi, \mathbf{x}) \models \varphi$. We construct a non-Zeno infinite schedule π' such that $\text{exec}(\pi', (\mathbf{x}, \dots, \mathbf{x})) \models \varphi'$. By definition of φ , for each $i \in [1..n]$, there exist $\tau_{i,0} < \tau_{i,1} < \dots \in \mathbb{R}_{\geq 0}$ and $\mathbf{x}_{i,0}, \mathbf{x}_{i,1}, \dots \in \mathbb{R}^d$ such that $\tau_{i,0} = 0$, $\mathbf{x}_{i,0} = \mathbf{x}$ and

$$\mathbf{x}_{i,j} \xrightarrow{\pi[\tau_{i,j}.. \tau_{i,j+1}]} \mathbf{x}_{i,j+1} \in Z_i \quad \text{for all } j \in \mathbb{N}.$$

Without loss of generality, we may assume that $\tau_{i,j} \leq \tau_{1,j} \leq \tau_{i,j+1}$ for all $i \in [1..n]$ and $j \in \mathbb{N}$. Indeed, the values can be chosen this way since each zone Z_i is visited infinitely often.

Let π_i denote the schedule obtained from π by replacing each mode $\mathbf{m} \in M$ with $\mathbf{m}_i \in M'$. Let $\pi' := u_0 v_0 u_1 v_1 \dots$ be the infinite schedule, where

$$\begin{aligned} u_j &:= \pi_1[\tau_{1,j}.. \tau_{1,j+1}] \quad \pi_2[\tau_{1,j}.. \tau_{2,j+1}] \quad \dots \quad \pi_n[\tau_{1,j}.. \tau_{n,j+1}], \\ v_j &:= \pi_1[\tau_{1,j+1}.. \tau_{1,j+1}] \pi_2[\tau_{2,j+1}.. \tau_{1,j+1}] \dots \pi_n[\tau_{n,j+1}.. \tau_{1,j+1}]. \end{aligned}$$

By definition, we have

$$\begin{aligned} (\mathbf{x}, \dots, \mathbf{x}) &\rightarrow_Z^{u_0} (\mathbf{x}_{1,1}, \dots, \mathbf{x}_{n,1}) \in X \\ &\rightarrow_Z^{v_0} (\mathbf{x}_{1,1}, \dots, \mathbf{x}_{1,1}) \in Y \\ &\rightarrow_Z^{u_1} (\mathbf{x}_{1,2}, \dots, \mathbf{x}_{n,2}) \in X \\ &\rightarrow_Z^{v_1} (\mathbf{x}_{1,2}, \dots, \mathbf{x}_{1,2}) \in Y \\ &\rightarrow_Z^{u_2} \dots \end{aligned}$$

Thus, it follows that $(\mathbf{x}, \dots, \mathbf{x}) \models_{M'} \varphi'$.

\Leftarrow Let π' be a non-Zeno infinite schedule that satisfies $\text{exec}(\pi', (\mathbf{x}, \dots, \mathbf{x})) \models \varphi'$. There exist $0 = \tau_0 < \gamma_1 < \tau_1 < \gamma_2 < \tau_2 < \dots \in \mathbb{R}_{\geq 0}$ and $\mathbf{x}_1, \mathbf{y}_1, \mathbf{x}_2, \mathbf{y}_2, \dots \in \mathbb{R}^{nd}$ such that

$$\begin{aligned} (\mathbf{x}, \dots, \mathbf{x}) &\rightarrow_Z^{\pi'[\tau_0.. \gamma_1]} \mathbf{x}_1 \in X \\ &\rightarrow_Z^{\pi'[\gamma_1.. \tau_1]} \mathbf{y}_1 \in Y \\ &\rightarrow_Z^{\pi'[\tau_1.. \gamma_2]} \mathbf{x}_2 \in X \\ &\rightarrow_Z^{\pi'[\gamma_2.. \tau_2]} \mathbf{y}_2 \in Y \\ &\rightarrow_Z^{\pi'[\tau_2.. \gamma_3]} \dots \end{aligned}$$

For every $i \in [1..n]$, let π'_i denote the schedule obtained from π' by keeping only the modes from $\{\mathbf{m}_i : \mathbf{m} \in M\}$, and changing each occurrence of \mathbf{m}_i to \mathbf{m} . Let

$$\begin{aligned} \pi &:= \pi_1[\tau_0.. \gamma_1] \quad \pi_1[\gamma_1.. \tau_1] \\ &\quad \pi_2[\tau_1.. \gamma_2] \quad \pi_2[\gamma_2.. \tau_2] \\ &\quad \dots \\ &\quad \pi_n[\tau_{n-1}.. \gamma_n] \quad \pi_n[\gamma_n.. \tau_n] \\ &\quad \pi_1[\tau_n.. \gamma_{n+1}] \quad \pi_1[\gamma_{n+1}.. \tau_{n+1}] \\ &\quad \dots \end{aligned}$$

Recall that $\mathbf{y}_j[i] = \mathbf{y}_j[i']$ for all $i, i' \in [1..n]$. So, by definition, we have

$$\begin{array}{lll} \mathbf{x} \rightarrow_{Z_0}^{\pi_1[\tau_0.. \gamma_1]} & \mathbf{x}_1[1] \rightarrow_{Z_0}^{\pi_1[\gamma_1.. \tau_1]} & \mathbf{y}_1[1] = \mathbf{y}_1[2] \\ \rightarrow_{Z_0}^{\pi_2[\tau_1.. \gamma_2]} & \mathbf{x}_2[2] \rightarrow_{Z_0}^{\pi_2[\gamma_2.. \tau_2]} & \mathbf{y}_2[2] = \mathbf{y}_2[3] \\ \dots & \dots & \\ \rightarrow_{Z_0}^{\pi_n[\tau_{n-1}.. \gamma_n]} & \mathbf{x}_n[n] \rightarrow_{Z_0}^{\pi_n[\gamma_n.. \tau_n]} & \mathbf{y}_n[n] = \mathbf{y}_n[1] \\ \rightarrow_{Z_0}^{\pi_1[\tau_n.. \gamma_{n+1}]} & \mathbf{x}_{n+1}[1] \rightarrow_{Z_0}^{\pi_1[\gamma_{n+1}.. \tau_{n+1}]} & \mathbf{y}_{n+1}[1] = \mathbf{y}_{n+1}[2] \\ \dots & \dots & \end{array}$$

Recall that $\mathbf{x}_j[i] \in Z_i$ for all $i \in [1..n]$ and $j \in \mathbb{N}$. Thus, each zone Z_i is visited infinitely often, and so $\text{exec}(\pi, \mathbf{x}) \models \varphi$. \square

Proof of Theorem 6 (correctness of reduction): Let us show that $\mathbf{0} \models_M \varphi$ iff there is a solution V to the SUBSET-SUM instance (S, t) . Recall that $\mathbf{0} \models_M \varphi$ holds iff M has a non-Zeno infinite schedule π such that $\text{exec}(\pi, \mathbf{0}) \models \varphi$.

\Leftarrow Let $V \subseteq S$ be such that $\sum_{v \in V} v = t$. We define a schedule π that satisfies φ . Let $\pi := \pi_1 \pi_2 \dots \pi_n \mathbf{y}_1^\omega$, where

$$\pi_i := \begin{cases} \mathbf{y}_i \bar{\mathbf{y}}_i & \text{if } s_i \in V, \\ \mathbf{n}_i \bar{\mathbf{n}}_i & \text{otherwise.} \end{cases}$$

\Rightarrow Let π be such that $\sigma := \text{exec}(\pi, \mathbf{0}) \models \varphi$. By definition of φ , there exist $\tau_T, \tau_{Y_1}, \tau_{N_1}, \dots, \tau_{Y_n}, \tau_{N_n} \in \mathbb{R}_{\geq 0}$ such that $\sigma(\tau_T) \in T$, $\sigma(\tau_{Y_i}) \in Y_i$ and $\sigma(\tau_{N_i}) \in N_i$ for all $i \in [1..n]$. Let all of these be minimal.

Let $\sigma' := \sigma[0.. \tau_T]$. Since π is a schedule for σ , there exists a schedule π' such that $\sigma' = \text{exec}(\pi', \mathbf{0})$. We will show that:

$$\text{time}_{\mathbf{y}_i}(\pi'), \text{time}_{\mathbf{n}_i}(\pi') \in \{0, 1\} \text{ for all } i \in [1..n]. \quad (*)$$

From (*), we can finish the proof. Indeed, by definition of π' and of the modes, we have

$$\sigma(\tau_T)(c^*) = \sum_{i=1}^n \text{time}_{\mathbf{y}_i}(\pi') \cdot s_i.$$

Additionally, $\sigma(\tau_T)(c^*) = t$ holds by definition of zone T . Since each $\text{time}_{\mathbf{y}_i}(\pi') \in \{0, 1\}$, we obtain a solution $V := \{s_i : \text{time}_{\mathbf{y}_i}(\pi') = 1\}$ to the SUBSET-SUM instance.

It remains to show (*). We first make the following claims for every $i \in [1..n]$:

- (1) $\tau_{Y_i} \leq \tau_T$ and $\tau_{N_i} \leq \tau_T$.
- (2) $\text{time}_{\mathbf{y}_i}(\pi') + \text{time}_{\mathbf{n}_i}(\pi') = 1$ and $\text{time}_{\bar{\mathbf{y}}_i}(\pi') + \text{time}_{\bar{\mathbf{n}}_i}(\pi') = 1$.

Let us prove these two claims.

- (1) We only show that $\tau_{Y_i} \leq \tau_T$, as $\tau_{N_i} \leq \tau_T$ is symmetric. We proceed by proving that for any $\mathbf{x} \in T \setminus Y_i$, it is the case that $\mathbf{x} \not\models_M \text{FY}_i$. By definition of T and Y_i , we have $\mathbf{x}(c_{i,1}) \in [-0.5, 0.5]$, $\mathbf{x}(c_{i,2}) = 2$ and $\mathbf{x}(c_{i,3}) = \mathbf{x}(c_{i,4}) = 1$. Note that the only modes affecting $\{c_{i,1}, \dots, c_{i,4}\}$ are $\{\mathbf{y}_i, \mathbf{n}_i, \bar{\mathbf{y}}_i, \bar{\mathbf{n}}_i\}$. The only modes affecting $c_{i,1}$ positively, and that could thus lead it to 0.5, are \mathbf{y}_i and $\bar{\mathbf{n}}_i$. However, both affect $c_{i,2}$ positively, and no mode decreases $c_{i,2}$. Hence, using either mode from \mathbf{x} can never lead to a point in Y_i . Thus, $\mathbf{x} \not\models_M \text{FY}_i$.

- (2) By $\sigma(\tau_T) \in T$, we have $\sigma(\tau_T)(c_{i,2}) = 2$. Since $\mathbf{y}_i(c_{i,2}) = \mathbf{n}_i(c_{i,2}) = \bar{\mathbf{y}}_i(c_{i,2}) = \bar{\mathbf{n}}_i(c_{i,2}) = 1$, and since no mode decreases $c_{i,2}$, it is the case that $\text{time}_{\mathbf{y}_i}(\pi') + \text{time}_{\mathbf{n}_i}(\pi') + \text{time}_{\bar{\mathbf{y}}_i}(\pi') + \text{time}_{\bar{\mathbf{n}}_i}(\pi') = 2$. Since T requires $c_{i,3} = 1$, and no mode decreases $c_{i,3}$, we have $\text{time}_{\mathbf{y}_i}(\pi') + \text{time}_{\mathbf{n}_i}(\pi') = 1$. Similarly, as T requires $c_{i,4} = 1$, we have $\text{time}_{\bar{\mathbf{y}}_i}(\pi') + \text{time}_{\bar{\mathbf{n}}_i}(\pi') = 1$.

It remains to use the above claims to prove (*). Note that $\tau_{Y_i} \neq \tau_{N_i}$ as the constraints of Y_i and N_i conflict on $c_{i,1}$. Let us assume that $\tau_{Y_i} < \tau_{N_i}$ (the other case is symmetric). Let ψ and ψ' be schedules such that $\text{exec}(\psi, \mathbf{0}) = \sigma[0.. \tau_{Y_i}]$ and $\text{exec}(\psi', \sigma(\tau_{Y_i})) = \sigma[\tau_{Y_i}.. \tau_{N_i}]$.

By definition of Y_i and N_i , we have $\sigma(\tau_{Y_i})(c_{i,3}) = 1$ and $\sigma(\tau_{N_i})(c_{i,3}) = 1$. Since no mode decreases $c_{i,3}$, modes \mathbf{y}_i and \mathbf{n}_i are not used in ψ' . Further, note that $\sigma(\tau_{Y_i})(c_{i,1}) = 0.5$ and $\sigma(\tau_{N_i})(c_{i,1}) = -0.5$. Therefore, it must be the case that

$$0.5 - \text{time}_{\bar{\mathbf{y}}_i}(\psi') + \text{time}_{\bar{\mathbf{n}}_i}(\psi') = -0.5.$$

Thus, $\text{time}_{\bar{y}_i}(\psi') - \text{time}_{\bar{n}_i}(\psi') = 1$. As ψ' arises from π' , Claim 2 yields

$$\text{time}_{\bar{y}_i}(\psi') + \text{time}_{\bar{n}_i}(\psi') \leq 1.$$

So, we have $\text{time}_{\bar{y}_i}(\psi') = 1$ and $\text{time}_{\bar{n}_i}(\psi') = 0$. From Claim 2, we further derive $\text{time}_{\bar{y}_i}(\psi) = \text{time}_{\bar{n}_i}(\psi) = 0$.

By definition of Y_i , we have $\sigma(\tau_{Y_i})(c_{i,1}) = 0.5$. Since \mathbf{y}_i and \mathbf{n}_i are the only modes possibly used in ψ to change $c_{i,1}$, we have $\frac{1}{2} \cdot \text{time}_{\mathbf{y}_i}(\psi) - \frac{1}{2} \cdot \text{time}_{\mathbf{n}_i}(\psi) = 0.5$. By Claim (2),

$$\text{time}_{\mathbf{y}_i}(\psi) + \text{time}_{\mathbf{n}_i}(\psi) \leq 1.$$

So, we have $\text{time}_{\mathbf{y}_i}(\psi) = 1$ and $\text{time}_{\mathbf{n}_i}(\psi) = 0$, which, by Claim (2), yields $\text{time}_{\mathbf{y}_i}(\pi') = 1$ and $\text{time}_{\mathbf{n}_i}(\pi') = 0$. \square

In Section VI, we implicitly assume that Petri nets with inhibitor arcs have no transition that consumes from, and produces in, the same place. We can make this assumption without loss of generality. Roughly, it is possible to split a transition t that consumes and produces in the same place into two transitions t_{pre} and t_{post} , while being equivalent with respect to reachability. We let t_{pre} consume from the place and realize the effect of t on all other places, and let t_{post} produce in the place. We can ensure that when t_{pre} , then immediately afterwards t_{post} is fired by adding a new place p_t , adding an arc from t_{pre} to p_t and an arc from p_t to t_{post} , and adding an inhibitor arc from p_t to all other transitions, thus preventing any transition other than t_{post} from being fired until t_{post} was fired to consume the token from p_t .

Now, let us prove the statements from Section VI.

Lemma 4: Let $\mathbf{x}_A, \mathbf{x}'_A \in A$ and let π be a finite schedule such that $\mathbf{x}_A(t_A) = 1$, $|\pi| > 0$, and $\mathbf{x}_A \xrightarrow{\pi}_{AP} \mathbf{x}'_A$ holds with no intermediate points in A , i.e. $\text{exec}(\pi, \mathbf{x}_A)(\tau) \in A$ iff $\tau \in \{0, \text{time}(\pi)\}$. It is the case that $\pi \equiv \mathbf{a}_t \mathbf{b}_t \mathbf{c}_t$ and there exist $\mathbf{x}_B \in B_t$, $\mathbf{x}_C \in C_t$ such that

$$\mathbf{x}_A \xrightarrow{\mathbf{a}_t}_{A'_t} \mathbf{x}_B \xrightarrow{\mathbf{b}_t}_{B'_t} \mathbf{x}_C \xrightarrow{\mathbf{c}_t}_{C'_t} \mathbf{x}'_A.$$

Proof: Let $\mathbf{x}_A, \mathbf{x}'_A \in A$ and π be as described.

Let $\pi = (\alpha, \mathbf{m})\pi'$. By definition of A , we must have $\mathbf{m} = \mathbf{a}_t$ for some $t \in T$. Since $\mathbf{x}_A(t_A) = 1$, we have $\alpha \in (0, 1]$. If $\alpha < 1$, then zone $A'_t \setminus A$ is reached, and the only mode that can be used is \mathbf{a}_t . Thus, we can assume w.l.o.g. that $\alpha = 1$. Let $\mathbf{x}_B := \mathbf{x}_A + \mathbf{a}_t$. We have

$$\mathbf{x}_A \xrightarrow{\mathbf{a}_t}_{A'_t} \mathbf{x}_B \text{ and } \mathbf{x}_B \in B_t.$$

Let $\pi' = (\alpha', \mathbf{m}')\pi''$. By definition of B_t , we have $\mathbf{m}' = \mathbf{b}_t$ or $\mathbf{m}' = \mathbf{a}_s$ for some $s \neq t$. Since no zone allows for $t_B > 0$ and $s_B > 0$, we must have $\mathbf{m}' = \mathbf{b}_t$. Since $\mathbf{x}_B(t_B) = 1$, we have $\alpha' \in (0, 1]$. If $\alpha' < 1$, then zone $B'_t \setminus B_t$ is reached, and, again, the only mode that can be used is \mathbf{b}_t . So, we can assume w.l.o.g. that $\alpha' = 1$. Let $\mathbf{x}_C := \mathbf{x}_B + \mathbf{b}_t$. We have

$$\mathbf{x}_B \xrightarrow{\mathbf{b}_t}_{B'_t} \mathbf{x}_C \text{ and } \mathbf{x}_C \in C_t.$$

Let $\pi'' = (\alpha'', \mathbf{m}'')\pi'''$. By definition of C_t , we have $\mathbf{m}'' = \mathbf{c}_t$ or $\mathbf{m}'' = \mathbf{a}_s$ for some $s \neq t$. Since no zone allows for $t_C > 0$ and $s_B > 0$, we must have $\mathbf{m}'' = \mathbf{c}_t$. Since $\mathbf{x}_C(t_C) = 1$, we have $\alpha'' \in (0, 1]$. If $\alpha'' < 1$, then zone $C'_t \setminus C_t$ is reached,

and, again, the only mode that can be used is \mathbf{c}_t . So, we can assume w.l.o.g. that $\alpha'' = 1$. Let $\mathbf{y} := \mathbf{x}_C + \mathbf{c}_t$. We have

$$\mathbf{x}_C \xrightarrow{\mathbf{c}_t}_{C'_t} \mathbf{y} \text{ and } \mathbf{y} \in A.$$

Since execution $\text{exec}(\pi, \mathbf{x}_A)$ does not contain intermediate points in A , we have $\mathbf{y} = \mathbf{x}'_A$ and hence π''' is empty. Consequently, $\pi \equiv \mathbf{a}_t \mathbf{b}_t \mathbf{c}_t$. \square

Proposition 25: Given a Petri net with inhibitor arcs \mathcal{N} and $\mathbf{x}_{\text{src}}, \mathbf{x}_{\text{tgt}}$, it is possible to compute an MMS M , points \mathbf{x}, \mathbf{x}' , and a finite set of zones AP closed under scaling, such that

- 1) $\mathbf{x}_{\text{src}} \xrightarrow{+} \mathbf{x}_{\text{tgt}}$ in \mathcal{N} iff $\mathbf{x} \xrightarrow{\pi}_{AP} \mathbf{x}'$ in M for some finite schedule π with $\text{time}(\pi) \geq 1$,
- 2) $\mathbf{0} \not\xrightarrow{+}_{AP} \mathbf{0}$ in M .

Proof: We show the proposition with M and AP described above. Let \mathbf{x} be the point such that $\mathbf{x}(t_A) := 1$ for all $t \in T$, $\mathbf{x}(\mathbf{p}) := \mathbf{x}_{\text{src}}$, and $\mathbf{x}(j) := 0$ for any other j . Let \mathbf{x}' be defined in the same way, but with \mathbf{x}_{tgt} rather than \mathbf{x}_{src} .

Item 2 follows from the fact that zones from AP are non-negative, and each mode of M decreases some dimension. It remains to show that $\mathbf{x}_{\text{src}} \xrightarrow{+} \mathbf{x}_{\text{tgt}}$ in \mathcal{N} iff $\mathbf{x} \xrightarrow{\pi}_{AP} \mathbf{x}'$ in M .

\Rightarrow) Let $\mathbf{x}_{\text{src}} \xrightarrow{+} \mathbf{x}_{\text{tgt}}$ in \mathcal{N} , where $\pi = t_1 \cdots t_k$. Let $\pi' := \mathbf{a}_{t_1} \mathbf{b}_{t_1} \mathbf{c}_{t_1} \cdots \mathbf{a}_{t_k} \mathbf{b}_{t_k} \mathbf{c}_{t_k}$. We have $\mathbf{x} \xrightarrow{\pi'}_{AP} \mathbf{x}'$ in M .

\Leftarrow) Let $\mathbf{x} \xrightarrow{\pi}_{AP} \mathbf{x}'$ in M where π is nonempty. As $\mathbf{x}, \mathbf{x}' \in A$, repeated applications of Lemma 4 yield $t_1, \dots, t_k \in T$ and points $\mathbf{y}_{A,1}, \mathbf{y}_{B,1}, \mathbf{y}_{C,1}, \dots, \mathbf{y}_{A,k}, \mathbf{y}_{B,k}, \mathbf{y}_{C,k}, \mathbf{y}_{A,k+1}$ such that $\pi \equiv \mathbf{a}_{t_1} \mathbf{b}_{t_1} \mathbf{c}_{t_1} \cdots \mathbf{a}_{t_k} \mathbf{b}_{t_k} \mathbf{c}_{t_k}$, and for all $i \in [1..k]$:

- 1) $\mathbf{y}_{A,i} \in A$, $\mathbf{y}_{B,i} \in B_{t_i}$ and $\mathbf{y}_{C,i} \in C_{t_i}$,
- 2) $\mathbf{y}_{A,1} = \mathbf{x}$ and $\mathbf{y}_{A,k+1} = \mathbf{x}'$, and
- 3) $\mathbf{y}_{A,i} \xrightarrow{\mathbf{a}_{t_i}}_{A'_{t_i}} \mathbf{y}_{B,i} \xrightarrow{\mathbf{b}_{t_i}}_{B'_{t_i}} \mathbf{y}_{C,i} \xrightarrow{\mathbf{c}_{t_i}}_{C'_{t_i}} \mathbf{y}_{A,i+1}$ in M .

By definition of the modes and zones, Item 3 yields

$$\mathbf{y}_{A,i}(\mathbf{p}) \xrightarrow{t_i} \mathbf{y}_{A,i+1}(\mathbf{p}) \text{ in } \mathcal{N} \quad \text{for all } i \in [1..k].$$

Thus, $\mathbf{x}_{\text{src}} = \mathbf{y}_{A,1}(\mathbf{p}) \xrightarrow{t_1 \cdots t_k} \mathbf{y}_{A,k+1}(\mathbf{p}) = \mathbf{x}_{\text{tgt}}$ in \mathcal{N} . \square

Lemma 5: Given $\psi_1, \dots, \psi_n, \varphi \in \text{LTL}_B(\{\mathbf{U}\})$, it is possible to compute a formula from $\text{LTL}_B(\{\mathbf{U}\})$ that is equivalent to formula $(\psi_1 \vee \dots \vee \psi_n) \mathbf{U} \varphi$.

Proof: The two following equivalences hold for LTL formulas interpreted over infinite words:

- 1) $(\psi \vee \psi') \mathbf{U} \varphi \equiv (\psi \mathbf{U} \psi') \mathbf{U} ((\psi' \mathbf{U} \psi) \mathbf{U} \varphi)$,
- 2) $\varphi \mathbf{U} (\psi \vee \psi') \equiv (\varphi \mathbf{U} \psi) \vee (\varphi \mathbf{U} \psi')$.

By Propositions 1 and 2, these equivalences also hold for negation-free LTL formulas interpreted over executions.

We proceed by induction on n . If $n = 1$, then the claim is trivial. Assume $n \geq 2$. Let $\psi' := \psi_2 \vee \dots \vee \psi_n$. We have:

$$\begin{aligned} & (\psi_1 \vee \psi_2 \vee \dots \vee \psi_n) \mathbf{U} \varphi \\ & \equiv (\psi_1 \vee \psi') \mathbf{U} \varphi \\ & \equiv (\psi_1 \mathbf{U} \psi') \mathbf{U} ((\psi' \mathbf{U} \psi_1) \mathbf{U} \varphi) \end{aligned} \tag{21}$$

$$\equiv (\psi_1 \mathbf{U} \psi') \mathbf{U} (\theta \mathbf{U} \varphi) \tag{22}$$

$$\equiv [(\psi_1 \mathbf{U} \psi_2) \vee \dots \vee (\psi_1 \mathbf{U} \psi_n)] \mathbf{U} (\theta \mathbf{U} \varphi) \tag{23}$$

$$\equiv \theta', \tag{24}$$

where (21) and (23) follow from Items 1 and 2, and where (22) and (24) yield $\theta, \theta' \in \text{LTL}_B(\{\mathbf{U}\})$ by induction hypothesis. \square

Theorem 7: $\text{LTL}_B(\{\mathbf{U}\})$ and $\text{LTL}_B(\{\mathbf{G}, \vee\})$ are undecidable.

Proof for $LTL_B(\{\mathbf{G}, \vee\})$: Let \mathcal{N} be a Petri net with inhibitor arcs and let $\mathbf{x}_{\text{src}}, \mathbf{x}_{\text{tgt}}$. Let $M, \mathbf{x}, \mathbf{x}'$ and AP be given by Proposition 24. We modify M and AP as follows. We add dimension \heartsuit to indicate that \mathbf{x}' was reached and extend the execution to an infinite one.

Each mode $\mathbf{m} \in M$ is extended with $\mathbf{m}(\heartsuit) := 0$. Each zone of AP is extended with the constraint $\heartsuit = 0$. We add modes $\{\mathbf{a}_\heartsuit, \overline{\mathbf{a}_\heartsuit}\}$ and zone A_\heartsuit defined by:

j	$\mathbf{a}_\heartsuit(j)$	$\overline{\mathbf{a}_\heartsuit}(j)$	A_\heartsuit
\heartsuit	1	-1	$\heartsuit \in [0, 1]$
else	0	0	rest = \mathbf{x}'

Let M' and AP' be the resulting MMS and set of zones. Let $\varphi := \mathbf{G}(\bigvee_{Z \in AP'} Z)$. Let $\mathbf{y} := (0, \mathbf{x})$. By Proposition 24, it suffices to show that $\mathbf{x} \rightarrow_{AP}^* \mathbf{x}'$ in M iff $\mathbf{y} \models_{M'} \varphi$.

\Rightarrow) Let π be a finite schedule such that $\mathbf{x} \rightarrow_{AP}^\pi \mathbf{x}'$ in M . Let $\pi' := \pi \mathbf{a}_\heartsuit \overline{\mathbf{a}_\heartsuit} \mathbf{a}_\heartsuit \overline{\mathbf{a}_\heartsuit} \dots$. We have $\mathbf{y} \rightarrow_{AP'}^{\pi'}$ in M' , and hence $\text{exec}(\pi', \mathbf{y}) \models_{M'} \varphi$, since π' is non-Zeno.

\Leftarrow) Let π be an infinite non-Zeno schedule of M' such that $\text{exec}(\pi, \mathbf{y}) \models \varphi$. Let $\sigma := \text{exec}(\pi, \mathbf{y})$. By Proposition 24, we have $\mathbf{y} \not\rightarrow_{AP' \setminus \{A_\heartsuit\}}^\pi$.

So, there exists $\tau \in \mathbb{R}_{\geq 0}$ such that $\sigma(\tau) \in A_\heartsuit$. By definition of M' , there is a minimal $\tau' \leq \tau$ with $\sigma(\tau') = (0, \mathbf{x}')$. Let π' be a finite schedule such that $\text{exec}(\pi', \mathbf{y}) = \sigma[0..\tau']$. We have $\mathbf{y} \rightarrow_{AP' \setminus \{A_\heartsuit\}}^{\pi'} (0, \mathbf{x}')$ in M' . Hence, $\mathbf{x} \rightarrow_{AP}^* \mathbf{x}'$ in M . \square