

Logics for Continuous Reachability in Petri Nets and Vector Addition Systems with States

Michael Blondin

Technische Universität München, Germany

blondin@in.tum.de

Christoph Haase

University of Oxford, United Kingdom

christoph.haase@cs.ox.ac.uk

Abstract—This paper studies sets of rational numbers definable by continuous Petri nets and extensions thereof. First, we identify a polynomial-time decidable fragment of existential $\text{FO}(\mathbb{Q}, +, <)$ and show that the sets of rationals definable in this fragment coincide with reachability sets of continuous Petri nets. Next, we introduce and study continuous vector addition systems with states (CVASS), which are vector addition systems with states in which counters may hold non-negative rational values, and in which the effect of a transition can be scaled by a positive rational number smaller or equal to one. This class strictly generalizes continuous Petri nets by additionally allowing for discrete control-state information. We prove that reachability sets of CVASS are equivalent to the sets of rational numbers definable in existential $\text{FO}(\mathbb{Q}, +, <)$ from which we can conclude that reachability in CVASS is NP-complete. Finally, our results explain and yield as corollaries a number of polynomial-time algorithms for decision problems that have recently been studied in the literature.

I. INTRODUCTION

Petri nets and *vector addition systems with states* (VASS) are versatile classes of infinite-state systems that find numerous applications in diverse areas of computer science, especially in formal verification. A Petri net comprises a finite number of places, which carry a non-negative number of tokens, and a finite set of transitions. A transition may add and remove tokens from places, provided that the resulting number of tokens remains non-negative in each place. Similarly, a VASS consists of a finite-state controller with a finite number of counters ranging over the natural numbers. When taking a transition, counters may be incremented and decremented, provided that the resulting counter values are non-negative. Petri nets and VASS are computationally equivalent, and their combinatorial structure is remarkably difficult to analyze. Despite significant progress and efforts that have been made over the course of the last forty years, reachability in Petri nets and VASS is only known to be EXPSpace-hard [1] and decidable in non-primitive recursive time [2]. In order to circumvent the computational difficulties encountered when dealing with these models, relaxed semantics have been proposed in the literature. An obvious relaxation consists in allowing the number of tokens or counter values to drop below zero. This is known as the *state equation* for Petri nets, and as \mathbb{Z} -VASS (see e.g. [3]) for vector addition systems with states. Since non-negativity constraints are usually crucial in modeling, researchers have also studied relaxed models in which the number of tokens or counter values must remain non-negative, but may range

over the real numbers. For example, David and Alla [4], [5] introduced *continuous Petri nets*, which are Petri nets in which a transition can be fired by a positive real factor, and in which places may contain a non-negative real number of tokens.

The computational and practical benefits obtained from relaxing the semantics of Petri nets and VASS are remarkable. Continuous Petri nets, in particular, are a unique class of infinite-state systems from a complexity-theoretic, logical and practical point of view: reachability in continuous Petri nets is only P-complete [6]; continuous Petri nets reachability relations over the rational numbers are definable in existential $\text{FO}(\mathbb{Q}, +, <)$ [7] as opposed to their discrete counterparts whose reachability sets are not even semi-linear [8]; and reachability in continuous Petri nets can consequently be decided using SMT solvers [7], while only limited tool support for reachability in discrete Petri nets exists. Moreover, continuous Petri nets as well as \mathbb{Z} -VASS, for which reachability is NP-complete [3], over-approximate the behavior and the reachability sets of their classical counterparts. Those over-approximations have successfully been employed inside decision procedures for discrete Petri nets and VASS in order to prune search spaces. For instance, in an approach proposed by the authors [7], which has recently been refined by Geffroy, Leroux and Sutre [9], reachability in continuous Petri nets is used as a pruning criterion inside the classical backward algorithm [10] for deciding coverability in Petri nets. Furthermore, integer over-approximations of VASS and Petri nets have been used as semi-decision procedures for coverability in VASS [11], [12]. In all such application scenarios, those over-approximations have led to decision procedures for coverability problems in VASS and Petri nets that yield an empirically outstanding performance on real-world instances [11], [7], [12], [9].

There are two aspects of continuous Petri nets that emerge from the aforementioned literature that have not yet been understood in a satisfactory way. First, whenever the authors presented their SMT-based decision procedure for reachability in continuous Petri nets [7], members of the audience often wondered why one would choose to solve a polynomial-time decidable problem by a translation into existential $\text{FO}(\mathbb{Q}, +, <)$, an NP-complete theory [13]. Even though there are good reasons for doing so, e.g. the solid engineering work that has gone into modern SMT solvers, fundamentally this is a valid concern. Phrased differently, the actual question is

whether there exists a polynomial-time decidable fragment of existential $\text{FO}(\mathbb{Q}, +, <)$ that can express reachability in continuous Petri nets, and on which SMT solvers are guaranteed to empirically perform well. Second, \mathbb{Z} -VASS have recently, and to some extent surprisingly, shown to provide a more suitable over-approximation than (subclasses of) continuous Petri nets when applied to VASS coverability instances obtained from the verification of concurrent Pthread-based C programs [12]. Essentially, the reason for this is that \mathbb{Z} -VASS allow for maintaining discrete control structures, whereas continuous Petri nets do not allow for *any* sort of discrete information. This shortcoming of continuous Petri nets has long been known, and led to the introduction of *hybrid Petri nets* in which only some places and transitions may operate under the relaxed continuous semantics [5]. However, reachability in hybrid Petri nets is, in the worst case, as difficult as in standard Petri nets. This raises the question of whether there exists an extension of continuous Petri nets that allows for maintaining discrete control structures while sharing the striking algorithmic and structural properties of continuous Petri nets.

Our contribution: We take a unified approach to the two questions identified in the previous paragraph. First, in Section III-A we identify a fragment of existential $\text{FO}(\mathbb{Q}, +, <)$ whose satisfiability problem is P-complete. Such fragments have already been studied in the community working on constraint satisfaction problems (CSPs). Our contribution is to introduce the CSP of convex semi-linear Horn formulas ($\text{CSP}(\Gamma_c)$), and to provide a polynomial-time decision procedure for it. We then show in Section III-B that reachability in continuous Petri nets reduces to $\text{CSP}(\Gamma_c)$, and conversely in Section III-C that the set of satisfying assignments of an instance of $\text{CSP}(\Gamma_c)$ can be obtained as the reachability set of a continuous Petri net. Thus, $\text{CSP}(\Gamma_c)$ fully characterizes the reachability sets of continuous Petri nets. In addition, in Section III-D we remark that our $\text{CSP}(\Gamma_c)$ -based approach to reachability in continuous Petri nets immediately recovers further polynomial-time results for which dedicated algorithms have been developed in the literature.

Our next contribution in Section IV is to introduce *continuous vector addition systems with states (CVASS)*, a class of hybrid Petri nets. Such CVASS are like VASS except that when a transition is taken, its effect on the counters can be scaled by any positive rational number less or equal to one, provided that the resulting counter values remain non-negative. We first show the not immediate fact that continuous Petri nets embed into CVASS in the same way as Petri nets embed into VASS. We then prove in Section IV-B that reachability relations of CVASS are definable in existential $\text{FO}(\mathbb{Q}, +, <)$, and hence that reachability in CVASS is decidable in NP. Analogously to the case of continuous Petri nets, we show in Section IV-C that the sets of solutions of an existential $\text{FO}(\mathbb{Q}, +, <)$ formula is definable as the reachability set of a CVASS. In Section IV-D, we give reductions which essentially show that adding any sort of discrete control structure to continuous Petri nets makes reachability NP-hard. Finally, we

discuss our results in Section VI and give an outlook on possible future work.

Rational versus real arithmetic: Continuous Petri nets are named this way because transitions may be fired by a non-negative *real* number. However, for technical convenience, we only consider firing transitions by rational numbers throughout this paper. That way, we avoid problems such as how to finitely represent arbitrary real numbers. Almost all of our constructions and proofs only depend on the denseness of the domain, and where this is not the case we explicitly remark how to make an argument work when allowing transitions to be fired a real number of times. Also note that all decision problems we consider in this paper reduce to existential $\text{FO}(\mathbb{Q}, +, <)$, and existential $\text{FO}(\mathbb{Q}, +, <)$ and existential $\text{FO}(\mathbb{R}, +, <)$ are equivalent theories with an NP-complete satisfiability problem [13].

II. PRELIMINARIES

We denote respectively by \mathbb{Z} , \mathbb{N} , \mathbb{Q} , \mathbb{Q}_+ and $\mathbb{Q}_{>0}$ the sets of integers, non-negative integers, rational numbers, non-negative rational numbers and positive rational numbers. For $a, b \in \mathbb{Z}$, we let $[a, b] \stackrel{\text{def}}{=} \{x \in \mathbb{Z} : a \leq x \leq b\}$, and $[a] \stackrel{\text{def}}{=} [1, a]$. Intervals over \mathbb{Q} are denoted in the standard way. We write vectors $(v_1, v_2, \dots, v_d) \in \mathbb{Q}^d$ in bold as \mathbf{v} , and write $\mathbf{v}(i)$ and v_i for v_i . Whenever \mathbf{v}_i is also a vector, we sometimes write $\mathbf{v}_{i,j}$ to denote $(\mathbf{v}_i)_j$. Vectors are ordered component-wise, i.e., $\mathbf{u} \leq \mathbf{v}$ if and only if $\mathbf{u}(i) \leq \mathbf{v}(i)$ for every $i \in [d]$. Given a finite index set I and $\mathbb{A} \subseteq \mathbb{Q}$, we denote by \mathbb{A}^I the set of all *I-indexed vectors* $\mathbf{v} : I \rightarrow \mathbb{A}$. With no loss of generality, we may assume a total order on I , and hence for I such that $|I| = d$ we sometimes write elements from \mathbb{A}^I as tuples $\mathbf{v} = (v_1, \dots, v_d) \in \mathbb{A}^d$. We write $\mathbf{0}$ for the null vector in any dimension, and \mathbf{e}_i for the *i*-th unit vector, i.e., \mathbf{e}_i is such that $\mathbf{e}_i(i) = 1$ and $\mathbf{e}_i(j) = 0$ for every $j \neq i$. The *support* of a vector $\mathbf{v} \in \mathbb{Q}^d$ is $\llbracket \mathbf{v} \rrbracket \stackrel{\text{def}}{=} \{i \in [d] : \mathbf{v}(i) \neq 0\}$. Similarly, the *positive support* and *negative support* of \mathbf{v} are defined respectively as

$$\llbracket \mathbf{v} \rrbracket^+ \stackrel{\text{def}}{=} \{i \in [d] : \mathbf{v}(i) > 0\} \quad \llbracket \mathbf{v} \rrbracket^- \stackrel{\text{def}}{=} \{i \in [d] : \mathbf{v}(i) < 0\}.$$

A. Continuous Petri nets

A *continuous Petri net* is a tuple $\mathcal{N} = (P, T, \mathbf{Pre}, \mathbf{Post})$, where P and T are disjoint finite sets, and $\mathbf{Pre}, \mathbf{Post} \in \mathbb{N}^{P \times T}$. The sets P and T are respectively called the *places* and *transitions* of \mathcal{N} . The *transpose* of \mathcal{N} is defined as $\mathcal{N}^\dagger \stackrel{\text{def}}{=} (P, T, \mathbf{Post}, \mathbf{Pre})$. The *size* of \mathcal{N} is the number of symbols required to write down \mathcal{N} , where we assume binary encoding of numbers and that transitions are represented as lists of non-negative components of $\mathbb{N}^{P \times T}$.

Let $p \in P$ and $t \in T$. The *pre-sets* of p and t are respectively defined as

- $p \stackrel{\text{def}}{=} \{t' \in T : \mathbf{Post}(p, t') > 0\}$,
- $t \stackrel{\text{def}}{=} \{p' \in P : \mathbf{Pre}(p', t) > 0\}$.

Similarly, the *post-sets* of p and t are respectively

$$\begin{aligned} p^\bullet &\stackrel{\text{def}}{=} \{t' \in T : \mathbf{Pre}(p, t') > 0\}, \\ t^\bullet &\stackrel{\text{def}}{=} \{p' \in P : \mathbf{Post}(p', t) > 0\}. \end{aligned}$$

We extend these definitions to sets, e.g., for every $Q \subseteq P$, we let $\bullet Q \stackrel{\text{def}}{=} \bigcup_{p \in Q} \bullet p$.

For every $t \in T$, \mathbf{Pre}_t and \mathbf{Post}_t denote the column of \mathbf{Pre} of and \mathbf{Post} associated to t . The *effect* of a transition is $\delta(t) \stackrel{\text{def}}{=} \mathbf{Post}_t - \mathbf{Pre}_t$. For every $\mathbb{A} \in \{\mathbb{Q}, \mathbb{Q}_+\}$, $\alpha \in (0, 1]$, and $t \in T$, we define the *transition relation* $\xrightarrow{\alpha t}_{\mathbb{A}}$ over \mathbb{A}^P by

$$\mathbf{u} \xrightarrow{\alpha t}_{\mathbb{A}} \mathbf{v} \stackrel{\text{def}}{\iff} \mathbf{u} \geq \alpha \cdot \mathbf{Pre}_t \text{ and } \mathbf{v} = \mathbf{u} + \alpha \cdot \delta(t).^1$$

For integral \mathbf{u} and \mathbf{v} , we write $\mathbf{u} \xrightarrow{t}_{\mathbb{N}} \mathbf{v}$ for $\mathbf{u} \xrightarrow{1t}_{\mathbb{Q}_+} \mathbf{v}$, which corresponds to the discrete Petri net semantics. To simplify the notation, we define $T^\star \stackrel{\text{def}}{=} ((0, 1] \times T)^\star$. The identity relation over \mathbb{A}^P is denoted by $\xrightarrow{\varepsilon}_{\mathbb{A}}$, and for $\pi = \alpha_1 t_1 \alpha_2 t_2 \cdots \alpha_k t_k \in T^\star$, we let

$$\xrightarrow{\pi}_{\mathbb{A}} \stackrel{\text{def}}{=} \xrightarrow{\alpha_k t_k}_{\mathbb{A}} \circ \cdots \circ \xrightarrow{\alpha_2 t_2}_{\mathbb{A}} \circ \xrightarrow{\alpha_1 t_1}_{\mathbb{A}}.$$

The *reachability relation* $\mathcal{R}(\mathcal{N})$ of \mathcal{N} is defined as

$$\xrightarrow{\ast}_{\mathbb{A}} \stackrel{\text{def}}{=} \bigcup_{\pi \in T^\star} \xrightarrow{\pi}_{\mathbb{A}}.$$

The *effect*, the (*rational*) *Parikh image* $\Psi(\pi): T \rightarrow \mathbb{Q}_+$, and the *support* of π are defined respectively as

$$\begin{aligned} \delta(\pi) &\stackrel{\text{def}}{=} \alpha_1 \cdot \delta(t_1) + \alpha_2 \cdot \delta(t_2) + \cdots + \alpha_n \cdot \delta(t_n), \text{ and} \\ \Psi(\pi)(t) &\stackrel{\text{def}}{=} \sum_{\substack{1 \leq i \leq n \\ t_i = t}} \alpha_i \text{ for every } t \in T, \\ \llbracket \pi \rrbracket &\stackrel{\text{def}}{=} \{t_i : i \in [n]\}. \end{aligned}$$

We write $\alpha\pi = (\alpha \cdot \alpha_1)t_1(\alpha \cdot \alpha_2)t_2 \cdots (\alpha \cdot \alpha_n)t_n$ for every $\alpha \in (0, 1]$. A sequence $\pi \in T^\star$ is a *firing sequence from* $\mathbf{u} \in \mathbb{Q}_+^P$ if $\mathbf{u} \xrightarrow{\pi}_{\mathbb{Q}_+} \mathbf{u} + \delta(\pi)$. We say that a subset of transitions $S \subseteq T$ is a *firing set* from $\mathbf{u} \in \mathbb{Q}_+$ if there exists a firing sequence π from \mathbf{u} such that $\llbracket \pi \rrbracket = S$.

At a later stage, we will be interested in sets $M \subseteq \mathbb{Q}_+^d$ definable by continuous Petri nets. In general, it is unlikely to expect that all places of a continuous Petri net can contribute towards defining M , as some of them may need to be used, for instance, for temporarily storing tokens. Therefore, we use the following definition: a set $M \subseteq \mathbb{Q}_+^n$ is *definable* by a continuous Petri net \mathcal{N} with places P such that $|P| = m + n$ whenever there exists a fixed $\mathbf{u} \in \mathbb{Q}_+^m$ such that $\mathbf{v} \in M$ if and only if $(\mathbf{v}, \mathbf{u}) \xrightarrow{\ast}_{\mathbb{Q}_+} (\mathbf{0}, \mathbf{0})$. Alternatively, M is definable by \mathcal{N} if it is a projection of a cross section of $\mathcal{R}(\mathcal{N}^\dagger)(\mathbf{0})$.

¹In the literature, α is often taken from $\mathbb{Q}_{>0}$ instead of $(0, 1]$. However, it is easily seen that for reachability purposes either choice gives the same reachability sets.

B. Linear rational arithmetic

We show decidability and complexity results via reductions into existential $\text{FO}(\mathbb{Q}, +, <)$, the existential fragment of the first-order theory of the rationals with addition and order. Formulas of this theory are built from positive Boolean combinations of atomic formulas of the form $\mathbf{a}^\top \cdot \mathbf{x} \sim c$, where $\mathbf{x} = (x_1, \dots, x_n)$ are first-order variables, $\mathbf{a} \in \mathbb{Q}^n$, $c \in \mathbb{Q}$, and $\sim \in \{>, \geq, =\}$. We write $\phi(\mathbf{x})$ in order to indicate that the variables $\mathbf{x} = (x_1, \dots, x_n)$ occur free in ϕ , and $\llbracket \phi \rrbracket \subseteq \mathbb{Q}^n$ for the set of rational numbers that can be extended to solutions of ϕ . The size of a formula is the number of symbols required to write it down, where we assume binary encoding of numbers.

Satisfiability in existential $\text{FO}(\mathbb{Q}, +, <)$ is NP-complete [13]. If we only consider conjunctions of atomic formulas, we actually decide the feasibility problem for linear programming, which is known to be hard for P [14, Prob. A.4.1] and decidable in polynomial time [15], even in the presence of strict inequalities, see e.g. [16, Chap. 8.7.1].

III. A LOGICAL CHARACTERIZATION OF REACHABILITY IN CONTINUOUS PETRI NETS

The goal of this section is to identify a fragment of existential $\text{FO}(\mathbb{Q}, +, <)$ that fully characterizes reachability in continuous Petri nets. We introduce this fragment in Section III-A below. It turns out that our fragment is closely related to semi-linear CSPs, and we present our fragment as a particular semi-linear CSP: the CSP of convex linear Horn constraints. In particular, we show that its feasibility problem is P-complete. Next, in Section III-B we provide a reduction from reachability in continuous Petri nets to feasibility in the CSP of semi-linear Horn constraints. This yields as a corollary an alternative proof that reachability in continuous Petri nets is decidable in P. Conversely, in Section III-C we show that the set of feasible solutions of an instance of a CSP of linear Horn constraints can be obtained as the set of reachable markings of a continuous Petri net. To round up the picture, in Section III-D we discuss applications of our results as well as relationships to other models that have been studied in the literature.

A. The CSP of convex linear Horn constraints

In this section, we identify a syntactic fragment of existential $\text{FO}(\mathbb{Q}, +, <)$ that generalizes linear programming, whose sets of solutions are convex, and for which satisfiability is P-complete. In fact, our fragment is closely related to the CSP of linear Horn constraints studied by Jonsson and Bäckström [17], and Koubarakis [18]. Due to the close relationship to CSPs and for technical convenience, we present our fragment of existential $\text{FO}(\mathbb{Q}, +, <)$ in the language of CSPs.

Given a *domain* \mathbb{D} , a set of *relation symbols* $\tau = \{R_1, R_2, \dots\}$, and a τ -*structure* $\Gamma = (\mathbb{D}; R_1^\Gamma, R_2^\Gamma, \dots)$ such that $R_i^\Gamma \subseteq \mathbb{D}^{k_i}$ is of arity k_i , the problem $\text{CSP}(\Gamma)$ is defined as follows. Given a finite set of variables $X = \{x_1, \dots, x_n\}$, and a finite set of expressions $R(x_{i_1}, \dots, x_{i_k})$ such that $R \in \tau$, is there an assignment $\sigma: X \rightarrow \mathbb{D}$ such that $(\sigma(x_{i_1}), \dots, \sigma(x_{i_k})) \in R^\Gamma$ for all constraints $R(x_{i_1}, \dots, x_{i_k})$

of the input? Whenever $D = \mathbb{Q}$ and all R_i are definable by a quantifier-free formula in $\text{FO}(\mathbb{Q}, +, <)$, Γ is called semi-linear.²

Convex semi-linear relations over the non-negative rational numbers are defined by *convex semi-linear Horn clauses*, which are quantifier-free $\text{FO}(\mathbb{Q}, +, <)$ formulas of the form

$$\psi = \mathbf{a}^\top \cdot \mathbf{x} \sim c \vee \bigvee_{1 \leq i \leq m} \theta_i, \text{ where } \theta_i = \bigwedge_{j \in J_i} x_j > 0$$

and where J_1, J_2, \dots, J_m are finite sets of indices. Given a convex semi-linear Horn clause ψ as above, the *negative literals* of ψ are $\text{neg}(\psi) \stackrel{\text{def}}{=} \{\theta_1, \dots, \theta_m\}$, and $\text{fact}(\psi) \stackrel{\text{def}}{=} \mathbf{a}^\top \cdot \mathbf{x} \sim c$. Given a negative literal $\theta \in \text{neg}(\psi)$, we denote by $\psi \setminus \theta$ the convex semi-linear Horn clause ψ' which is obtained from ψ by removing θ from its set of negative literals. A *convex semi-linear Horn formula* is a finite set Φ of convex semi-linear Horn clauses. We define $\text{facts}(\Phi)$ to be the set of all *semi-linear facts* in Φ such that $\psi \in \text{facts}(\Phi)$ if and only if $\text{neg}(\psi) = \emptyset$. Whenever it is convenient, we also write convex semi-linear Horn formulas as conjunctions of convex semi-linear Horn clauses. In the following, we denote by $\Gamma_c = (\mathbb{Q}_+; R_1, R_2, \dots)$ the structure whose domain is the set of non-negative rational numbers and whose relations are all convex semi-linear relations.

We now study some properties of $\text{CSP}(\Gamma_c)$. First of all, in order to make our life a bit easier, we introduce the following abbreviations that will later become handy:

$$x > 0 \rightarrow \bigvee_{1 \leq i \leq m} \theta_i \iff x = 0 \vee \bigvee_{1 \leq i \leq m} \theta_i \quad (1)$$

$$x = 0 \rightarrow \bigwedge_{1 \leq i \leq m} y_i = 0 \iff \bigwedge_{1 \leq i \leq m} y_i = 0 \vee x > 0. \quad (2)$$

As the name suggests, solutions of the CSP of convex semi-linear Horn constraints are convex, as we show in the following lemma.

Lemma 3.1: The set of solutions of an instance Φ of $\text{CSP}(\Gamma_c)$ is convex.

Proof: Let $\mathbf{y}, \mathbf{z} \in \mathbb{Q}_+^n$ be two solutions of Φ . We claim that $\mathbf{x}^* = \lambda \cdot \mathbf{y} + (1 - \lambda) \cdot \mathbf{z}$ is a solution of Φ for all rational $\lambda \in [0, 1]$. This is clear for all semi-linear facts. Now consider a general semi-linear Horn clause ψ in which, with no loss of generality, \mathbf{z} makes some negative literal $\theta_i = \bigwedge_{j \in J_i} x_j > 0$ true. Note that for any $\lambda \in [0, 1)$, we have $\mathbf{x}_j^* > 0$ as well due to interpreting variables only over the non-negative integers. Hence for such choices of λ , θ_i also evaluates to true under \mathbf{x}^* , and so does ψ . If $\lambda = 1$, then $\mathbf{x}^* = \mathbf{y}$ is a solution of ψ by assumption. \square

Even though the sets of rationals definable in $\text{CSP}(\Gamma_c)$ are convex, in general they cannot be obtained as points inside a convex polytope.

Lemma 3.2: There exists an infinite family of convex linear Horn constraints $(\psi_n(x, y_1, \dots, y_n))_{n > 0}$ such that no convex polytope defines any $\llbracket \psi_n \rrbracket$.

²Here, we adopt the notion of semi-linearity used in the CSP community. Probably more commonly, semi-linear sets are also known as the sets of natural numbers definable in $\text{FO}(\mathbb{N}, +, <)$.

Algorithm 1 Decision procedure for $\text{CSP}(\Gamma_c)$.

INPUT: A convex semi-linear Horn formula Φ .

```

1: repeat
2:    $\Psi := \text{facts}(\Phi)$ ;  $\text{halt} := \text{true}$ 
3:   if  $\Psi$  is unsatisfiable in  $\mathbb{Q}_+$  then return unsatisfiable
4:   for all  $\psi \in \Phi \setminus \text{facts}(\Phi)$  do
5:      $\psi' := \psi$ 
6:     for all  $\theta \in \text{neg}(\psi)$  do
7:       if  $\Psi \cup \{\theta\}$  is unsatisfiable in  $\mathbb{Q}_+$  then
8:          $\psi' := \psi' \setminus \theta$ 
9:          $\text{halt} := \text{false}$ 
10:     $\Phi := (\Phi \setminus \{\psi\}) \cup \{\psi'\}$ 
11: until  $\text{halt}$ 
12: return satisfiable

```

Proof: It is easily checked that $\psi_n(x, y_1, \dots, y_n) \stackrel{\text{def}}{=} (x = 0 \vee y_1 > 0 \vee \dots \vee y_n > 0)$ is a family of convex linear Horn constraints with the desired properties. \square

Finally, we show that $\text{CSP}(\Gamma_c)$ is decidable in polynomial time via Algorithm 1, an adaptation of the classical algorithm deciding satisfiability of Boolean Horn formulas.

Theorem 3.3: $\text{CSP}(\Gamma_c)$ is P-complete.

Proof: Hardness follows from P-hardness of linear programming feasibility [14, Prob. A.4.1]. It thus remains to be shown that Algorithm 1 is a decision procedure for $\text{CSP}(\Gamma_c)$ that runs in polynomial time.

We begin with showing that the algorithm runs in polynomial time. The satisfiability checks in Lines 3 and 7 can be performed in polynomial time due to linear programming being decidable in polynomial time [15], [16]. Thus, it remains to show that the repeat-loop is traversed at most a polynomial number of times. In every iteration of the loop, at least one negative literal of some semi-linear Horn clause gets removed, and thus the loop is iterated at most $O(m)$ times, where m is the number of negative literals in Φ . This in turn implies an overall polynomial running time of Algorithm 1.

Regarding the correctness of the algorithm, if the algorithm returns *unsatisfiable* in Line 3, then Φ contains an unsatisfiable subset of facts, and hence is unsatisfiable. Moreover, if $\Psi \cup \{\theta\}$ is unsatisfiable in Line 7, then Ψ implies $-\theta$, and hence the updated Φ in Line 10 in which all such θ have been removed has the same set of solutions as before. Let Φ' and Ψ' be respectively the convex semi-linear Horn formula and the set of facts obtained in the algorithm once Line 12 has been reached. For every convex semi-linear Horn clause $\psi_i = \mathbf{a}_i^\top \cdot \mathbf{x} \sim c_i \vee \bigvee_{1 \leq j \leq m_i} \theta_{i,j}$ in Φ' , the algorithm guarantees that there are solutions $\mathbf{x}_{i,j}$ of $\Psi' \cup \{\theta_{i,j}\}$. We claim that

$$\mathbf{x}^* \stackrel{\text{def}}{=} \frac{1}{n} \cdot \sum_{\substack{\psi_i \in \Phi' \\ 1 \leq j \leq m_i}} \mathbf{x}_{i,j}$$

is a solution of Φ' . Indeed, \mathbf{x}^* is a solution of all facts in Ψ' , cf. Lemma 3.1. Moreover, for any $x > 0$ of some $\theta_{i,j}$, we have $x_{i,j} > 0$ and consequently $x^* > 0$, since variables are

interpreted over the non-negative rational numbers, resembling the same argument used in the proof of Lemma 3.1. Consequently, \mathbf{x}^* is a solution of Φ' and hence of Φ . Thus, Φ is satisfiable when the algorithm return *satisfiable* in Line 12. \square

Before we conclude this section, let us briefly discuss the relationship of our CSP of convex linear Horn constraints to other CSPs that have been studied in the literature. To the best of our knowledge, no convex tractable extension of linear programming has been identified so far. The tractable CSP of linear Horn constraints studied in [17], [18], [19] allows for constraints of the form $\mathbf{a}^T \cdot \mathbf{x} \sim c \vee \bigvee_{1 \leq i \leq m} \mathbf{a}_i^T \cdot \mathbf{x} \neq c_i$, where \mathbf{x} is interpreted over the rational numbers. While this CSP can express constraints of $\text{CSP}(\Gamma_c)$ in which every negative literal is a single strict inequality, it is not obvious to us whether the CSP of linear Horn constraints could be generalized such that $\text{CSP}(\Gamma_c)$ fully embeds into it. In any case, solutions to instances of the CSP of linear Horn constraints are not convex in general. Finally, it is worth mentioning that our convex semi-linear Horn formulas can also be viewed as convex Horn formulas, defined in [19] and allowing for restricted classes of polynomials as negative literals, by observing that over \mathbb{Q}_+ , $\bigwedge_{1 \leq i \leq n} x_i > 0 \equiv \prod_{1 \leq i \leq n} x_i > 0$. However, a P upper bound cannot be derived from the results in [19].

B. Reducing reachability in continuous Petri nets to $\text{CSP}(\Gamma_c)$

We now show that reachability sets of continuous Petri nets are definable in $\text{CSP}(\Gamma_c)$. Given a continuous Petri net \mathcal{N} , we show how to construct in logarithmic space an instance Φ of $\text{CSP}(\Gamma_c)$ such that a projection of $\llbracket \Phi \rrbracket$ corresponds to the reachability set of \mathcal{N} . Our reduction is based on the following characterization of reachability in continuous Petri nets.

Proposition 3.4 ([6, Thm. 20]): Let $\mathcal{N} = (P, T, \text{Pre}, \text{Post})$ be a continuous Petri net. Let $\mathbf{u}, \mathbf{v} \in \mathbb{Q}_+^P$ and $\mathbf{x} \in \mathbb{Q}_+^T$. There exists $\pi \in T^*$ such that $\mathbf{u} \xrightarrow{\pi} \mathbb{Q}_+ \mathbf{v}$ and $\Psi(\pi) = \mathbf{x}$ if and only if

- (i) $\mathbf{v} = \mathbf{u} + (\text{Post} - \text{Pre}) \cdot \mathbf{x}$;
- (ii) $\llbracket \mathbf{x} \rrbracket$ is a firing set from \mathbf{u} in \mathcal{N} ; and
- (iii) $\llbracket \mathbf{x} \rrbracket$ is a firing set from \mathbf{v} in \mathcal{N}^\dagger .

The first item requires the existence of a solution \mathbf{x} to the Petri net state equation, and the subsequent two items require that the support of \mathbf{x} has a corresponding firing sequence in both \mathcal{N} and its transpose. In [7], the authors developed a translation of the conditions of Proposition 3.4 into existential $\text{FO}(\mathbb{Q}, +, <)$. There, the encoding of the Conditions (ii) and (iii) of Proposition 3.4 is achieved by adapting an approach proposed by Verma, Seidl and Schwentick for characterizing the Parikh image of the language of a context-free grammar in existential Presburger arithmetic [20]. However, the construction in [7] does not yield a formula in $\text{CSP}(\Gamma_c)$. While the Petri net state equation clearly is expressible in $\text{CSP}(\Gamma_c)$, the encoding used for Conditions (ii) and (iii) of Proposition 3.4 does not lie in $\text{CSP}(\Gamma_c)$.

Lemma 3.5: There exists a logarithmic-space computable instance $\Phi(\mathbf{u}, \mathbf{y})$ of $\text{CSP}(\Gamma_c)$ such that $(\mathbf{u}, \mathbf{y}) \in \llbracket \Phi \rrbracket$ if and only if $\llbracket \mathbf{y} \rrbracket$ is a firing set from \mathbf{u} in \mathcal{N} .

Proof: The $\text{CSP}(\Gamma_c)$ instance $\Phi(\mathbf{u}, \mathbf{y})$ that we define uses auxiliary vectors of variables $\mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_{|T|} \in \mathbb{Q}_+^{P \cup T}$. Similarly as in [20], the idea behind our construction is to successively explore which places could carry tokens and which transitions could in turn be enabled after a certain number of steps. By $x_{i,p} > 0$ and $x_{i,t} > 0$ we indicate that after at most i explorations steps the place p could carry tokens, respectively that the transition t could be enabled. In the definitions of the subsequent formulas, we make heavy use of the abbreviations (1) and (2) in order to stay inside $\text{CSP}(\Gamma_c)$.

Initially only places in $\llbracket \mathbf{u} \rrbracket$ carry tokens and no transition is enabled:

$$\psi_0(\mathbf{u}, \mathbf{x}_0) \stackrel{\text{def}}{=} \bigwedge_{p \in P} (x_{0,p} > 0 \leftrightarrow u_p > 0) \wedge \bigwedge_{t \in T} x_{0,t} = 0.$$

Next, we declare that after i steps, places that carried tokens after $i-1$ steps continue to have tokens, as well as places that receive tokens from transitions that were enabled after $i-1$ steps. Moreover, if a transition t is declared enabled after i steps, we require that all places from which t consumes tokens carry tokens after $i-1$ steps:

$$\begin{aligned} \psi_i(\mathbf{x}_i, \mathbf{x}_{i-1}) \stackrel{\text{def}}{=} & \bigwedge_{p \in P} (x_{i,p} > 0 \leftrightarrow (x_{i-1,p} > 0 \vee \bigvee_{t \in \bullet p} x_{i,t} > 0)) \wedge \\ & \bigwedge_{t \in T} (x_{i,t} > 0 \rightarrow \bigwedge_{p \in \bullet t} x_{i-1,p} > 0). \end{aligned} \quad (3)$$

Finally, the instance $\Phi(\mathbf{u}, \mathbf{y})$ of $\text{CSP}(\Gamma_c)$ required in the lemma is obtained as the following conjunction:

$$\begin{aligned} \Phi(\mathbf{u}, \mathbf{y}) \stackrel{\text{def}}{=} & \psi_0(\mathbf{u}, \mathbf{x}_0) \wedge \bigwedge_{1 \leq i \leq |T|} \psi_i(\mathbf{x}_i, \mathbf{x}_{i-1}) \wedge \\ & \bigwedge_{t \in T} (\mathbf{y}_t > 0 \leftrightarrow \bigvee_{1 \leq i \leq |T|} x_{i,t} > 0). \end{aligned}$$

It is not difficult to check that $\Phi(\mathbf{u}, \mathbf{y})$ is an instance of $\text{CSP}(\Gamma_c)$. We can prove the correctness of the construction by induction on $\llbracket \mathbf{y} \rrbracket$. This is not straightforward since (3) does not enforce a transition t to be enabled whenever there are tokens in every place of $\bullet t$. We omit details for brevity and refer the reader to Appendix A-A for further details. \square

The reason for the seemingly cumbersome definition of when transitions become enabled in (3) is that we cannot express $x_{i,t} > 0 \leftrightarrow (\bigwedge_{p \in \bullet t} x_{i-1,p} > 0)$ in $\text{CSP}(\Gamma_c)$. Intuitively this is the main reason why our previously developed existential $\text{FO}(\mathbb{Q}, +, <)$ characterization of continuous Petri net reachability [7] does not yield a formula in $\text{CSP}(\Gamma_c)$. The price to pay for staying inside $\text{CSP}(\Gamma_c)$ is that the formula constructed in Lemma 3.5 is quadratic in the size of \mathcal{N} , whereas in [7] the formula is linear. Combining Proposition 3.4 with Lemma 3.5 now gives us the following theorem.

Theorem 3.6: There exists a logarithmic-space computable instance $\Phi(\mathbf{u}, \mathbf{v}, \mathbf{w})$ of $\text{CSP}(\Gamma_c)$ such that $(\mathbf{u}, \mathbf{v}, \mathbf{w}) \in \llbracket \Phi \rrbracket$ if and only if there exists $\pi \in T^*$ such that $\mathbf{u} \xrightarrow{\pi} \mathbb{Q}_+ \mathbf{v}$ and $\Psi(\pi) = \mathbf{w}$.




Fig. 1. Continuous Petri net for $\psi(\mathbf{x}) = (a_1 \cdot x_1 + \dots + a_n \cdot x_n > c) \vee (x_1 > 0 \wedge x_2 > 0) \vee \dots \vee (x_n > 0)$ where $a_1, a_2, c > 0$ and $a_n < 0$.

C. Reducing $\text{CSP}(\Gamma_c)$ to reachability in continuous Petri nets

We now establish the converse direction of the previous section. Given an instance $\Phi(\mathbf{x})$ of $\text{CSP}(\Gamma_c)$, we show that the sets of all non-negative rational solutions \mathbf{x}^* of $\Phi(\mathbf{x})$ can be obtained as subsets of a cross section of reachable markings of a continuous Petri net. The key point is to construct a continuous Petri net gadget for simulating a single convex linear Horn constraint.

Lemma 3.7: Let $\psi(\mathbf{x}) = \mathbf{a}^\top \cdot \mathbf{x} \sim c \vee \bigvee_{1 \leq i \leq m} \theta_i$ be a convex linear Horn constraint. There exist a continuous Petri net \mathcal{N} and $\mathbf{y} \in \{0, 1\}^6$, computable in logarithmic space, such that $\psi(\mathbf{x})$ holds if and only if $(\mathbf{x}, \mathbf{y}) \xrightarrow{*}_{\mathbb{Q}_+} \mathbf{0}$.

Proof: We assume w.l.o.g. that \mathbf{a} and c have integer values. The construction of the continuous Petri net \mathcal{N} for $\psi(\mathbf{x}) = (a_1 \cdot x_1 + \dots + a_n \cdot x_n > c) \vee \theta_1 \vee \dots \vee \theta_m$ is illustrated in Figure 1, where $\theta_1 = (x_1 > 0 \wedge x_2 > 0)$, $\theta_m = (x_n > 0)$, $a_1, a_2, c > 0$ and $a_n < 0$. The places p_1, p_2, \dots, p_n initially contain the values of \mathbf{x} . One of the blue transitions is enabled if and only if θ_i holds for some $i \in [n]$. If this is the case, then firing a blue transition brings a token in q_θ which allows to empty all places using the gray transitions. Thus it remains to consider the case where no blue transition is enabled. Let $\text{val}(\mathcal{N})$ be the “current” difference between the content of q_{pos} and q_{neg} . In order to empty p_1, p_2, \dots, p_n , the black transitions must be fired until $\text{val}(\mathcal{N}) = a_1 \cdot x_1 + \dots + a_n \cdot x_n$. Similarly, q_{\geq} can only be emptied by fully firing the green transition, leading to $\text{val}(\mathcal{N}) = a_1 \cdot x_1 + \dots + a_n \cdot x_n - c$. The leftmost magenta transition allows to decrease $\text{val}(\mathcal{N})$ by any $\lambda \geq 0$. If \sim equals $>$, then a token is initially placed in $g_{>}$, and it can only be removed by decreasing $\text{val}(\mathcal{N})$ by some $\lambda > 0$. At this point, we obtain $\text{val}(\mathcal{N}) = a_1 \cdot x_1 + \dots + a_n \cdot x_n - c - \lambda$. The orange transition “normalizes” the representation of $\text{val}(\mathcal{N})$. Repeatedly firing this transition allows to empty both q_{pos} and q_{neg} if and only if $\text{val}(\mathcal{N}) = 0$. Altogether, if $\text{val}(\mathcal{N}) = 0$, we obtain $a_1 \cdot x_1 + \dots + a_n \cdot x_n = c + \lambda$ for some $\lambda > 0$ as desired. Even though the transitions need not to be fired in this specific way, it can be proven that any firing sequence can be rearranged in this fashion. \square

Since instances of $\text{CSP}(\Gamma_c)$ are sets of convex linear Horn constraints, taking the disjoint union of the gadgets constructed in Lemma 3.7 allows us to conclude that solutions of $\text{CSP}(\Gamma_c)$ are definable by continuous Petri nets.

Theorem 3.8: Let $\Phi(\mathbf{x})$ be an instance of $\text{CSP}(\Gamma_c)$ with free variables \mathbf{x} . There exists a logarithmic-space computable continuous Petri net \mathcal{N} defining the set $\llbracket \Phi \rrbracket$ of solutions of $\Phi(\mathbf{x})$.

Note that Theorem 3.8 implies that the satisfiability problem for $\text{CSP}(\Gamma_c)$ reduces to continuous Petri net reachability. Indeed, the continuous Petri net obtained from Theorem 3.8 can be extended with transitions non deterministically guessing an assignment of the free variables.

D. Applications and observations

We close our investigation of continuous Petri nets with a short account on some observations and P upper bounds that follow as corollaries of our results, and on relationships of continuous Petri nets to other models studied in the literature.

Properties of continuous Petri net reachability sets: It is well-known that reachability sets of continuous Petri nets are convex, see e.g. [6]. It is thus a legitimate question to ask whether it is possible to obtain reachability sets of continuous Petri nets as (projections of) points inside a convex polytope given by a system of strict and non-strict linear inequalities. Combining Theorem 3.6 with Lemma 3.2 shows that this is not the case.

Proposition 3.9: Reachability sets of continuous Petri nets are not definable as (projections of) points inside a convex polytope.

Boundedness, coverability and structural cyclicity: Given a Petri net \mathcal{N} and an initial marking $\mathbf{u} \in \mathbb{Q}_+^P$, *continuous boundedness* is the problem to decide whether there exists $b \in \mathbb{Q}_+$ such that for every marking \mathbf{v} reachable from \mathbf{u} and, for every place p , $v(p) \leq b$ holds. Note that \mathcal{N} is unbounded if and only if there exist $\mathbf{v}, \mathbf{w} \in \mathbb{Q}_+^P$ such that $\mathbf{u} \xrightarrow{*}_{\mathbb{Q}_+} \mathbf{v} \xrightarrow{*}_{\mathbb{Q}_+} \mathbf{w}$ and $\mathbf{v} < \mathbf{w}$. Given markings $\mathbf{u}, \mathbf{v} \in \mathbb{Q}_+$, *continuous coverability* is to decide whether $\mathbf{u} \xrightarrow{*}_{\mathbb{Q}_+} \mathbf{w}$ for some $\mathbf{w} \geq \mathbf{v}$. Finally, *structural cyclicity* is to decide whether $\mathbf{0} \xrightarrow{\pi}_{\mathbb{N}} \mathbf{0}$ for some path π such that $|\pi| > 0$. Observe that \mathcal{N} is structurally cyclic if and only if $\mathbf{0} \xrightarrow{\pi}_{\mathbb{Q}_+} \mathbf{0}$ for some π such that $|\pi| > 0$.

Continuous boundedness and continuous coverability were shown decidable in P in [6] via dedicated algorithms that rely on specific analyses of runs in continuous Petri nets. Structural cyclicity was recently shown P-complete in a dedicated paper

by Drewes and Leroux [21]. Note that the reachability algorithm presented in [6] does not immediately yield membership in P of structural cyclicity.

Corollary 3.10: Continuous boundedness, continuous coverability and structural cyclicity for Petri nets belong to P .

Proof: Rephrase the definitions of the respective decision problems in $\text{CSP}(\Gamma_c)$ and apply Theorem 3.6. \square

Constant-rate multi-mode systems: In [22], Alur, Trivedi and Wojtczak study so-called *constant-rate multi-mode systems (MMS)*, a class of hybrid systems that can switch among a finite set of modes, and whose dynamics is specified by a finite number of real-valued variables with mode-dependent constant rates. Such systems can, for instance, be used in order to model the behavior of heating, ventilation and air conditioning (HVAC) systems, or two tank systems [23]. Formally, a MMS is a tuple $\mathcal{H} = (M, n, R)$, where M is a finite non-empty set of *modes*, n is the number of continuous variables of \mathcal{H} , and $R: M \rightarrow \mathbb{Q}^n$ defines the *rates* of \mathcal{H} that specify for every mode the change in value per time unit in every component. Given an initial configuration $\mathbf{u}_0 \in \mathbb{Q}^n$, a *run* of \mathcal{H} is a possibly infinite sequence $\rho: (\mathbf{u}_0, (m_1, t_1), \mathbf{u}_1, (m_2, t_2), \mathbf{u}_2, \dots)$ such that $m_i \in M$, $t_i \in \mathbb{Q}_+$ and $\mathbf{u}_i = \mathbf{u}_{i-1} + t_i \cdot R(m_i)$ for all $i > 0$. Given a system of linear inequalities $S: A \cdot \mathbf{x} \geq \mathbf{b}$ such that $A \in \mathbb{Q}^{m \times n}$ and $\mathbf{b} \in \mathbb{Q}^m$, ρ is safe with respect to S if $A \cdot \mathbf{u}_i \geq \mathbf{b}$ for all $i \geq 0$. In [22], Alur et al. study safe schedulability and safe reachability in MMS. Without going into any further details, in its essence safe schedulability asks whether there is a finite non-empty run starting and ending in the same \mathbf{u} that is safe with respect to S , and safe reachability asks whether there is a finite run that is safe with respect to S , starts in a designated \mathbf{u} , and ends in some designated $\mathbf{v} \in \mathbb{Q}^n$. Both problems are shown decidable in [22] in polynomial time via sophisticated analyses and dedicated algorithms. However, with no further efforts, the same results can easily be obtained as a corollary of Theorem 3.6 by viewing constant-rate MMS as continuous Petri nets by an application of some insights from Lemma 3.7.

Corollary 3.11: Safe schedulability and safe reachability in constant-rate MMS are logarithmic-space reducible to continuous Petri net reachability problems and thus belong to P .

IV. CONTINUOUS VECTOR ADDITION SYSTEMS WITH STATES

We now introduce CVASS. The difference to their classical discrete counterparts is that in CVASS counters may take non-negative rational values, and the effect of a transition can arbitrarily be scaled by any rational number in $(0, 1]$. Formally, a *continuous d -dimensional vector addition system with states (d -CVASS)* is a pair $\mathcal{V} = (Q, T)$ such that Q is a finite set of *control states*, and $T \subseteq Q \times \mathbb{Z}^d \times Q$ is a finite set of *transitions*. For notational convenience, we write a *configuration* $(q, \mathbf{u}) \in Q \times \mathbb{Q}_+^d$ as $q(\mathbf{u})$, where every component of \mathbf{u} is referred to as a *counter*.

For every transition $t = (p, \mathbf{z}, q)$, the *effect* of t is $\delta(t) \stackrel{\text{def}}{=} \mathbf{z}$. For every $\mathbb{A} \in \{\mathbb{Q}_+, \mathbb{Q}\}$, $\alpha \in (0, 1]$, and $t \in T$, we define the

transition relation $\xrightarrow{\alpha t}_{\mathbb{A}}$ over $Q \times \mathbb{A}^d$ by

$$p(\mathbf{u}) \xrightarrow{\alpha t}_{\mathbb{A}} q(\mathbf{v}) \stackrel{\text{def}}{\iff} \mathbf{u} + \alpha \cdot \delta(t) = \mathbf{v}.$$

For integral \mathbf{u} and \mathbf{v} , we write $p(\mathbf{u}) \xrightarrow{t}_{\mathbb{N}} q(\mathbf{v})$ and $p(\mathbf{u}) \xrightarrow{t}_{\mathbb{Z}} q(\mathbf{v})$ for $p(\mathbf{u}) \xrightarrow{1t}_{\mathbb{Q}_+} q(\mathbf{v})$ and $p(\mathbf{u}) \xrightarrow{1t}_{\mathbb{Q}} q(\mathbf{v})$ in order to denote respectively the VASS and \mathbb{Z} -VASS transition relations. As for continuous Petri nets, we write T^\star for $((0, 1] \times T)^\star$. Transition relations can be extended to sequences, and we denote the *reachability relation* by $\xrightarrow{*}_{\mathbb{A}}$. We say that a sequence $\pi \in T^\star$ is *q -admissible* from $p(\mathbf{u}) \in Q \times \mathbb{Q}_+^d$ if $p(\mathbf{u}) \xrightarrow{\pi}_{\mathbb{Q}_+} q(\mathbf{v})$ for some $\mathbf{v} \in \mathbb{Q}_+^d$. *Reachability in CVASS* is to decide whether $p(\mathbf{u}) \xrightarrow{*}_{\mathbb{Q}_+} q(\mathbf{v})$, and *cyclic reachability* is to decide whether $q(\mathbf{u}) \xrightarrow{*}_{\mathbb{Q}_+} q(\mathbf{v})$. Finally, *\mathbb{Q} -reachability* and *cyclic \mathbb{Q} -reachability* are to decide $p(\mathbf{u}) \xrightarrow{*}_{\mathbb{Q}} q(\mathbf{v})$ and $q(\mathbf{u}) \xrightarrow{*}_{\mathbb{Q}} q(\mathbf{v})$, respectively. Reachability sets of CVASS are defined analogously as for continuous Petri nets.

For every sequence $\pi = \alpha_1 t_1 \alpha_2 t_2 \dots \alpha_n t_n \in T^\star$, the *effect* $\delta(\pi) \in \mathbb{Q}^d$, the (*rational*) *Parikh image* $\Psi(\pi): T \rightarrow \mathbb{Q}_+$, the *support* $\llbracket \pi \rrbracket \subseteq T$ and the *underlying path* of π are defined respectively as

$$\begin{aligned} \delta(\pi) &\stackrel{\text{def}}{=} \alpha_1 \cdot \delta(t_1) + \alpha_2 \cdot \delta(t_2) + \dots + \alpha_n \cdot \delta(t_n), \\ \Psi(\pi)(t) &\stackrel{\text{def}}{=} \sum_{\substack{1 \leq i \leq n \\ t_i = t}} \alpha_i \text{ for every } t \in T, \\ \llbracket \pi \rrbracket &\stackrel{\text{def}}{=} \{t_i : i \in [n]\}, \text{ and} \\ \text{path}(\pi) &\stackrel{\text{def}}{=} t_1 t_2 \dots t_n. \end{aligned}$$

Note that $\delta(\pi) = \sum_{t \in \llbracket \pi \rrbracket} \Psi(t) \cdot \delta(t)$, and if π is q -admissible from $p(\mathbf{u})$, then $p(\mathbf{u}) \xrightarrow{\pi}_{\mathbb{Q}_+} q(\mathbf{u} + \delta(\pi))$. Effects, Parikh images, supports and underlying paths are naturally extended to languages over T^\star .

We define the *transpose* of a transition $t = (p, \mathbf{z}, q) \in T$ as $t^\dagger \stackrel{\text{def}}{=} (q, -\mathbf{z}, p)$, and the *transpose* of a sequence $\pi = \alpha_1 t_1 \alpha_2 t_2 \dots \alpha_n t_n \in T^\star$ as $\pi^\dagger \stackrel{\text{def}}{=} \alpha_n t_n^\dagger \alpha_{n-1} t_{n-1}^\dagger \dots \alpha_1 t_1^\dagger$. Note that $\delta(\pi) = -\delta(\pi^\dagger)$, $\Psi(\pi)(t) = \Psi(\pi^\dagger)(t^\dagger)$, and $t \in \llbracket \pi \rrbracket \iff t^\dagger \in \llbracket \pi^\dagger \rrbracket$. The *transpose* of a d -CVASS $\mathcal{V} = (Q, T)$ is $\mathcal{V}^\dagger \stackrel{\text{def}}{=} (Q, T^\dagger)$ where $T^\dagger \stackrel{\text{def}}{=} \{t^\dagger : t \in T\}$.

The *underlying directed graph* of \mathcal{V} is $G_{\mathcal{V}} \stackrel{\text{def}}{=} (Q, S)$ where $S \stackrel{\text{def}}{=} \{(p, q) : (p, \mathbf{z}, q) \in T\}$. The Parikh image and support of sequences over S^* are defined in the same way as over T^\star . For every $s = (p, q) \in S$, $\text{in}(s) \stackrel{\text{def}}{=} p$ and $\text{out}(s) \stackrel{\text{def}}{=} q$. For every $q \in Q$, $\text{in}(q) \stackrel{\text{def}}{=} \{s \in S : \text{out}(s) = q\}$, $\text{out}(q) \stackrel{\text{def}}{=} \{s \in S : \text{in}(s) = q\}$ and $\text{edges}(q) \stackrel{\text{def}}{=} \text{in}(q) \cup \text{out}(q)$. Let $G = (V, E)$ be a directed graph and $E' \subseteq E$, we define $G[E']$ as the subgraph induced by E' , i.e. $G[E'] \stackrel{\text{def}}{=} (\{v \in V : \text{edges}(v) \cap E' \neq \emptyset\}, E')$.

In order to justify our definition of CVASS, our first result shows that continuous Petri nets can be embedded into CVASS. The standard construction for the discrete case cannot directly be used here since transitions associated to **Pre** and **Post** could be scaled by different factors. We add extra counters in order to rule out such undesired behavior.

Proposition 4.1: Let $\mathcal{N} = (P, T, \mathbf{Pre}, \mathbf{Post})$ be a continuous Petri net. There exist a logarithmic-space computable $(|T|+|P|)$ -CVASS $\mathcal{V} = (Q, S)$ and $q \in Q$ such that

$$\mathbf{u} \xrightarrow{*}_{\mathbb{Q}_+} \mathbf{v} \text{ in } \mathcal{N} \text{ if and only if } q(\mathbf{0}, \mathbf{u}) \xrightarrow{*}_{\mathbb{Q}_+} q(\mathbf{0}, \mathbf{v}) \text{ in } \mathcal{V}.$$

Proof: Let $T = \{t_1, t_2, \dots, t_n\}$. We let $\mathcal{V} \stackrel{\text{def}}{=} (Q, S)$ be the $(|T|+|P|)$ -CVASS such that $Q \stackrel{\text{def}}{=} \{q\} \cup \{p_i : i \in [n]\}$ and $S \stackrel{\text{def}}{=} \{s_i : i \in [n]\} \cup \{t_i : i \in [n]\}$, where $s_i \stackrel{\text{def}}{=} (q, (e_i, -\mathbf{Pre}_i), p_i)$ and $t_i \stackrel{\text{def}}{=} (p_i, (-e_i, \mathbf{Post}_i), q)$ for every $i \in [n]$.

If $\mathbf{u} \xrightarrow{*}_{\mathbb{Q}_+} \mathbf{v}$ in \mathcal{N} , then it is clear that $q(\mathbf{0}, \mathbf{u}) \xrightarrow{*}_{\mathbb{Q}_+} q(\mathbf{0}, \mathbf{v})$ in \mathcal{V} . Let us show the converse. Suppose there exists $\pi \in S^*$ such that $q(\mathbf{0}, \mathbf{u}) \xrightarrow{\pi}_{\mathbb{Q}_+} q(\mathbf{0}, \mathbf{v})$ in \mathcal{V} . First note that π is of the form $\pi = \alpha_1 s_{\ell_1} \beta_1 t_{\ell_1} \dots \alpha_k s_{\ell_k} \beta_k t_{\ell_k}$. If $\alpha_i = \beta_i$, we say that i is *good*, otherwise we say that i is *bad*. In Appendix B, we show that π can be rearranged into an equivalent run with fewer bad indices. Formally, it can be shown that there exists $\pi' \in S^*$ such that $\delta(\pi') = \delta(\pi)$, $\text{path}(\pi') = \text{path}(\pi)$, π' is q -admissible, and π' has fewer bad indices than π . The claim implies that the number of bad indices can be decreased until there are only good indices, in which case we obtain a run of \mathcal{V} corresponding to a run of \mathcal{N} . \square

A. A characterization of cyclic reachability in CVASS

In this section, we establish a characterization of cyclic reachability in CVASS in a similar vein to Proposition 3.4. The latter proposition constitutes the basis of the P upper bound for reachability in continuous Petri nets shown by Fraca and Haddad in [6]. In order to establish an analogous characterization for CVASS, it turns out that we can closely follow and adapt the approach described in [6] to work for CVASS. While the formal technical aspects are non-trivial to establish, the underlying idea is simple to understand. Basically, in order to show $q(\mathbf{u}) \xrightarrow{*}_{\mathbb{Q}_+} q(\mathbf{v})$, we show that it suffices to check whether $q(\mathbf{u}) \xrightarrow{\pi}_{\mathbb{Q}} q(\mathbf{v})$ for some path π , and whether π can be scaled in a way such that it does not force any counter to go below zero when starting in $q(\mathbf{u})$ and when ending in $q(\mathbf{v})$, as illustrated in Figure 2.

For brevity, we only provide all necessary intermediate statements required for our characterization; all proofs can be found in Appendix B-A. For the remainder of this section, we fix a d -CVASS $\mathcal{V} = (Q, T)$. We first establish a couple of facts on runs of CVASS whose proofs are immediate.

Lemma 4.2: The following statements hold for every $\pi \in T^*$ and $\alpha, \alpha' \in (0, 1]$:

- $p(\mathbf{u}) \xrightarrow{\pi}_{\mathbb{Q}_+} q(\mathbf{v})$ if and only if $p(\alpha \cdot \mathbf{u}) \xrightarrow{\alpha \cdot \pi}_{\mathbb{Q}_+} q(\alpha \cdot \mathbf{v})$,
- if $p(\mathbf{u}) \xrightarrow{\alpha \pi}_{\mathbb{Q}_+} q(\mathbf{v})$, $p(\mathbf{u}') \xrightarrow{\alpha' \pi}_{\mathbb{Q}_+} q(\mathbf{v}')$ and $\alpha + \alpha' \leq 1$, then $p(\mathbf{u} + \mathbf{u}') \xrightarrow{(\alpha + \alpha') \pi}_{\mathbb{Q}_+} q(\mathbf{v} + \mathbf{v}')$,
- if $p(\mathbf{u}) \xrightarrow{\pi}_{\mathbb{Q}_+} q(\mathbf{v})$, then $p(\mathbf{u}) \xrightarrow{\alpha \pi}_{\mathbb{Q}_+} q((1 - \alpha) \cdot \mathbf{u} + \alpha \cdot \mathbf{v})$,
- $p(\mathbf{u}) \xrightarrow{\pi}_{\mathbb{Q}_+} q(\mathbf{v})$ in \mathcal{V} if and only if $q(\mathbf{v}) \xrightarrow{\pi^\dagger}_{\mathbb{Q}_+} p(\mathbf{u})$ in \mathcal{V}^\dagger .

Next, we establish a sufficient condition to apply the transformation illustrated in Figure 2, i.e. to convert a cyclic run over \mathbb{Q} into a cyclic run over \mathbb{Q}_+ with the same Parikh image.




Fig. 2. Two runs of $\mathcal{V} = (Q, T)$, where $Q = \{p, q\}$ and $T = \{s : (q, (-2, 4), p), t : (p, (4, -2), q)\}$. State q and transition s appear in blue, and p and t appear in magenta. Left: $q(1, 1) \xrightarrow{s^t}_{\mathbb{Q}} q(3, 3)$; right: an equivalent run $q(1, 1) \xrightarrow{((1/4)s(1/4)t)^4}_{\mathbb{Q}_+} q(3, 3)$ that remains above $\mathbf{0}$.

Lemma 4.3: Let $q \in Q$, $\mathbf{u}, \mathbf{v} \in \mathbb{Q}_+^d$, and $\pi \in T^*$ be such that

- $q(\mathbf{u}) \xrightarrow{\pi}_{\mathbb{Q}} q(\mathbf{v})$,
- $\llbracket \delta(t) \rrbracket^- \subseteq \llbracket \mathbf{u} \rrbracket$ for every $t \in \llbracket \pi \rrbracket$, and
- $\llbracket \delta(t) \rrbracket^+ \subseteq \llbracket \mathbf{v} \rrbracket$ for every $t \in \llbracket \pi \rrbracket$.

There exists $\pi' \in T^*$ such that $q(\mathbf{u}) \xrightarrow{\pi'}_{\mathbb{Q}_+} q(\mathbf{v})$ and $\Psi(\pi') = \Psi(\pi)$.

The subsequent lemma now shows that if we assume the existence of a path π witnessing \mathbb{Q}_+ -reachability then we can extract a path π' from π that ends in a configuration that fulfills some of the criteria of Lemma 4.3.

Lemma 4.4: Let $p, q \in Q$, $\mathbf{u}, \mathbf{v} \in \mathbb{Q}_+^d$ and $\pi \in T^*$. If $p(\mathbf{u}) \xrightarrow{\pi}_{\mathbb{Q}_+} q(\mathbf{v})$, then there exists $\pi' \in T^*$ such that

- π' is q -admissible from $p(\mathbf{u})$,
- $\text{path}(\pi') = \text{path}(\pi)$, and
- $\llbracket \delta(t) \rrbracket \subseteq \llbracket \mathbf{v}' \rrbracket$ for every $t \in \llbracket \pi' \rrbracket$, where $\mathbf{v}' \stackrel{\text{def}}{=} \mathbf{u} + \delta(\pi')$.

We have now established all intermediate results that we require in order to obtain our characterization of CVASS cyclic reachability:

Proposition 4.5: Let $q \in Q$, $\mathbf{u}, \mathbf{v} \in \mathbb{Q}_+^d$ and $\mathbf{w} \in \mathbb{Q}_+^T$. There exists $\pi \in T^*$ such that $q(\mathbf{u}) \xrightarrow{\pi}_{\mathbb{Q}_+} q(\mathbf{v})$ and $\Psi(\pi) = \mathbf{w}$ if and only if there exist $\pi, \pi_{\text{fwd}}, \pi_{\text{bwd}} \in T^*$ such that

- $q(\mathbf{u}) \xrightarrow{\pi}_{\mathbb{Q}} q(\mathbf{v})$,
- π_{fwd} is q -admissible from $q(\mathbf{u})$ in \mathcal{V} ,
- π_{bwd}^\dagger is q -admissible from $q(\mathbf{v})$ in \mathcal{V}^\dagger ,
- $\Psi(\pi) = \mathbf{w}$ and $\llbracket \pi \rrbracket = \llbracket \pi_{\text{fwd}} \rrbracket = \llbracket \pi_{\text{bwd}} \rrbracket = \llbracket \mathbf{w} \rrbracket$.

Proof: The “only if” direction is immediate. Therefore, we assume (a–d) hold, and we show that $q(\mathbf{u}) \xrightarrow{*}_{\mathbb{Q}_+} q(\mathbf{v})$.

By Lemma 4.4, there exists $\pi'_{\text{fwd}} \in T^*$ such that

- π'_{fwd} is q -admissible from $q(\mathbf{u})$ in \mathcal{V} ,
- $\text{path}(\pi'_{\text{fwd}}) = \text{path}(\pi_{\text{fwd}})$, and
- $\llbracket \delta(t) \rrbracket \subseteq \llbracket \mathbf{x} \rrbracket$ for every $t \in \llbracket \pi'_{\text{fwd}} \rrbracket$, where $\mathbf{x} \stackrel{\text{def}}{=} \mathbf{u} + \delta(\pi'_{\text{fwd}})$.

In particular, note that $q(\mathbf{u}) \xrightarrow{\pi'_{\text{fwd}}}_{\mathbb{Q}_+} q(\mathbf{x})$.

By Lemma 4.4, there exists $\pi'_{\text{bwd}} \in T^*$ such that

- $\pi'_{\text{bwd}}^\dagger$ is q -admissible from $q(\mathbf{v})$ in \mathcal{V}^\dagger ,
- $\text{path}(\pi'_{\text{bwd}}^\dagger) = \text{path}(\pi_{\text{bwd}}^\dagger)$, and
- $\llbracket \delta(t) \rrbracket \subseteq \llbracket \mathbf{y} \rrbracket$ for every $t \in \llbracket \pi'_{\text{bwd}}^\dagger \rrbracket$, where $\mathbf{y} \stackrel{\text{def}}{=} \mathbf{v} + \delta(\pi'_{\text{bwd}}^\dagger)$.

Note that $q(\mathbf{v}) \xrightarrow{\pi'_{\text{bwd}}}_{\mathbb{Q}_+} q(\mathbf{y})$, which implies, together with Lemma 4.2(d), that $q(\mathbf{y}) \xrightarrow{\pi'_{\text{bwd}}}_{\mathbb{Q}_+} q(\mathbf{v})$.

Let $\pi = \alpha_1 t_1 \alpha_2 t_2 \cdots \alpha_k t_k$. For every $j \in [k]$, let

$$\begin{aligned} a_j &\stackrel{\text{def}}{=} \Psi(\pi'_{\text{fwd}})(t_j) / \Psi(\text{path}(\pi))(t_j) \\ b_j &\stackrel{\text{def}}{=} \Psi(\pi'_{\text{bwd}})(t_j) / \Psi(\text{path}(\pi))(t_j). \end{aligned}$$

Let $0 < \lambda \leq 1$ be small enough so that $\alpha_j - \lambda(a_j + b_j) \geq 0$ for every $j \in [k]$. Let $c_j \stackrel{\text{def}}{=} \alpha_j - \lambda(a_j + b_j)$ for every $j \in [k]$, and let $\sigma \stackrel{\text{def}}{=} c_1 t_1 c_2 t_2 \cdots c_k t_k$. By Lemma 4.2(c), we have

$$q(\mathbf{u}) \xrightarrow{\lambda \pi'_{\text{fwd}}}_{\mathbb{Q}_+} \overbrace{q(\mathbf{u} + \lambda \delta(\pi'_{\text{fwd}}))}^{\mathbf{x}'}$$

Similarly, by applying Lemma 4.2(d), Lemma 4.2(c) and again Lemma 4.2(d), we obtain

$$q(\underbrace{\mathbf{v} - \lambda \delta(\pi'_{\text{bwd}})}_{\mathbf{y}'}) \xrightarrow{\lambda \pi'_{\text{bwd}}}_{\mathbb{Q}_+} q(\mathbf{v}).$$

Moreover,

$$\begin{aligned} \delta(\sigma) &= \sum_{j \in [k]} c_j \cdot \delta(t_j) \\ &= \sum_{j \in [k]} (\alpha_j - \lambda(a_j + b_j)) \cdot \delta(t_j) \\ &= \sum_{j \in [k]} \alpha_j \cdot \delta(t_j) - \lambda \sum_{j \in [k]} (a_j + b_j) \cdot \delta(t_j) \\ &= \delta(\pi) - \lambda \sum_{j \in [k]} (a_j + b_j) \cdot \delta(t_j) \\ &= \delta(\pi) - \lambda \sum_{j \in [k]} \frac{\Psi(\pi'_{\text{fwd}})(t_j) + \Psi(\pi'_{\text{bwd}})(t_j)}{\Psi(\text{path}(\pi))(t_j)} \cdot \delta(t_j) \\ &= \delta(\pi) - \lambda \sum_{t \in [\pi]} (\Psi(\pi'_{\text{fwd}})(t) + \Psi(\pi'_{\text{bwd}})(t)) \cdot \delta(t) \\ &= \delta(\pi) - \lambda \delta(\pi'_{\text{fwd}}) - \lambda \delta(\pi'_{\text{bwd}}) \\ &= \mathbf{v} - \mathbf{u} - \lambda \delta(\pi'_{\text{fwd}}) - \lambda \delta(\pi'_{\text{bwd}}) \\ &= (\mathbf{v} - \lambda \delta(\pi'_{\text{bwd}})) - (\mathbf{u} + \lambda \delta(\pi'_{\text{fwd}})) \\ &= \mathbf{y}' - \mathbf{x}'. \end{aligned}$$

Therefore, $q(\mathbf{x}') \xrightarrow{\sigma}_{\mathbb{Q}} q(\mathbf{y}')$. By definition of \mathbf{x}' and \mathbf{y}' , and since $0 < \lambda \leq 1$, we have $[\mathbf{x}] \subseteq [\mathbf{x}']$ and $[\mathbf{y}] \subseteq [\mathbf{y}']$. Thus,

$$\begin{aligned} [\delta(t)] &\subseteq [\mathbf{x}'] \text{ for every } t \in [\pi'_{\text{fwd}}], \\ [\delta(t)] &\subseteq [\mathbf{y}'] \text{ for every } t \in [\pi'_{\text{bwd}}], \text{ and} \end{aligned}$$

hence σ satisfies all of the conditions of Lemma 4.3. Thus, there exists $\sigma' \in T^\star$ such that $\Psi(\sigma') = \Psi(\sigma)$ and $q(\mathbf{x}') \xrightarrow{\sigma'}_{\mathbb{Q}_+} q(\mathbf{y}')$. We obtain

$$q(\mathbf{u}) \xrightarrow{\lambda \pi'_{\text{fwd}}}_{\mathbb{Q}_+} q(\mathbf{x}') \xrightarrow{\sigma'}_{\mathbb{Q}_+} q(\mathbf{y}') \xrightarrow{\lambda \pi'_{\text{bwd}}}_{\mathbb{Q}_+} q(\mathbf{v}).$$

It remains to observe that $\Psi(\lambda \pi'_{\text{fwd}} \cdot \sigma' \cdot \lambda \pi'_{\text{bwd}}) = \Psi(\pi)$. \square

B. Reducing reachability in CVASS to existential FO($\mathbb{Q}, +, <$)

We now show that the characterization of CVASS cyclic reachability established in Proposition 4.5 enables us to define reachability relations of CVASS in existential FO($\mathbb{Q}, +, <$). This in turn shows that reachability in CVASS can be decided in NP. For the remainder of this section, let $\mathcal{V} = (Q, T)$ be an arbitrary but fixed d -CVASS.

We first make the conditions of Proposition 4.5 effective and then develop an existential FO($\mathbb{Q}, +, <$) characterization of CVASS cyclic reachability. Our first step is to provide an effective characterization of Proposition 4.5(a). Essentially, Lemma 4.6 below is an adaptation to CVASS of the classical Euler-Hierholzer theorem characterizing directed graphs containing Eulerian cycles.

Lemma 4.6: Let $q \in Q$, $\mathbf{u}, \mathbf{v} \in \mathbb{Q}^d$ and $\mathbf{w} \in \mathbb{Q}_+^T$. There exists $\pi \in T^\star$ such that $q(\mathbf{u}) \xrightarrow{\pi}_{\mathbb{Q}} q(\mathbf{v})$ and $\Psi(\pi) = \mathbf{w}$ if and only if there exist $\mathbf{g} \in \mathbb{Q}_+^T$ such that

- (a) $\mathbf{v} = \mathbf{u} + \sum_{t \in T} \mathbf{w}(t) \cdot \delta(t)$,
- (b) $[\mathbf{g}] = [\mathbf{w}]$ and $\mathbf{g}(t) \geq \mathbf{w}(t)$ for every $t \in S$,
- (c) $G_{\mathcal{V}}[[\mathbf{g}]]$ is strongly connected and contains q , and
- (d) $\sum_{t \in \text{in}(p)} \mathbf{g}(t) = \sum_{t \in \text{out}(p)} \mathbf{g}(t)$ for every $p \in Q$.

Proof: \Rightarrow) Let $\pi = \alpha_1 t_1 \alpha_2 t_2 \cdots \alpha_k t_k$ be such that $q(\mathbf{u}) \xrightarrow{\pi}_{\mathbb{Q}} q(\mathbf{v})$ and $\Psi(\pi) = \mathbf{w}$. Let $\mathbf{g}(t) \stackrel{\text{def}}{=} \Psi(\text{path}(\pi))(t)$ for every $t \in T$. Note that (a) immediately holds, and that (b) holds since $0 < \alpha_i \leq 1$ for every $i \in [k]$. Let $S \stackrel{\text{def}}{=} [\mathbf{g}]$. Since $\text{path}(\pi)$ is a path from q to q , $G_{\mathcal{V}}[S]$ is strongly connected and contains q . It remains to show (d). Let H be the directed multigraph obtained from $G_{\mathcal{V}}[S]$ by replacing each edge $t \in S$ by $\Psi(\text{path}(\pi))(t)$ many arcs from $\text{in}(t)$ to $\text{out}(t)$. It is readily seen that $\text{path}(\pi)$ yields an Eulerian cycle in H . Therefore, $|\text{in}(H, p)| = |\text{out}(H, p)|$ for every $p \in Q$, which implies (d).

\Leftarrow) Let $S \stackrel{\text{def}}{=} [\mathbf{g}]$. For every $t \in S$, let $a_t \in \mathbb{N}$ and $b_t \in \mathbb{N}_{>0}$ be such that $\mathbf{g}(t) = a_t / b_t$. Let $\lambda \stackrel{\text{def}}{=} \text{lcm}\{b_t : t \in S\}$, and let H be the directed multigraph obtained from $G_{\mathcal{V}}[S]$ by replacing each arc $t \in S$ by $\lambda \cdot \mathbf{g}(t)$ arcs from $\text{in}(t)$ to $\text{out}(t)$. By (c), $G_{\mathcal{V}}[S]$ is strongly connected and contains q . Thus, H is strongly connected and contains q . Moreover, for every $p \in Q$,

$$\begin{aligned} \text{in}(H, p) &= \sum_{t \in \text{in}(G_{\mathcal{V}}, p)} \lambda \cdot \mathbf{g}(t) \\ &= \lambda \cdot \sum_{t \in \text{in}(G_{\mathcal{V}}, p)} \mathbf{g}(t) \\ &= \lambda \cdot \sum_{t \in \text{out}(G_{\mathcal{V}}, p)} \mathbf{g}(t) \quad (\text{by (d)}) \\ &= \sum_{t \in \text{out}(G_{\mathcal{V}}, p)} \lambda \cdot \mathbf{g}(t) = \text{out}(H, p). \end{aligned}$$

Therefore, H has an Eulerian cycle σ from q to q . Let $\sigma' = t_1 t_2 \cdots t_k$ be the path of $G_{\mathcal{V}}$ obtained from σ . Let $i \in [k]$, and $\alpha_i \stackrel{\text{def}}{=} \mathbf{w}(t_i) / \Psi(\sigma')(t_i)$. Note that $\Psi(\sigma')(t_i) = \lambda \cdot \mathbf{g}(t_i)$.

Moreover, $0 < \mathbf{w}(t_i) \leq \mathbf{g}(t_i)$ by (b). Thus, $0 < \alpha_i \leq 1$. Let $\pi \stackrel{\text{def}}{=} \alpha_1 t_1 \alpha_2 t_2 \cdots \alpha_k t_k$. We have

$$\begin{aligned}
\delta(\pi) &= \sum_{i \in [k]} \alpha_i \cdot \delta(t_i) \\
&= \sum_{i \in [k]} (\mathbf{w}(t_i) / \Psi(\sigma')(t_i)) \cdot \delta(t_i) \\
&= \sum_{t \in \llbracket \pi \rrbracket} \Psi(\sigma')(t) \cdot (\mathbf{w}(t) / \Psi(\sigma')(t)) \cdot \delta(t) \\
&= \sum_{t \in \llbracket \pi \rrbracket} \mathbf{w}(t) \cdot \delta(t) \\
&= \mathbf{v} - \mathbf{u} \tag{by (a)}.
\end{aligned}$$

Therefore, we obtain $q(\mathbf{u}) \xrightarrow{\pi} q(\mathbf{v})$. Moreover, a simple calculation shows that $\Psi(\pi) = \mathbf{w}$. \square

Note that if CVASS over \mathbb{R} are considered instead, \mathbf{g} can still be chosen to be rational in the previous proof, and hence taking the least common multiple makes sense. See Appendix B-B for more details.

We now turn towards an effective characterization of Proposition 4.5(b) and (c). Basically, the next lemma states that there exists a q -admissible path π starting in $p(\mathbf{u})$ whenever we can find an ordering of the transitions visited by π such that all counters which are decremented by a transition have previously been incremented.

Lemma 4.7: Let $p(\mathbf{u}) \in Q \times \mathbb{Q}_+^d$, $q \in Q$ and $S \subseteq T$. There exists $\pi \in T^*$ such that

- π is q -admissible from $p(\mathbf{u})$, and
- $\llbracket \pi \rrbracket = S$.

if and only if there exists an injection $f: S \rightarrow \mathbb{Q}_{>0}$ such that for every $t, t' \in S$,

- (a) if $f(t) = \min\{f(s) : s \in S\}$, then $\text{in}(t) = p$,
- (b) if $f(t) < f(t')$, then there exists a path π in $G_{\mathcal{V}}$ from $\text{out}(t)$ to $\text{in}(t')$ such that $\llbracket \pi \rrbracket \subseteq \{s \in S : f(s) < f(t')\}$,
- (c) if $f(t) = \max\{f(s) : s \in S\}$, then there exists a path π in $G_{\mathcal{V}}$ from $\text{out}(t)$ to q such that $\llbracket \pi \rrbracket \subseteq S$, and
- (d) $\llbracket \delta(t) \rrbracket^- \subseteq \llbracket \mathbf{u} \rrbracket \cup \bigcup_{\substack{s \in S \\ f(s) < f(t)}} \llbracket \delta(s) \rrbracket^+$.

It is now not so difficult to see that we can translate CVASS cyclic reachability into $\text{FO}(\mathbb{Q}, +, <)$ via the effective characterizations of the conditions of Proposition 4.5 provided by Lemma 4.6 and Lemma 4.7:

Lemma 4.8: There exists a logarithmic-space computable existential $\text{FO}(\mathbb{Q}, +, <)$ formula $\varphi_{\mathcal{V}, q}(\mathbf{x}, \mathbf{y}, \mathbf{z})$ such that $\varphi_{\mathcal{V}, q}(\mathbf{u}, \mathbf{v}, \mathbf{w})$ holds if and only if there exists $\pi \in T^*$ such that $q(\mathbf{u}) \xrightarrow{\pi} q(\mathbf{v})$ and $\Psi(\pi) = \mathbf{w}$.

We may now define CVASS reachability in existential $\text{FO}(\mathbb{Q}, +, <)$. Given $p(\mathbf{u}), q(\mathbf{v}) \in Q \times \mathbb{Q}_+^d$, we observe that any run from $p(\mathbf{u})$ to $q(\mathbf{v})$ can be decomposed as a run alternating between simple paths and cycles of arbitrarily length. More formally, $p(\mathbf{u}) \xrightarrow{*} q(\mathbf{v})$ if and only if there exist $0 \leq k \leq |Q|$, $q_0, q_1, \dots, q_{k+1} \in Q$, $\mathbf{u}_0, \mathbf{u}_1, \dots, \mathbf{u}_k, \mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_{k+1} \in \mathbb{Q}_+^d$ and $\pi_0, \pi_1, \dots, \pi_k \in T^*$ such that

- (a) $p(\mathbf{u}) = q_0(\mathbf{u}_0)$, $q(\mathbf{v}) = q_{k+1}(\mathbf{v}_{k+1})$,

(b) $q_i(\mathbf{u}_i) \xrightarrow{\pi_i} q_{i+1}(\mathbf{v}_{i+1})$ for every $0 \leq i \leq k$,

(c) $q_i(\mathbf{v}_i) \xrightarrow{*} q_i(\mathbf{u}_i)$ for every $0 < i \leq k$,

(d) $0 \leq |\pi_i| < |Q|$ for every $0 \leq i \leq k$.

Since the intermediate paths π_i in (b) are of bounded length, they can be guessed in $\text{FO}(\mathbb{Q}, +, <)$, and Condition (c) can be encoded using the formulas established in Lemma 4.8. Those observations enable us to prove the main theorem of this section.

Theorem 4.9: Let $p, q \in Q$. There exists a logarithmic-space computable existential $\text{FO}(\mathbb{Q}, +, <)$ formula $\varphi_{\mathcal{V}, p, q}(\mathbf{x}, \mathbf{y}, \mathbf{z})$ such that $\varphi_{\mathcal{V}, p, q}(\mathbf{u}, \mathbf{v}, \mathbf{w})$ holds if and only if there exists $\pi \in T^*$ such that $p(\mathbf{u}) \xrightarrow{\pi} q(\mathbf{v})$ and $\Psi(\pi) = \mathbf{w}$.

Since satisfiability in existential $\text{FO}(\mathbb{Q}, +, <)$ is NP-complete [13], we obtain an upper bound for reachability in CVASS as an immediate consequence.

Corollary 4.10: Reachability in CVASS is in NP.

C. Reducing existential $\text{FO}(\mathbb{Q}, +, <)$ to reachability in CVASS

In this section, we briefly discuss that the sets of rational numbers definable in existential $\text{FO}(\mathbb{Q}, +, <)$ are, in a certain sense, equivalent to reachability sets of CVASS. Since counters of CVASS cannot hold negative numbers, we encode arbitrary rational numbers as pairs of non-negative rational numbers. Given $u \in \mathbb{Q}$, we define $\text{enc}(u) \stackrel{\text{def}}{=} (u, 0)$ if $u \geq 0$, and $\text{enc}(u) \stackrel{\text{def}}{=} (0, |u|)$ if $u < 0$. The value represented by $(u, v) \in \mathbb{Q}_+^2$ is $\text{val}(u, v) \stackrel{\text{def}}{=} u - v$. We naturally extend this notation to vectors. For every $\mathbf{u} \in \mathbb{Q}^d$, $\text{enc}(\mathbf{u}) \in \mathbb{Q}_+^{2d}$ is the vector whose $(2 \cdot i - 1)$ -th and $(2 \cdot i)$ -th components correspond to $\text{enc}(u(i))$. The value of $\mathbf{u} \in \mathbb{Q}_+^{2d}$ is the vector $\text{val}(\mathbf{u}) \in \mathbb{Q}^d$ such that $\text{val}(\mathbf{u})(i) \stackrel{\text{def}}{=} \mathbf{u}(2 \cdot i - 1) - \mathbf{u}(2 \cdot i)$.

Lemma 4.11: Let $\psi(\mathbf{x})$ be an atomic formula of $\text{FO}(\mathbb{Q}, +, <)$ and $n = |\mathbf{x}|$. There is a logarithmic-space computable $(2n+3)$ -CVASS $\mathcal{V} = (Q, T)$ and states $p, r \in Q$ such that $\psi(\mathbf{u})$ holds if and only if $p(\text{enc}(\mathbf{u}), 0, 0, 1) \xrightarrow{*} r(\mathbf{0})$.

Proof: The proof consists of a more or less straightforward adaptation of Lemma 3.7. \square

Using Lemma 4.11 as a gadget, and by simulating conjunctions by sequential composition and disjunction by branching, we can lift Lemma 4.11 to arbitrary existential $\text{FO}(\mathbb{Q}, +, <)$ formulas. The construction is similar to the one for \mathbb{Z} -VASS established in [3, Lem. 12].

Theorem 4.12: Let $\varphi(\mathbf{x})$ be a formula of existential $\text{FO}(\mathbb{Q}, +, <)$ and $n = |\mathbf{x}|$. There exist $k > 0$, an $(n+k)$ -CVASS $\mathcal{V} = (Q, T)$, $p, q \in Q$ and $\mathbf{v} \in \{0, 1\}^k$, computable in logarithmic space, such that $\varphi(\mathbf{u})$ holds if and only if $p(\text{enc}(\mathbf{u}), \mathbf{v}) \xrightarrow{*} q(\mathbf{0})$.

D. Complexity of decision problems for CVASS

Having already established an NP upper bound in Section IV-B, here we provide strong complementary lower bounds for reachability in CVASS, and briefly discuss the complexity of related decision problems. The next lemma essentially shows that allowing for any sort of non-trivial control structure renders reachability in CVASS NP-hard.

Lemma 4.13: CVASS reachability is already NP-hard for the following cases:

- (a) CVASS in which the number of counters is two and the underlying control graph is acyclic, and
- (b) cyclic reachability in CVASS whose constants are encoded in unary.

Proof: In order to show (a), we reduce from an instance $A \subseteq \{a_1, a_2, \dots, a_n\} \subseteq \mathbb{N}$, $t \in \mathbb{N}$ of the classical NP-complete subset sum problem. The reduction is similar to the reduction showing NP-hardness of reachability in two-clock timed automata [24]. We define $\mathcal{V} = (Q, T)$ such that $Q = \{p_0, p_1, \dots, p_n\}$ and

$$T \stackrel{\text{def}}{=} \{(p_{i-1}, (a_i, 1), p_i), (p_{i-1}, (0, 1), p_i) : i \in [n]\}.$$

Clearly, $p_0(0, 0) \xrightarrow{*}_{\mathbb{Q}_+} p_n(t, n)$ if and only if there exists $A' \subseteq A$ such that $\sum_{a \in A'} a = t$.

Regarding (b), we reduce from an instance $\varphi(x_1, \dots, x_n)$ of Boolean 3-SAT. We construct a $(2 \cdot n)$ -CVASS \mathcal{V} from φ . Intuitively, counter $2 \cdot i - 1$ (respectively $2 \cdot i$) contains a non-zero value if and only if x_i (respectively $\neg x_i$) is set to true. Let C_i be the set of the three literals occurring in clause i of φ . First, \mathcal{V} guesses an assignment of the variables by increasing the counter associated to x_i or $\neg x_i$. Then, \mathcal{V} ensures that each clause is satisfied by subtracting a small amount from the counter associated to some literal $\ell \in C_i$. Finally \mathcal{V} allows to decrease all counters and to go back to the starting state. Formal details can be found in Appendix B-D. \square

Together with Corollary 4.10, we obtain one of the main theorems of this paper.

Theorem 4.14: Reachability in CVASS is NP-complete.

Since reachability in continuous Petri nets is decidable in polynomial time and reduces to CVASS cyclic reachability (see Proposition 4.1), one may have hoped CVASS cyclic reachability to be decidable in polynomial time as well. Lemma 4.13(b) shows that this is not the case. The only P upper bound that we can establish in the presence of control states is for cyclic \mathbb{Q} -reachability.

Theorem 4.15: Cyclic \mathbb{Q} -reachability is in P.

Proof: We show that the conditions of Lemma 4.6 can be expressed in $\text{CSP}(\Gamma_c)$. This is clear for all conditions except for (c), which requires that a directed graph is strongly connected and contains a certain vertex. Thus, let $G = (V, E)$ be a directed graph, $u \in V$, and let $\mathbf{x} = (x_e)_{e \in E}$ and $\mathbf{r} = (r_{u,v})_{u,v \in V}$. Our first step is to define $\Phi_p(\mathbf{x}, \mathbf{r})$ such that $\llbracket \Phi_p \rrbracket = \{(\mathbf{x}, \mathbf{r}) : r_{u,v} > 0 \text{ if and only if } G[\llbracket \mathbf{x} \rrbracket] \text{ contains a path from } u \text{ to } v\}$. To this end, we introduce additional variables $\mathbf{y} = (y_{u,v,i})_{u,v \in V, i \in [0, |E|]}$ capturing reachability relations after i steps and define:

$$\begin{aligned} \Phi_p \stackrel{\text{def}}{=} & \bigwedge_{u \in V} (y_{u,u,0} > 0 \wedge \bigwedge_{v \in V \setminus \{u\}} y_{u,v,0} = 0) \wedge \\ & \bigwedge_{\substack{u,v \in V \\ 1 \leq i \leq |E|}} (y_{u,v,i} > 0 \rightarrow \bigvee_{e \in \text{in}(v)} (x_e > 0 \wedge y_{u, \text{in}(e), i-1} > 0)) \wedge \\ & \bigwedge_{u,v \in V} r_{u,v} > 0 \leftrightarrow \bigvee_{0 \leq i \leq |E|} y_{u,v,i} > 0. \end{aligned}$$

Using Φ_p , we can now define $\Phi_{sc}(\mathbf{x})$ such that $\mathbf{x} \in \llbracket \Phi_{sc} \rrbracket$ if and only if $G[\llbracket \mathbf{x} \rrbracket]$ is strongly connected and contains u . Additionally, Φ_{sc} contains variables $\mathbf{g} = (g_v)_{v \in V}$ such that $g_v > 0$ indicates that v is contained in $G[\llbracket \mathbf{x} \rrbracket]$:

$$\begin{aligned} \Phi_{sc} \stackrel{\text{def}}{=} & \Phi_p \wedge \left(\bigwedge_{v \in V} g_v > 0 \leftrightarrow \bigvee_{e \in \text{edges}(v)} x_e > 0 \right) \wedge g_u > 0 \wedge \\ & \bigwedge_{v \in V} g_v > 0 \rightarrow (r_{u,v} > 0 \wedge r_{v,u} > 0). \end{aligned}$$

It is not difficult to verify that Φ_{sc} has the desired properties. A full construction defining all conditions of Lemma 4.6 can be found in Appendix B-B. \square

Moreover, as an easy consequence of the previously established lower bounds and the existential $\text{FO}(\mathbb{Q}, +, <)$ characterization of reachability in CVASS, we obtain NP-completeness results for other decision problems for CVASS. Given $p(\mathbf{u}) \in \mathbb{Q}_+^d$, *boundness* asks whether there exists \mathbf{w} such that $p(\mathbf{u}) \xrightarrow{*}_{\mathbb{Q}_+} q(\mathbf{v})$ implies $\mathbf{v} \leq \mathbf{w}$. Given an additional configuration $q(\mathbf{v}) \in \mathbb{Q}_+^d$, *coverability* asks whether $p(\mathbf{v}) \xrightarrow{*}_{\mathbb{Q}_+} q(\mathbf{w})$ for some $\mathbf{w} \geq \mathbf{v}$.

Theorem 4.16: Coverability, q -admissibility and unboundness in CVASS are all NP-complete.

Proof: For the upper bounds, we respectively translate the decision problems into existential $\text{FO}(\mathbb{Q}, +, <)$ as follows:

- $\exists \mathbf{w} \in \mathbb{Q}_+^d \exists \mathbf{p} \in \mathbb{Q}_+^T : \varphi_{\mathcal{V}, p, q}(\mathbf{u}, \mathbf{w}, \mathbf{p}) \wedge \mathbf{w} \geq \mathbf{v}$,
- $\exists \mathbf{v} \in \mathbb{Q}_+^d \exists \mathbf{p} \in \mathbb{Q}_+^T : \varphi_{\mathcal{V}, p, q}(\mathbf{u}, \mathbf{v}, \mathbf{p})$,
- $\exists \mathbf{v}, \mathbf{w} \in \mathbb{Q}_+^d \exists \mathbf{p}, \mathbf{p}' \in \mathbb{Q}_+^T : \bigvee_{q \in Q} \varphi_{\mathcal{V}, p, q}(\mathbf{u}, \mathbf{v}, \mathbf{p}) \wedge \varphi_{\mathcal{V}, q}(\mathbf{v}, \mathbf{w}, \mathbf{p}') \wedge \mathbf{w} > \mathbf{v}$.

For the NP-hardness, it is possible to adapt the proof of Lemma 4.13(b). \square

In contrast to the good complexity-theoretic properties of CVASS for the aforementioned decision problems, as in classical VASS, the introduction of zero-test renders reachability in CVASS undecidable.

Theorem 4.17: Reachability in CVASS with zero-tests is undecidable.

Proof: Recall that reachability for 2-VASS with zero-tests is undecidable [25]. We sketch how to simulate a VASS \mathcal{V} with zero-tests by a CVASS \mathcal{V}' with zero-tests. First, two new counters are added. Each non zero-test transition of \mathcal{V} is converted into the CVASS gadget of Figure 3 to ensure that each transition is “fully taken”. We add a transition from q




Fig. 3. Gadget to transform a VASS into a CVASS with zero-tests with the same reachability set.

to a “trap state” r which allows to decrement the two new counters. It readily seen that $p(\mathbf{u}) \xrightarrow{*}_{\mathbb{Q}_+} q(\mathbf{v})$ in \mathcal{V} if and only if $p(1, 0, \mathbf{u}) \xrightarrow{*}_{\mathbb{Q}_+} r(0, 0, \mathbf{v})$ in \mathcal{V}' . \square





Fig. 4. Example of a 2-VASS (left) and its equivalent Petri net (right).

V. OVER-APPROXIMATING VASS REACHABILITY SETS BY CVASS

As discussed in the introduction, continuous Petri nets, \mathbb{Z} -VASS and variants thereof have been used in the literature in order to over-approximate reachability sets of discrete Petri nets and VASS [11], [7], [12], [9]. In general, the finer an over-approximation, the more useful it is for pruning search spaces in discrete Petri nets and VASS. The goal of this section is to briefly discuss the over-approximation of VASS reachability sets achieved by CVASS in comparison with continuous Petri nets and \mathbb{Z} -VASS.

As we showed in Proposition 4.1, continuous Petri nets embed into CVASS, and hence CVASS subsume reachability sets of continuous Petri nets. Since CVASS define all existential $\text{FO}(\mathbb{Q}, +, <)$ -definable sets and continuous Petri nets $\text{CSP}(\Gamma_c)$ -definable sets, it is clear that there exist CVASS whose reachability sets cannot be defined by continuous Petri nets. Figure 4 provides an example, which additionally illustrates the benefits of over-approximating discrete VASS with CVASS instead of continuous Petri nets. Recall from Section IV-D that $q(\mathbf{v})$ is coverable from $p(\mathbf{u})$ if $p(\mathbf{u}) \xrightarrow{*} \mathbb{Q}_+ q(\mathbf{w})$ for some $\mathbf{w} \geq \mathbf{v}$. It is readily seen that $p(1, 1)$ cannot be covered from $p(0, 0)$ in the VASS \mathcal{V} on the left-hand side of Figure 4, under both the discrete and continuous semantics. However, if we turn \mathcal{V} into a Petri net \mathcal{N} such as the one shown on the right-hand side of Figure 4, the control states of \mathcal{V} become places, and hence interpreting \mathcal{N} under continuous semantics allows for, informally speaking, being in fractions of the control states of the original VASS \mathcal{V} .³ Consequently, the following firing sequence is possible in \mathcal{N} :

$$\{p\} \xrightarrow{\pi = \frac{1}{2} \cdot t_1 \frac{1}{2} \cdot t_5 \frac{1}{2} \cdot t_2 \frac{1}{2} \cdot t_4 t_3 t_6} \mathbb{Q}_+ \{p, 2 \cdot c_1, 2 \cdot c_2\}.$$

Hence, $\{p, c_1, c_2\}$ is coverable in \mathcal{N} . It is worth noting that the Parikh image of π above has non-integer values, i.e., $\Psi(\pi) = (1/2, 1/2, 1, 1/2, 1/2, 1)$. Considering only executions with integer Parikh images, as it is for instance done in [7], would still be insufficient in order to prevent \mathcal{N} from covering $\{p, c_1, c_2\}$. Indeed, $\pi' \stackrel{\text{def}}{=} \pi\pi$ leads to $\{p, 4 \cdot c_1, 4 \cdot c_2\}$ and $\Psi(\pi') = (1, 1, 2, 1, 1, 2)$. This “erroneous behavior” is due to the fact that both q_a and q_b can continuously be reached simultaneously in \mathcal{N} .

³Of course, there are many different ways to transform a VASS into a Petri net, but in any such transformation, control states become places.

If we interpret \mathcal{V} as a \mathbb{Z} -VASS, $p(1, 1)$ could still be covered, since $p(0, 0) \xrightarrow{t_1 t_2 t_3 t_6} \mathbb{Z} p(3, 1)$. This example nicely illustrates the advantage CVASS have over \mathbb{Z} -VASS by not allowing counters to drop below zero, while still being decidable in NP. To be fair, there are also instances in which a CVASS can cover a configuration of an over-approximated VASS which is not coverable in a \mathbb{Z} -VASS. Consider Figure 5. We have $p(1) \xrightarrow{\frac{1}{2} \cdot t_1 t_2} \mathbb{Q} r(1)$, while $r(1)$ cannot be covered from $p(1)$ in the \mathbb{Z} -VASS semantics.




Fig. 5. Example where over-approximating over \mathbb{Z} is better than over \mathbb{Q}_+ .

VI. CONCLUDING REMARKS

While discrete Petri nets and VASS are known to be equivalent and often used as synonyms, our results show that this is not the case in the continuous setting: continuous Petri nets and CVASS define different sets of rationals, and have respectively P-complete and NP-complete reachability problems. In particular, we showed that reachability sets of CVASS are equivalent to those definable in existential $\text{FO}(\mathbb{Q}, +, <)$, and reachability sets of continuous Petri nets coincide with the sets of non-negative rational numbers definable in $\text{CSP}(\Gamma_c)$, a polynomial-time decidable fragment of existential $\text{FO}(\mathbb{Q}, +, <)$. The latter characterization might contribute towards understanding why reachability in continuous Petri nets with tens of thousands of places can still efficiently be decided by SMT solvers, as demonstrated in [7]. There has been some work showing that abstracted models of clause-learning based SAT solvers run in expected polynomial time on finite max-closed CSPs [26]. Though presumably challenging, it would be striking if variants of such results could be lifted to CSPs over infinite domains such as $\text{CSP}(\Gamma_c)$.

As for future work, in Section III-D we showed that continuous Petri nets are closely related to constant-rate MMS. In [22], the problem of finding runs in constant-rate MMS which are optimal with respect to a linear objective function was also shown to be decidable in polynomial time. It seems conceivable that an optimization variant of $\text{CSP}(\Gamma_c)$ could also be decidable in polynomial time, which would enable us to compute optimal runs in continuous Petri nets in polynomial

time. Finally, Le Roux, Pauly and Raskin have recently studied Minkowski games [27], which can be viewed as a game-theoretic variant of constant-rate MMS. Investigating various types of games on continuous Petri nets and CVASS appears to be a promising direction of future research.

ACKNOWLEDGMENTS

The authors are grateful to Dmitry Chistikov and Peter Jeavons for various insightful discussions related to $\text{CSP}(\Gamma_c)$.

REFERENCES

- [1] E. Cardoza, R. J. Lipton, and A. R. Meyer, “Exponential space complete problems for Petri nets and commutative semigroups: Preliminary report,” in *Proc. Symposium on Theory of Computing, STOC*. ACM, 1976, pp. 50–54.
- [2] J. Leroux and S. Schmitz, “Demystifying reachability in vector addition systems,” in *Proc. Logic in Computer Science, LICS*. IEEE Computer Society, 2015, pp. 56–67.
- [3] C. Haase and S. Halfon, “Integer vector addition systems with states,” in *Proc. Reachability Problems, RP*, ser. Lect. Notes Comp. Sci., vol. 8762. Springer, 2014, pp. 112–124.
- [4] R. David and H. Alla, “Continuous Petri nets,” in *Proc. Application and Theory of Petri nets, ATPN*, 1987, pp. 275–294.
- [5] —, *Discrete, continuous, and hybrid Petri nets*. Springer Science & Business Media, 2010.
- [6] E. Fraca and S. Haddad, “Complexity analysis of continuous Petri nets,” *Fundam. Inform.*, vol. 137, no. 1, pp. 1–28, 2015.
- [7] M. Blondin, A. Finkel, C. Haase, and S. Haddad, “Approaching the coverability problem continuously,” in *Proc. Tools and Algorithms for the Construction and Analysis of Systems, TACAS*, ser. Lect. Notes Comp. Sci., vol. 9636. Springer, 2016, pp. 480–496.
- [8] J. E. Hopcroft and J. Pansiot, “On the reachability problem for 5-dimensional vector addition systems,” *Theor. Comput. Sci.*, vol. 8, pp. 135–159, 1979.
- [9] T. Geffroy, J. Leroux, and G. Sutre, “Occam’s razor applied to the Petri net coverability problem,” in *Proc. Reachability Problems, RP*, ser. Lect. Notes Comp. Sci., vol. 9899. Springer, 2016, pp. 77–89.
- [10] P. A. Abdulla, K. Cerans, B. Jonsson, and Y. Tsay, “General decidability theorems for infinite-state systems,” in *Proc. Logic in Computer Science, LICS*. IEEE Computer Society, 1996, pp. 313–321.
- [11] J. Esparza, R. Ledesma-Garza, R. Majumdar, P. J. Meyer, and F. Niksic, “An SMT-based approach to coverability analysis,” in *Proc. Computer Aided Verification, CAV*, ser. Lect. Notes Comp. Sci., vol. 8559. Springer, 2014, pp. 603–619.
- [12] K. Athanasiou, P. Liu, and T. Wahl, “Unbounded-thread program verification using thread-state equations,” in *Automated Reasoning, IJCAR*, ser. Lect. Notes Comp. Sci., vol. 9706. Springer, 2016, pp. 516–531.
- [13] E. D. Sontag, “Real addition and the polynomial hierarchy,” *Inform. Process. Lett.*, vol. 20, no. 3, pp. 115–120, 1985.
- [14] R. Greenlaw, H. J. Hoover, and W. L. Ruzzo, *Limits to Parallel Computation: P-completeness Theory*. Oxford University Press, 1995.
- [15] L. G. Khachiyan, “A polynomial algorithm in linear programming,” *Doklady Akademii Nauk SSSR*, vol. 244, pp. 1093–1096, 1979.
- [16] C. Papadimitriou and K. Steiglitz, *Combinatorial Optimization: Algorithms and Complexity*. Dover Publications, 1982.
- [17] P. Jonsson and C. Bäckström, “A unifying approach to temporal constraint reasoning,” *Artif. Intell.*, vol. 102, no. 1, pp. 143–155, 1998.
- [18] M. Koubarakis, “Tractable disjunctions of linear constraints: basic results and applications to temporal reasoning,” *Theor. Comput. Sci.*, vol. 266, no. 1-2, pp. 311–339, 2001.
- [19] M. Bodirsky, P. Jonsson, and T. von Oertzen, “Essential convexity and complexity of semi-algebraic constraints,” *Log. Meth. Comput. Sci.*, vol. 8, no. 4, 2012.
- [20] K. N. Verma, H. Seidl, and T. Schwentick, “On the complexity of equational Horn clauses,” in *Proc. Automated Deduction - CADE*, vol. 3632. Springer, 2005, pp. 337–352.
- [21] F. Drewes and J. Leroux, “Structurally cyclic Petri nets,” *Log. Meth. Comput. Sci.*, vol. 11, no. 4, 2015.
- [22] R. Alur, A. Trivedi, and D. Wojtczak, “Optimal scheduling for constant-rate multi-mode systems,” in *Proc. Hybrid Systems: Computation and Control, HSCC*. ACM, 2012, pp. 75–84.
- [23] M. Heymann, F. Lin, G. Meyer, and S. Resmerita, “Analysis of Zeno behaviors in a class of hybrid systems,” *IEEE Trans. Automat. Contr.*, vol. 50, no. 3, pp. 376–383, 2005.
- [24] F. Laroussinie, N. Markey, and P. Schnoebelen, “Model checking timed automata with one or two clocks,” in *Proc. Concurrency Theory, CONCUR*, ser. Lect. Notes Comp. Sci., vol. 3170. Springer, 2004, pp. 387–401.
- [25] M. L. Minsky, *Computation: finite and infinite machines*. Prentice Hall, 1967.
- [26] P. Jeavons and J. Petke, “Local consistency and SAT-solvers,” *J. Artif. Intell. Res.*, vol. 43, pp. 329–351, 2012.
- [27] S. Le Roux, A. Pauly, and J.-F. Raskin, “Minkowski games,” in *Proc. Symposium on Theoretical Aspects of Computer Science, STACS*, vol. 66. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2017, pp. 50:1–50:13.
- [28] A. Schrijver, *Theory of Linear and Integer Programming*. John Wiley & Sons, 1986.

APPENDIX A
MISSING PROOFS FROM SECTION III

A. Missing proofs from Section III-B

In this section, we formally prove Lemma 3.5. Let us fix a continuous Petri net $\mathcal{N} = (P, T, \text{Pre}, \text{Post})$ for the remainder of this section. Recall that we want to build a formula expressing firing sets of \mathcal{N} . More formally, we build a formula φ such that $\varphi(\mathbf{u}, \mathbf{y})$ holds if and only if there exists some $\pi \in T^\star$ such that π is a firing sequence from \mathbf{u} and $\llbracket \pi \rrbracket = \llbracket \mathbf{y} \rrbracket$.

Let

$$\varphi(\mathbf{u}, \mathbf{y}) \stackrel{\text{def}}{=} \exists \mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_{|T|} \in \mathbb{Q}_+^{P \cup T} \quad \psi_0(\mathbf{x}_0) \wedge \bigwedge_{1 \leq i \leq |T|} \psi_i(\mathbf{x}_i, \mathbf{x}_{i-1}) \wedge \bigwedge_{t \in T} \mathbf{y}_t > 0 \leftrightarrow \bigvee_{1 \leq i \leq |T|} \mathbf{x}_{i,t} > 0$$

where

$$\psi_0(\mathbf{u}, \mathbf{x}_0) \stackrel{\text{def}}{=} \bigwedge_{p \in P} (\mathbf{x}_{0,p} > 0 \leftrightarrow \mathbf{u}(p) > 0) \wedge \bigwedge_{t \in T} \mathbf{x}_{0,t} = 0$$

$$\psi_i(\mathbf{x}_i, \mathbf{x}_{i-1}) \stackrel{\text{def}}{=} \bigwedge_{p \in P} (\mathbf{x}_{i,p} > 0 \leftrightarrow (\mathbf{x}_{i-1,p} > 0 \vee \bigvee_{t \in \bullet p} \mathbf{x}_{i,t} > 0)) \wedge \bigwedge_{t \in T} (\mathbf{x}_{i,t} > 0 \rightarrow \bigwedge_{p \in \bullet t} \mathbf{x}_{i-1,p} > 0).$$

In order to prove that φ is correct, let us define the following intermediary formula:

$$\Phi_n(\mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_n) \stackrel{\text{def}}{=} \psi_0(\mathbf{u}, \mathbf{x}_0) \wedge \bigwedge_{1 \leq i \leq n} \psi_i(\mathbf{x}_i, \mathbf{x}_{i-1}).$$

Lemma A.1: Let $\mathbf{u} \in \mathbb{Q}_+^P$ and $\mathbf{y} \in \mathbb{Q}_+^T$. For every $\pi \in T^\star$, if $\llbracket \pi \rrbracket = \llbracket \mathbf{y} \rrbracket$ and π is a firing sequence from \mathbf{u} , then there exist $\mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_n \in \mathbb{Q}_+^{P \cup T}$, where $n = \llbracket \pi \rrbracket$, such that

- (a) $\Phi_n(\mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_n)$ is satisfied,
- (b) $t \in \llbracket \pi \rrbracket$ if and only if $\mathbf{x}_{i,t} > 0$ for some $i \in [n]$,
- (c) $\mathbf{x}_{n,p} > 0$ for every $p \in \llbracket \mathbf{u} \rrbracket \cup \llbracket \pi \rrbracket^\bullet$.

Proof: We prove the claim by induction on n . If $n = 0$, then $\pi = \varepsilon$. For every $q \in P \cup T$, we let

$$\mathbf{x}_{0,q} \stackrel{\text{def}}{=} \begin{cases} 1 & \text{if } q \in P \text{ and } \mathbf{u}(q) > 0, \\ 0 & \text{otherwise.} \end{cases}$$

The validity of (a), (b) and (c) follows immediately.

Assume $n > 0$ and that the claim holds for every sequence of length smaller than n . There exist $\sigma, \sigma' \in T^\star$, $\alpha \in (0, 1]$, $s \in T$ such that

- $\pi = \sigma \alpha s \sigma'$, and
- $\llbracket \sigma \rrbracket = \llbracket \pi \rrbracket \setminus \{s\}$.

Let $\pi' \stackrel{\text{def}}{=} \sigma \alpha s$. Note that π' is a firing sequence from \mathbf{u} , and $\llbracket \pi' \rrbracket = \llbracket \pi \rrbracket$. Thus, we can build an assignment from π' instead of π . By induction hypothesis, there exists an assignment $\mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_{n-1} \in \mathbb{Q}_+^{P \cup T}$ such that

- (1) $\Phi_{n-1}(\mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_{n-1})$ is satisfied,
- (2) $t \in \llbracket \sigma \rrbracket$ if and only if $\mathbf{x}_{i,t} > 0$ for some $i \in [n-1]$,
- (3) $\mathbf{x}_{n-1,p} > 0$ for every $p \in \llbracket \mathbf{u} \rrbracket \cup \llbracket \sigma \rrbracket^\bullet$.

It remains to choose the values of \mathbf{x}_n . For every $p \in P$ and $t \in T$, we let

$$\mathbf{x}_{n,p} \stackrel{\text{def}}{=} \begin{cases} 1 & \text{if } \mathbf{x}_{n-1,p} > 0 \text{ or } s \in \bullet p, \\ 0 & \text{otherwise.} \end{cases}$$

and

$$\mathbf{x}_{n,t} \stackrel{\text{def}}{=} \begin{cases} 1 & \text{if } t = s, \\ 0 & \text{otherwise.} \end{cases}$$

It is readily seen that for every $p \in P$,

$$\mathbf{x}_{n,p} > 0 \leftrightarrow (\mathbf{x}_{n-1,p} > 0 \vee \bigvee_{t \in \bullet p} \mathbf{x}_{n,t} > 0).$$

Let us show that the following holds:

$$\bigwedge_{p \in \bullet s} \mathbf{x}_{n-1,p} > 0. \quad (4)$$

Since π' is a firing sequence from \mathbf{u} , there exist $\mathbf{v}, \mathbf{w} \in \mathbb{Q}_+^P$ such that

$$\mathbf{u} \xrightarrow{\sigma}_{\mathbb{Q}_+} \mathbf{v} \xrightarrow{\alpha s}_{\mathbb{Q}_+} \mathbf{w}.$$

We have $\bullet s \subseteq \llbracket \mathbf{v} \rrbracket \subseteq \llbracket \mathbf{u} \rrbracket \cup \llbracket \sigma \rrbracket$. By (3), this implies that $\mathbf{x}_{n-1,p} > 0$ for every $p \in \bullet s$. Therefore, (4) holds. This means that ψ_n is satisfied, and, altogether with (1), that Φ_n is satisfied. Thus, (a) holds. Moreover, by (2), (3) and the assignment of \mathbf{x} , (b) and (c) hold. \square

Corollary A.2: Let $\mathbf{u} \in \mathbb{Q}_+^P$, $\mathbf{y} \in \mathbb{Q}_+^T$, and let $\pi \in T^\star$ be a firing sequence from \mathbf{u} such that $\llbracket \pi \rrbracket = \llbracket \mathbf{y} \rrbracket$. Formula φ is satisfied by (\mathbf{u}, \mathbf{y}) .

Proof: Let $n \stackrel{\text{def}}{=} \llbracket \llbracket \pi \rrbracket \rrbracket$. By Lemma A.1, there exist $\mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_n \in \mathbb{Q}_+^{P \cup T}$ such that

$$\Phi_n(\mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_n) \text{ is satisfied.} \quad (5)$$

and

$$t \in \llbracket \pi \rrbracket \text{ if and only if } \mathbf{x}_{i,t} > 0 \text{ for some } i \in [n]. \quad (6)$$

For every $n < i \leq |T|$, $q \in P \cup T$, let

$$\mathbf{x}_{i,q} \stackrel{\text{def}}{=} \begin{cases} 1 & \text{if } q \in P \text{ and } \mathbf{x}_{i-1,q} > 0, \\ 0 & \text{otherwise.} \end{cases}$$

For every $n < i \leq |T|$, $\psi_i(\mathbf{x}_i, \mathbf{x}_{i-1})$ holds. Therefore, with (5), this implies that $\Phi_{|T|}(\mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_{|T|})$ is satisfied. By (6), we conclude that $\varphi(\mathbf{u}, \mathbf{v})$ is satisfiable. \square

Lemma A.3: Let $\mathbf{u} \in \mathbb{Q}_+^P$ and $\mathbf{y} \in \mathbb{Q}_+^T$. Let $0 \leq n \leq |T|$ and let $\mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_n \in \mathbb{Q}_+^{P \cup T}$. If $\Phi_n(\mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_n)$ is satisfied, then there exist $\pi \in T^\star$ and $\mathbf{v} \in \mathbb{Q}_+^P$ such that

- (a) $\mathbf{u} \xrightarrow{\pi}_{\mathbb{Q}_+} \mathbf{v}$,
- (b) $\llbracket \pi \rrbracket = T_n$,
- (c) $\llbracket \mathbf{v} \rrbracket = P_n \cup T_n^\bullet$.

where

$$P_n = \{p \in P; \mathbf{x}_{n,p} > 0\} \text{ and} \\ T_n = \{t \in T; \mathbf{x}_{i,t} > 0 \text{ for some } i \in [n]\}.$$

Proof: We prove the claim by induction on n . Suppose $n = 0$. We have $P_0 = \llbracket \mathbf{u} \rrbracket$ and $T_0 = \emptyset$. Thus, $\mathbf{v} \stackrel{\text{def}}{=} \mathbf{u}$ and $\pi \stackrel{\text{def}}{=} \varepsilon$ satisfy (a), (b) and (c).

Assume $n > 0$ and that the claim holds for $n - 1$. Since $\Phi_n(\mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_n)$ is satisfied, $\Phi_{n-1}(\mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_{n-1})$ is also satisfied. By induction hypothesis, there exist $\mathbf{w} \in \mathbb{Q}_+^P$ and $\sigma \in T^\star$ such that

$$\mathbf{u} \xrightarrow{\sigma}_{\mathbb{Q}_+} \mathbf{w}, \llbracket \sigma \rrbracket = T_{n-1} \text{ and } \llbracket \mathbf{w} \rrbracket = P_{n-1} \cup T_{n-1}^\bullet. \quad (7)$$

Let $S \stackrel{\text{def}}{=} T_n \setminus T_{n-1} = \{t_1, t_2, \dots, t_k\}$. Note that $\mathbf{x}_{n,t_i} > 0$ for every $i \in [k]$. Thus, by definition of ψ_n , we have $\mathbf{x}_{i-1,p} > 0$ for every $p \in \bullet S$. By 7, this implies that

$$\bullet S \subseteq P_{n-1} \subseteq \llbracket \mathbf{w} \rrbracket. \quad (8)$$

Let

$$\gamma \stackrel{\text{def}}{=} \max(1, \max\{\lceil 2/\mathbf{w}(p) \rceil : p \in \llbracket \mathbf{w} \rrbracket\}).$$

For every $i \in [k]$, let $\beta_i \stackrel{\text{def}}{=} \max(1, \max\{\mathbf{Pre}(p, t_i) : p \in \bullet t_i\})$ and $\alpha_i \stackrel{\text{def}}{=} 1/(\gamma \cdot \beta_i)$. Let $i \in [k]$ and $p \in \bullet t_i$. We have

$$\begin{aligned} \alpha_i \cdot \mathbf{Pre}(p, t_i) &= \mathbf{Pre}(p, t_i) / (\gamma \cdot \beta_i) \\ &\leq \mathbf{Pre}(p, t_i) / (\lceil 2/\mathbf{w}(p) \rceil \cdot \beta_i) \\ &\leq \mathbf{Pre}(p, t_i) / (2/\mathbf{w}(p) \cdot \beta_i) \\ &\leq \mathbf{Pre}(p, t_i) / (2/\mathbf{w}(p) \cdot \mathbf{Pre}(p, t_i)) \\ &= \mathbf{w}(p) / 2. \end{aligned}$$

Therefore,

$$\begin{aligned} \mathbf{w}(p) - \alpha_i \cdot \mathbf{Pre}(p, t_i) &\geq \mathbf{w}(p) - \mathbf{w}(p)/2 \\ &= \mathbf{w}(p)/2 \\ &> 0 \end{aligned} \quad (\text{by (8)}).$$

This implies that $\mathbf{w} \xrightarrow{\alpha_i t_i}_{\mathbb{Q}_+} \mathbf{v}_i$ for some $\mathbf{v}_i \in \mathbb{Q}_+^P$ such that $\llbracket \mathbf{v}_i \rrbracket = \llbracket \mathbf{w} \rrbracket \cup t_i^\bullet$. Let

$$\sigma' \stackrel{\text{def}}{=} (\alpha_1/k)t_1 (\alpha_2/k)t_2 \cdots (\alpha_k/k)t_k$$

and $\mathbf{v} \stackrel{\text{def}}{=} (\mathbf{v}_1 + \mathbf{v}_2 + \cdots + \mathbf{v}_k)/k$. By [6, Lemma 12], we have $\mathbf{w} \xrightarrow{\sigma'}_{\mathbb{Q}_+} \mathbf{v}$. Therefore, $\mathbf{u} \xrightarrow{\sigma\sigma'}_{\mathbb{Q}_+} \mathbf{v}$, which implies that (a) and (b) hold.

It remains to show (c). First note that

$$\begin{aligned} \llbracket \mathbf{v} \rrbracket &= \llbracket \mathbf{v}_1 \rrbracket \cup \llbracket \mathbf{v}_2 \rrbracket \cup \cdots \cup \llbracket \mathbf{v}_k \rrbracket \\ &= \llbracket \mathbf{w} \rrbracket \cup t_1^\bullet \cup t_2^\bullet \cup \cdots \cup t_k^\bullet \\ &= (P_{n-1} \cup T_{n-1}^\bullet) \cup S^\bullet && (\text{by (7)}) \\ &= P_{n-1} \cup T_n^\bullet && (\text{by def. of } S). \end{aligned}$$

By definition of ψ_n , $P_{n-1} \subseteq P_n$. Thus, $\llbracket \mathbf{v} \rrbracket \subseteq P_n \cup T_n^\bullet$.

We conclude by proving that $P_n \subseteq P_{n-1} \cup T_n^\bullet$. Let $p \in P_n$. By definition of ψ_n , either $p \in P_{n-1}$ or $\mathbf{x}_{n,t} > 0$ for some $t \in \bullet p$. If $p \in P_{n-1}$, then we are done. If $\mathbf{x}_{n,t} > 0$ for some $t \in \bullet p$, then $t \in T_n$ which implies that $p \in t^\bullet \subseteq T_n^\bullet$. \square

Corollary A.4: Let $\mathbf{u} \in \mathbb{Q}_+^P$ and $\mathbf{y} \in \mathbb{Q}_+^T$. If φ is satisfied by (\mathbf{u}, \mathbf{y}) , then there exists $\pi \in T^\star$ such that π is a firing sequence from \mathbf{u} , and $\llbracket \pi \rrbracket = \llbracket \mathbf{y} \rrbracket$.

Proof: Assume φ is satisfied by (\mathbf{u}, \mathbf{y}) . There exist $\mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_{|T|} \in \mathbb{Q}_+^{P \cup T}$ such that $\Phi_n(\mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_{|T|})$ is satisfied. By Lemma A.3, there exists $\pi \in T^\star$ such that π is a firing sequence from \mathbf{u} , and

$$\llbracket \pi \rrbracket = \{t \in T : \mathbf{x}_{|T|,n} > i \text{ for some } 0 \leq i \leq |T|\}.$$

We are done since by definition of φ , $\llbracket \mathbf{y} \rrbracket = \{t \in T : \mathbf{x}_{|T|,n} > i \text{ for some } 0 \leq i \leq |T|\}$. \square

Lemma A.5: φ belongs to $\text{CSP}(\Gamma_c)$.

Proof: By definition, φ is a conjunction of constraints of the form

- $x = 0$,
- $x > 0 \leftrightarrow \bigvee_{y \in Y} y > 0$, and
- $x > 0 \rightarrow \bigwedge_{y \in Y} y > 0$.

Constraints of the first type can be written as $-x \geq 0$. The two other types of constraints can be rewritten by using the following equivalences:

$$x > 0 \leftrightarrow \bigvee_{y \in Y} y > 0 \equiv (x = 0 \vee \bigvee_{y \in Y} y > 0) \wedge \bigwedge_{y \in Y} (y = 0 \vee x > 0),$$

and

$$x > 0 \rightarrow \bigwedge_{y \in Y} y > 0 \equiv x = 0 \vee \bigwedge_{y \in Y} y > 0. \quad \square$$

B. Missing proofs from Section III-C

We detail the construction described in the proof of Lemma 3.7:

Formal construction of Lemma 3.7: First note that we can assume coefficients appearing in \mathbf{a} and \mathbf{c} to have integer values, since $(a_1/d_1)x_1 + \cdots + (a_m/d_m)x_m \sim c$ has the same solutions as $\lambda(a_1/d_1)x_1 + \cdots + \lambda(a_m/d_m)x_m \sim \lambda \cdot c$ where $\lambda \stackrel{\text{def}}{=} d_1 \cdots d_m$. Note that λ can be computed in logarithmic space.

Formally, \mathcal{N} is defined as $\mathcal{N} \stackrel{\text{def}}{=} (P, T, \mathbf{Pre}, \mathbf{Post})$ where

$$\begin{aligned} P &\stackrel{\text{def}}{=} \{p_i : i \in [n]\} \cup \{q_{\text{pos}}, q_{\text{neg}}, q_{\geq}, q_{>}, q'_{>}, q_{\theta}\}, \\ T &\stackrel{\text{def}}{=} \{s_i : i \in [n]\} \cup \{s'_i : i \in [m]\} \cup \{t_p : p \in P\} \cup \{u_1, u_2, u_3, u_4, u_5\}. \end{aligned}$$

For every $i \in [n]$, we have

$$\begin{aligned} \mathbf{Post}(q_{\text{pos}}, s_i) &= a_i & \text{if } a_i \geq 0, \\ \mathbf{Post}(q_{\text{pos}}, u_1) &= c & \text{if } c \geq 0, \\ \mathbf{Post}(q_{\text{neg}}, s_i) &= |a_i| & \text{if } a_i < 0, \\ \mathbf{Post}(q_{\text{neg}}, u_1) &= |c| & \text{if } c < 0. \end{aligned}$$

Moreover,

$$\begin{aligned} \mathbf{Post}(q_{\text{pos}}, u_5) &= 1, & \mathbf{Post}(q_{\text{pos}}^\bullet, t_{q_{\text{pos}}}) &= 1, \\ \mathbf{Post}(q_{\text{neg}}, u_5) &= 1, & \mathbf{Post}(q_{\text{neg}}^\bullet, t_{q_{\text{neg}}}) &= 1. \end{aligned}$$

The remaining arcs, labeled by 1, are defined by:

$$\begin{aligned} \bullet p_i &= \{s'_j : j \in [m] \text{ and } x_i \in \theta_j\}, & p_i^\bullet &= \bullet p_i \cup \{s_i, t_{p_i}\}, \\ \bullet q_{\geq} &= \emptyset, & q_{\geq}^\bullet &= \{u_1, t_{q_{\geq}}\}, \\ \bullet q'_{>} &= \{u_2, u_3\}, & q'_{>}^\bullet &= \{u_3, u_4, t_{q'_{>}}\}, \\ \bullet q_{>} &= \emptyset, & q_{>}^\bullet &= \{u_3, t_{q_{>}}\}, \\ \bullet q_\theta &= \{t_p : p \in P \setminus \{q_\theta\}\}, & q_\theta^\bullet &= \{t_p : p \in P\}. \end{aligned}$$

It can be proven that any firing sequence π not using transitions from $\{s'_i : i \in [m]\} \cup \{t_p : p \in P\}$ can be rearranged into a firing sequence π' such that $\delta(\pi') = \delta(\pi)$ and $\pi' \in s_1^\star s_2^\star \cdots s_n^\star u_1^\star u_2^\star u_3^\star u_4^\star u_5^\star$. \square

Theorem 3.8: Let $\Phi(\mathbf{x})$ be an instance of $\text{CSP}(\Gamma_c)$ with free variables \mathbf{x} . There exists a logarithmic-space computable continuous Petri net \mathcal{N} defining the set $[\![\Phi]\!]$ of solutions of $\Phi(\mathbf{x})$.

Proof: Let $\mathbf{x} = (x_1, x_2, \dots, x_n)$ and $\Phi = \{\psi_1, \psi_2, \dots, \psi_k\}$. For every $i \in [k]$, let \mathcal{N}_i and \mathbf{y}_i be respectively the continuous Petri net and marking obtained from Lemma 3.7 for ψ_i . We construct the required continuous Petri net \mathcal{N} as follows. First, we take the disjoint union of $\mathcal{N}_1, \mathcal{N}_2, \dots, \mathcal{N}_k$. Next, we add new places p_1, p_2, \dots, p_n and new transitions t_1, t_2, \dots, t_n . Each place p_i initially contains the value of x_i , and t_i allows to “transfer” the content of p_i simultaneously to the places of $\mathcal{N}_1, \mathcal{N}_2, \dots, \mathcal{N}_k$ that correspond to x_i . It then follows from Lemma 3.7 that $\Phi(\mathbf{x})$ holds if and only if $(\mathbf{x}, \mathbf{0}, \mathbf{y}_1, \mathbf{0}, \mathbf{y}_2, \dots, \mathbf{0}, \mathbf{y}_k) \xrightarrow{*}_{\mathbb{Q}_+} \mathbf{0}$. \square

C. Missing proofs from Section III-D

Corollary 3.11: Safe schedulability and safe reachability in constant-rate MMS are logarithmic-space reducible to continuous Petri net reachability problems and thus belong to P.

Proof: An MMS $\mathcal{H} = (M, n, R)$ with an initial configuration \mathbf{u} can be viewed as a continuous Petri net \mathcal{N} in which there are two places per continuous variable such that the value of a continuous variable is encoded as the difference between the corresponding two places. At any time, \mathcal{N} can invoke transitions that simultaneously decrease the two places simulating a continuous variable, analogously to the orange transitions in the continuous Petri nets constructed in Lemma 3.7. A safety region $S: A \cdot \mathbf{x} \geq \mathbf{b}$ can be incorporated into this Petri net by adding one place for each row $\mathbf{a}_i^\top \cdot \mathbf{x} \geq b_i$ of S which initially carries $\mathbf{a}_i^\top \cdot \mathbf{u} - b_i$ many tokens. Any transition corresponding to a mode m then also needs to push or respectively consume $|\mathbf{a}_i^\top \cdot R(m)|$ from those additional places. Theorem 3.6 then yields the desired upper bound. \square

APPENDIX B MISSING PROOFS FROM SECTION IV

Proposition 4.1: Let $\mathcal{N} = (P, T, \mathbf{Pre}, \mathbf{Post})$ be a continuous Petri net. There exist a logarithmic-space computable $(|T|+|P|)$ -CVASS $\mathcal{V} = (Q, S)$ and $q \in Q$ such that

$$\mathbf{u} \xrightarrow{*}_{\mathbb{Q}_+} \mathbf{v} \text{ in } \mathcal{N} \text{ if and only if } q(\mathbf{0}, \mathbf{u}) \xrightarrow{*}_{\mathbb{Q}_+} q(\mathbf{0}, \mathbf{v}) \text{ in } \mathcal{V}.$$

Proof: Let $T = \{t_1, t_2, \dots, t_n\}$. We let $\mathcal{V} \stackrel{\text{def}}{=} (Q, S)$ be the continuous $(|T|+|P|)$ -VASS such that

$$Q \stackrel{\text{def}}{=} \{q\} \cup \{p_i : i \in [n]\} \text{ and } S \stackrel{\text{def}}{=} \{s_i, t_i : i \in [n]\},$$

where $s_i \stackrel{\text{def}}{=} (q, [e_i, -\mathbf{Pre}_i], p_i)$ and $t_i \stackrel{\text{def}}{=} (p_i, [-e_i, \mathbf{Post}_i], q)$ for every $i \in [n]$. We illustrate the construction in Figure 6.




Fig. 6. Example of conversion from a continuous Petri net to a CVASS.

If $\mathbf{u} \xrightarrow{*}_{\mathbb{Q}_+} \mathbf{v}$ in \mathcal{N} , then it is clear that $q(\mathbf{0}, \mathbf{u}) \xrightarrow{*}_{\mathbb{Q}_+} q(\mathbf{0}, \mathbf{v})$ in \mathcal{V} . Let us show the converse. Suppose there exists $\pi \in S^\star$ such that $q(\mathbf{0}, \mathbf{u}) \xrightarrow{\pi}_{\mathbb{Q}_+} q(\mathbf{0}, \mathbf{v})$ in \mathcal{V} . First note that π is of the form $\pi = \alpha_1 s_{\ell_1} \beta_1 t_{\ell_1} \cdots \alpha_k s_{\ell_k} \beta_k t_{\ell_k}$. If $\alpha_i = \beta_i$, we say that i is *good*, otherwise we say that i is *bad*. We claim that π can be rearranged into an equivalent run with fewer bad indices. Formally, we show that there exists $\pi' \in S^\star$ such that $\delta(\pi') = \delta(\pi)$, $\text{path}(\pi') = \text{path}(\pi)$, π' is q -admissible, and π' has fewer bad indices than π . The claim implies that the number of bad indices can be decreased until there are only good indices, in which case we obtain a run in \mathcal{V} corresponding to a run of \mathcal{N} .

Let us prove the claim. Let $i \in [k]$ be the smallest bad index. We must have $\alpha_i > \beta_i$, otherwise counter ℓ_i would drop below zero. Since $\delta(\pi)(\ell_i) = 0$, there exists $j \in (i, k]$ such that $\ell_j = \ell_i$ and $\beta_j > \alpha_j$. Let j be the smallest such index. We consider two cases.

- Suppose $\beta_j - \alpha_j \geq \alpha_i - \beta_i$. Let

$$\pi' \stackrel{\text{def}}{=} \alpha'_1 s_{\ell_1} \beta'_1 t_{\ell_1} \cdots \alpha'_k s_{\ell_k} \beta'_k t_{\ell_k}$$

where $\alpha'_x \stackrel{\text{def}}{=} \alpha'_x$ for every $x \in [k]$, and

$$\beta'_x \stackrel{\text{def}}{=} \begin{cases} \alpha_i & \text{if } x = i, \\ \beta_j - (\alpha_i - \beta_i) & \text{if } x = j, \\ \beta_x & \text{otherwise.} \end{cases}$$

Note that $\beta'_i, \beta'_j \in (0, 1]$, hence $\pi' \in S^\star$. We have $\delta(\pi') = \delta(\pi)$ since $\beta'_i + \beta'_j = \beta_i + \beta_j$. Moreover, $\text{path}(\pi') = \text{path}(\pi)$, and i is now good in π' . It remains to argue that π' is q -admissible. Recall that $\delta(t_{\ell_i}) = [-e_{\ell_i}, \text{Post}_{\ell_i}]$. Since $\ell_i = \ell_j$, $\text{Post}_{\ell_i} \geq \mathbf{0}$ and β_i is increased at the expense of decreasing β_j , the n last counters still never drop below zero. By minimality of j , we have $\alpha_x \geq \beta_x$ for every $x \in (i, j)$. Thus, for every $y \in [i, j)$, we have

$$\sum_{x=1}^y (\alpha'_x - \beta'_x) = \sum_{\substack{x=1 \\ x \neq i}}^y (\alpha_x - \beta_x) + (\alpha_i - \beta'_i) = \sum_{x=i+1}^y (\alpha_x - \beta_x) + (\alpha_i - \beta'_i) \geq \alpha_i - \beta'_i = 0.$$

Moreover,

$$\sum_{x=1}^j (\alpha'_x - \beta'_x) = \sum_{x=1}^{j-1} (\alpha'_x - \beta'_x) + (\alpha'_j + \beta'_j) \geq \alpha'_j - \beta'_j = \alpha_j - (\beta_j - (\alpha_i - \beta_i)) = (\alpha_j - \beta_j) - (\alpha_i - \beta_i) \geq 0.$$

Therefore, counter ℓ_i never drops below zero, which implies that π' is q -admissible.

- Suppose $\beta_j - \alpha_j < \alpha_i - \beta_i$. We can make j good by choosing $\pi' \stackrel{\text{def}}{=} \alpha'_1 s_{\ell_1} \beta'_1 t_{\ell_1} \cdots \alpha'_k s_{\ell_k} \beta'_k t_{\ell_k}$ where $\alpha'_x \stackrel{\text{def}}{=} \alpha'_x$ for every $x \in [k]$, and

$$\beta'_x \stackrel{\text{def}}{=} \begin{cases} \beta_i + (\beta_j - \alpha_j) & \text{if } x = i, \\ \alpha_j & \text{if } x = j, \\ \beta_x & \text{otherwise.} \end{cases}$$

The proof is symmetric to the previous case. □

A. Missing proofs from Section IV-A

We first introduce some notation for this subsection. Let $\pi = \alpha_1 t_1 \alpha_2 t_2 \cdots \alpha_n t_n \in T^\star$. For every $1 \leq i \leq j \leq n$, we define $\pi[i..j]$ as the subsequence $\pi[i..j] \stackrel{\text{def}}{=} \alpha_i t_i \alpha_{i+1} t_{i+1} \cdots \alpha_j t_j$.

Lemma 4.2: The following statements hold for every $\pi \in T^\star$ and $\alpha, \alpha' \in (0, 1]$:

- (a) $p(\mathbf{u}) \xrightarrow{\pi}_{\mathbb{Q}_+} q(\mathbf{v})$ if and only if $p(\alpha \cdot \mathbf{u}) \xrightarrow{\alpha \cdot \pi}_{\mathbb{Q}_+} q(\alpha \cdot \mathbf{v})$,
- (b) if $p(\mathbf{u}) \xrightarrow{\alpha \pi}_{\mathbb{Q}_+} q(\mathbf{v})$, $p(\mathbf{u}') \xrightarrow{\alpha' \pi}_{\mathbb{Q}_+} q(\mathbf{v}')$ and $\alpha + \alpha' \leq 1$, then $p(\mathbf{u} + \mathbf{u}') \xrightarrow{(\alpha + \alpha')\pi}_{\mathbb{Q}_+} q(\mathbf{v} + \mathbf{v}')$,
- (c) if $p(\mathbf{u}) \xrightarrow{\pi}_{\mathbb{Q}_+} q(\mathbf{v})$, then $p(\mathbf{u}) \xrightarrow{\alpha \pi}_{\mathbb{Q}_+} q((1 - \alpha) \cdot \mathbf{u} + \alpha \cdot \mathbf{v})$,
- (d) $p(\mathbf{u}) \xrightarrow{\pi}_{\mathbb{Q}_+} q(\mathbf{v})$ in \mathcal{V} if and only if $q(\mathbf{v}) \xrightarrow{\pi^\dagger}_{\mathbb{Q}_+} p(\mathbf{u})$ in \mathcal{V}^\dagger .

Proof: All claims are immediate for $\pi = \varepsilon$. Let $\pi = \beta_1 t_1 \beta_2 t_2 \cdots \beta_n t_n$ where $n > 0$. We prove all claims by induction on n .

- (a) Follows by a simple induction on n using the fact that $a \cdot \mathbf{x} + a \cdot b \cdot \delta(t) \geq \mathbf{0} \iff \mathbf{x} + b \cdot \delta(t) \geq \mathbf{0}$ for every $\mathbf{x} \in \mathbb{Q}_+^d$, $a, b \in (0, 1]$ and $t \in T$.
- (b) Assume the claim holds for sequences of length less than n . Let $r(\mathbf{w}), r(\mathbf{w}') \in Q \times \mathbb{Q}_+^d$ be such that

$$\begin{aligned} p(\mathbf{u}) &\xrightarrow{\alpha \pi[1..n-1]}_{\mathbb{Q}_+} r(\mathbf{w}) \xrightarrow{\alpha \beta_n t_n}_{\mathbb{Q}_+} q(\mathbf{v}), \\ p(\mathbf{u}') &\xrightarrow{\alpha' \pi[1..n-1]}_{\mathbb{Q}_+} r(\mathbf{w}') \xrightarrow{\alpha' \beta_n t_n}_{\mathbb{Q}_+} q(\mathbf{v}'). \end{aligned}$$

By induction hypothesis,

$$p(\mathbf{u} + \mathbf{u}') \xrightarrow{(\alpha + \alpha')\pi[1..n-1]}_{\mathbb{Q}_+} r(\mathbf{w} + \mathbf{w}').$$

Now, observe that $(\alpha + \alpha') \cdot \beta_n \in (0, 1]$, hence $(\alpha + \alpha') \cdot \beta_n t_n \in T^\star$. We are done since,

$$\begin{aligned} \mathbf{w} + \mathbf{w}' + (\alpha + \alpha') \cdot \beta_n \cdot \delta(t_n) &= (\mathbf{w} + \alpha \cdot \beta_n \cdot \delta(t_n)) + (\mathbf{w}' + \alpha' \cdot \beta_n \cdot \delta(t_n)) \\ &= \mathbf{v} + \mathbf{v}' \\ &\geq \mathbf{0}. \end{aligned}$$

- (c) Assume the claim holds for sequences of length less than n . Let $r(\mathbf{w}) \in Q \times \mathbb{Q}_+^d$ be such that

$$p(\mathbf{u}) \xrightarrow{\pi[1..n-1]}_{\mathbb{Q}_+} r(\mathbf{w}) \xrightarrow{\beta_n t_n}_{\mathbb{Q}_+} q(\mathbf{v}).$$

By induction hypothesis, $p(\mathbf{u}) \xrightarrow{\alpha \pi[1..n-1]}_{\mathbb{Q}_+} r(\mathbf{w}')$ where $\mathbf{w}' \stackrel{\text{def}}{=} (1 - \alpha) \cdot \mathbf{u} + \alpha \cdot \mathbf{w}$. Since, $\mathbf{w} + \beta_n \cdot \delta(t_n) = \mathbf{v} \geq \mathbf{0}$ and $\mathbf{w}' \geq \alpha \cdot \mathbf{w}$, we obtain

$$r(\mathbf{w}') \xrightarrow{\alpha \cdot \beta_n t_n}_{\mathbb{Q}_+} q(\mathbf{w}' + \alpha \cdot \beta_n \cdot \delta(t_n))$$

Now, we are done since

$$\begin{aligned} \mathbf{w}' + \alpha \cdot \beta_n \cdot \delta(t_n) &= (1 - \alpha) \cdot \mathbf{u} + \alpha \cdot \mathbf{w} + \alpha \cdot \beta_n \cdot \delta(t_n) \\ &= (1 - \alpha) \cdot \mathbf{u} + \alpha \cdot (\mathbf{w} + \beta_n \cdot \delta(t_n)) \\ &= (1 - \alpha) \cdot \mathbf{u} + \alpha \cdot \mathbf{v}. \end{aligned}$$

- (d) Follows by a simple induction on n using the fact that $r(\mathbf{x}) \xrightarrow{at}_{\mathbb{Q}_+} r'(\mathbf{x}')$ if and only if $r'(\mathbf{x}') \xrightarrow{at^\dagger}_{\mathbb{Q}_+} r(\mathbf{x})$ for every $r(\mathbf{x}), r'(\mathbf{x}') \in Q \times \mathbb{Q}_+^d$ and $at \in T^\star$. □

Lemma 4.3: Let $q \in Q$, $\mathbf{u}, \mathbf{v} \in \mathbb{Q}_+^d$, and $\pi \in T^\star$ be such that

- (a) $q(\mathbf{u}) \xrightarrow{\pi}_{\mathbb{Q}} q(\mathbf{v})$,
- (b) $\llbracket \delta(t) \rrbracket^- \subseteq \llbracket \mathbf{u} \rrbracket$ for every $t \in \llbracket \pi \rrbracket$, and
- (c) $\llbracket \delta(t) \rrbracket^+ \subseteq \llbracket \mathbf{v} \rrbracket$ for every $t \in \llbracket \pi \rrbracket$.

There exists $\pi' \in T^\star$ such that $q(\mathbf{u}) \xrightarrow{\pi'}_{\mathbb{Q}_+} q(\mathbf{v})$ and $\Psi(\pi') = \Psi(\pi)$.

Proof: Let

$$\begin{aligned}\text{neg}(i) &\stackrel{\text{def}}{=} \sum_{\substack{t \in T \\ i \in \llbracket \delta(t) \rrbracket^-}} \Psi(\pi)(t) \cdot |\delta(t)(i)|, \\ \text{pos}(i) &\stackrel{\text{def}}{=} \sum_{\substack{t \in T \\ i \in \llbracket \delta(t) \rrbracket^+}} \Psi(\pi)(t) \cdot \delta(t)(i).\end{aligned}$$

Let

$$\begin{aligned}\alpha &\stackrel{\text{def}}{=} \min(1, \min\{\mathbf{u}(i)/\text{neg}(i) : i \in [d] \text{ and } \text{neg}(i) > 0\}), \\ \beta &\stackrel{\text{def}}{=} \min(1, \min\{\mathbf{v}(i)/\text{pos}(i) : i \in [d] \text{ and } \text{pos}(i) > 0\}).\end{aligned}$$

By (b) and (c), $\alpha > 0$ and $\beta > 0$. Let $n \stackrel{\text{def}}{=} \max(\lceil \frac{1}{\alpha} \rceil, \lceil \frac{1}{\beta} \rceil, 2)$. Let $\sigma \stackrel{\text{def}}{=} \frac{1}{n}\pi$ and $\pi' \stackrel{\text{def}}{=} \sigma^n$. We show that π' is as required. It is readily seen that $\Psi(\pi) = \Psi(\pi')$. It remains to show that $q(\mathbf{u}) \xrightarrow{\pi'}_{\mathbb{Q}_+} q(\mathbf{v})$.

For every $i \in [n]$, let $\mathbf{u}_i \stackrel{\text{def}}{=} \mathbf{u} + \frac{i}{n}\delta(\pi)$. We claim that $q(\mathbf{u}_0) \xrightarrow{\sigma}_{\mathbb{Q}_+} q(\mathbf{u}_1)$. Let $1 \leq \ell \leq |\pi|$ and $i \in [d]$. If $\text{neg}(i) = 0$, then

$$\mathbf{u}_0(i) + \delta(\sigma[1..\ell])(i) \geq \mathbf{u}_0(i).$$

If $\text{neg}(i) > 0$, then

$$\begin{aligned}\mathbf{u}_0(i) + \delta(\sigma[1..\ell])(i) &= \mathbf{u}(i) + \sum_{\alpha t \in \pi[1..\ell]} \frac{\alpha}{n} \cdot \delta(t)(i) \\ &\geq \mathbf{u}(i) + \frac{1}{n} \cdot \sum_{\substack{\alpha t \in \pi[1..\ell] \\ i \in \llbracket \delta(t) \rrbracket^-}} \alpha \cdot \delta(t)(i) \\ &\geq \mathbf{u}(i) - \frac{1}{n} \cdot \text{neg}(i) \\ &\geq \mathbf{u}(i) - \frac{1}{\lceil 1/\alpha \rceil} \cdot \text{neg}(i) \\ &\geq \mathbf{u}(i) - \alpha \cdot \text{neg}(i) \\ &\geq \mathbf{u}(i) - (\mathbf{u}(i)/\text{neg}(i)) \cdot \text{neg}(i) \\ &\geq 0.\end{aligned}$$

Similarly, $q(\mathbf{u}_{n-1}) \xrightarrow{\sigma}_{\mathbb{Q}_+} q(\mathbf{u}_n)$. Let $i \in [2, n-1]$, $\gamma_i \stackrel{\text{def}}{=} \frac{n-1-i}{n-1}$ and $\gamma'_i \stackrel{\text{def}}{=} \frac{i}{n-1}$. By Lemma 4.2(a),

$$q(\gamma_i \mathbf{u}_0) \xrightarrow{\gamma_i \sigma}_{\mathbb{Q}_+} q(\gamma_i \mathbf{u}_1) \text{ and } q(\gamma'_i \mathbf{u}_{n-1}) \xrightarrow{\gamma'_i \sigma}_{\mathbb{Q}_+} q(\gamma'_i \mathbf{u}_n).$$

Hence, by Lemma 4.2(b) and $\gamma_i + \gamma'_i = 1$, we obtain

$$q(\gamma_i \mathbf{u}_0 + \gamma'_i \mathbf{u}_{n-1}) \xrightarrow{\sigma}_{\mathbb{Q}_+} q(\gamma_i \mathbf{u}_1 + \gamma'_i \mathbf{u}_n). \quad (9)$$

Now,

$$\begin{aligned}\gamma_i \mathbf{u}_0 + \gamma'_i \mathbf{u}_{n-1} &= \gamma_i \mathbf{u} + \gamma'_i \left(\mathbf{u} + \frac{n-1}{n} \cdot \delta(\pi) \right) \\ &= \mathbf{u} + \gamma'_i \cdot \frac{n-1}{n} \cdot \delta(\pi) \\ &= \mathbf{u} + \frac{i}{n-1} \cdot \frac{n-1}{n} \cdot \delta(\pi) \\ &= \mathbf{u} + \frac{i}{n} \cdot \delta(\pi) \\ &= \mathbf{u}_i.\end{aligned}$$

Similarly, $\gamma_i \mathbf{u}_1 + \gamma'_i \mathbf{u}_n = \mathbf{u}_{i+1}$. Therefore, by (9), we conclude that $q(\mathbf{u}_i) \xrightarrow{\sigma}_{\mathbb{Q}_+} q(\mathbf{u}_{i+1})$, which hence $q(\mathbf{u}) \xrightarrow{\pi'}_{\mathbb{Q}_+} q(\mathbf{v})$. \square

Lemma 4.4: Let $p, q \in \mathcal{Q}$, $\mathbf{u}, \mathbf{v} \in \mathbb{Q}_+^d$ and $\pi \in T^\star$. If $p(\mathbf{u}) \xrightarrow{\pi}_{\mathbb{Q}_+} q(\mathbf{v})$, then there exists $\pi' \in T^\star$ such that (a) π' is q -admissible from $p(\mathbf{u})$,

- (b) $\text{path}(\pi') = \text{path}(\pi)$, and
(c) $\llbracket \delta(t) \rrbracket \subseteq \llbracket \mathbf{v}' \rrbracket$ for every $t \in \llbracket \pi' \rrbracket$, where $\mathbf{v}' \stackrel{\text{def}}{=} \mathbf{u} + \delta(\pi')$.

Proof: Let $\pi = \alpha_1 t_1 \alpha_2 t_2 \cdots \alpha_k t_k$. We have,

$$p_0(\mathbf{u}_0) \xrightarrow{\alpha_1 t_1} \mathbb{Q}_+ p_1(\mathbf{u}_1) \xrightarrow{\alpha_2 t_2} \mathbb{Q}_+ \cdots \xrightarrow{\alpha_k t_k} \mathbb{Q}_+ p_k(\mathbf{u}_k)$$

for some $p_0(\mathbf{u}_0), p_1(\mathbf{u}_1), \dots, p_k(\mathbf{u}_k) \in Q \times \mathbb{Q}_+^d$ such that $p_0(\mathbf{u}_0) = p(\mathbf{u})$ and $p_k(\mathbf{u}_k) = p(\mathbf{v})$.

We claim that $\pi' = \frac{\alpha_1}{2} t_1 \frac{\alpha_2}{4} t_2 \cdots \frac{\alpha_k}{2^k} t_k$ is as required. By definition of π' , (b) holds. To see that (a) and (c) also hold, let us show, by induction on $\ell \in [0, k]$, that

$$p_0(\mathbf{u}_0) \xrightarrow{\pi' [1.. \ell]} \mathbb{Q}_+ p_\ell \left(\frac{1}{2^\ell} \mathbf{u}_\ell + \sum_{j=0}^{\ell-1} \frac{1}{2^{j+1}} \mathbf{u}_j \right). \quad (10)$$

The base case $\ell = 0$ is immediate. Assume (10) holds for some $\ell > 0$. Since $p_\ell(\mathbf{u}_\ell) \xrightarrow{\alpha_\ell t_\ell} \mathbb{Q}_+ p_{\ell+1}(\mathbf{u}_{\ell+1})$, by Lemma 4.2(a) we have

$$p_\ell \left(\frac{1}{2^\ell} \mathbf{u}_\ell \right) \xrightarrow{\frac{1}{2^\ell} \alpha_\ell t_\ell} \mathbb{Q}_+ p_{\ell+1} \left(\frac{1}{2^\ell} \mathbf{u}_{\ell+1} \right).$$

By applying Lemma 4.2(c) with $\alpha = 1/2$, we obtain

$$p_\ell \left(\frac{1}{2^\ell} \mathbf{u}_\ell \right) \xrightarrow{\frac{1}{2^{\ell+1}} \alpha_\ell t_\ell} \mathbb{Q}_+ p_{\ell+1} \left(\frac{1}{2^{\ell+1}} \mathbf{u}_\ell + \frac{1}{2^{\ell+1}} \mathbf{u}_{\ell+1} \right). \quad (11)$$

By adding $\sum_{j=0}^{\ell-1} \frac{1}{2^{j+1}} \mathbf{u}_j$ on both sides of 11, we obtain

$$p_\ell \left(\frac{1}{2^\ell} \mathbf{u}_{\ell+1} + \sum_{j=0}^{\ell-1} \frac{1}{2^{j+1}} \mathbf{u}_j \right) \xrightarrow{\frac{1}{2^{\ell+1}} \alpha_\ell t_\ell} \mathbb{Q}_+ p_{\ell+1} \left(\frac{1}{2^{\ell+1}} \mathbf{u}_\ell + \sum_{j=0}^{\ell} \frac{1}{2^{j+1}} \mathbf{u}_j \right). \quad (12)$$

Therefore, by combining (10) and (12), we conclude that

$$p_0(\mathbf{u}_0) \xrightarrow{\pi' [1.. \ell+1]} \mathbb{Q}_+ p_{\ell+1} \left(\frac{1}{2^{\ell+1}} \mathbf{u}_{\ell+1} + \sum_{j=0}^{\ell} \frac{1}{2^{j+1}} \mathbf{u}_j \right).$$

Now, (a) holds since (10) is a run over \mathbb{Q}_+^d . It remains to prove (c). For the sake of contradiction, suppose there exist $t \in \llbracket \pi' \rrbracket$ and $i \in [d]$ such that $\delta(t)(i) \neq 0$ and $\mathbf{u}(i) + \delta(\pi')(i) = 0$. By (10), this implies that

$$\frac{1}{2^k} \mathbf{u}_k(i) + \sum_{j=0}^{k-1} \frac{1}{2^{j+1}} \mathbf{u}_j(i) = 0.$$

Hence, $\mathbf{u}_0(i) = \mathbf{u}_1(i) = \cdots = \mathbf{u}_k(i) = 0$, and we must have $\delta(t)(i) = 0$. \square

Proposition 4.5: Let $q \in Q$, $\mathbf{u}, \mathbf{v} \in \mathbb{Q}_+^d$ and $\mathbf{w} \in \mathbb{Q}_+^T$. There exists $\pi \in T^\star$ such that $q(\mathbf{u}) \xrightarrow{\pi} \mathbb{Q}_+ q(\mathbf{v})$ and $\Psi(\pi) = \mathbf{w}$ if and only if there exist $\pi, \pi_{\text{fwd}}, \pi_{\text{bwd}} \in T^\star$ such that

- $q(\mathbf{u}) \xrightarrow{\pi} \mathbb{Q}_+ q(\mathbf{v})$,
- π_{fwd} is q -admissible from $q(\mathbf{u})$ in \mathcal{V} ,
- π_{bwd}^\dagger is q -admissible from $q(\mathbf{v})$ in \mathcal{V}^\dagger ,
- $\Psi(\pi) = \mathbf{w}$ and $\llbracket \pi \rrbracket = \llbracket \pi_{\text{fwd}} \rrbracket = \llbracket \pi_{\text{bwd}} \rrbracket = \llbracket \mathbf{w} \rrbracket$.

Proof: The ‘‘only if’’ direction is immediate. Therefore, we assume (a–d) hold, and we show that $q(\mathbf{u}) \xrightarrow{*} \mathbb{Q}_+ q(\mathbf{v})$.

By Lemma 4.4, there exists $\pi'_{\text{fwd}} \in T^\star$ such that

- π'_{fwd} is q -admissible from $q(\mathbf{u})$ in \mathcal{V} ,
- $\text{path}(\pi'_{\text{fwd}}) = \text{path}(\pi_{\text{fwd}})$, and
- $\llbracket \delta(t) \rrbracket \subseteq \llbracket \mathbf{x} \rrbracket$ for every $t \in \llbracket \pi'_{\text{fwd}} \rrbracket$, where $\mathbf{x} \stackrel{\text{def}}{=} \mathbf{u} + \delta(\pi'_{\text{fwd}})$

Note that $q(\mathbf{u}) \xrightarrow{\pi'_{\text{fwd}}} \mathbb{Q}_+ q(\mathbf{x})$.

By Lemma 4.4, there exists $\pi'_{\text{bwd}} \in T^\star$ such that

- $\pi'_{\text{bwd}}^\dagger$ is q -admissible from $q(\mathbf{v})$ in \mathcal{V}^\dagger ,
- $\text{path}(\pi'_{\text{bwd}}^\dagger) = \text{path}(\pi_{\text{bwd}}^\dagger)$, and

- $[\delta(t)] \subseteq [\mathbf{y}]$ for every $t \in [\pi'_{\text{bwd}}]$, where $\mathbf{y} \stackrel{\text{def}}{=} \mathbf{v} + \delta(\pi'_{\text{bwd}})$.

Note that $q(\mathbf{v}) \xrightarrow{\pi'_{\text{bwd}}}_{\mathbb{Q}_+} q(\mathbf{y})$, which implies, together with Lemma 4.2(d), that $q(\mathbf{y}) \xrightarrow{\pi'_{\text{bwd}}}_{\mathbb{Q}_+} q(\mathbf{v})$.

Let $\pi = \alpha_1 t_1 \alpha_2 t_2 \cdots \alpha_k t_k$. For every $j \in [k]$, let

$$\begin{aligned} a_j &\stackrel{\text{def}}{=} \Psi(\pi'_{\text{fwd}})(t_j) / \Psi(\text{path}(\pi))(t_j) \\ b_j &\stackrel{\text{def}}{=} \Psi(\pi'_{\text{bwd}})(t_j) / \Psi(\text{path}(\pi))(t_j). \end{aligned}$$

Let $0 < \lambda \leq 1$ be small enough so that $\alpha_j - \lambda(a_j + b_j) \geq 0$ for every $j \in [k]$. Let $c_j \stackrel{\text{def}}{=} \alpha_j - \lambda(a_j + b_j)$ for every $j \in [k]$, and let $\sigma \stackrel{\text{def}}{=} c_1 t_1 c_2 t_2 \cdots c_k t_k$. By Lemma 4.2(c), we have

$$q(\mathbf{u}) \xrightarrow{\lambda \pi'_{\text{fwd}}}_{\mathbb{Q}_+} q(\mathbf{u} + \lambda \delta(\pi'_{\text{fwd}})).$$

Similarly, by applying Lemma 4.2(d), Lemma 4.2(c) and again Lemma 4.2(d), we obtain

$$q(\mathbf{v} - \lambda \delta(\pi'_{\text{bwd}})) \xrightarrow{\lambda \pi'_{\text{bwd}}}_{\mathbb{Q}_+} q(\mathbf{v}).$$

Moreover,

$$\begin{aligned} \delta(\sigma) &= \sum_{j \in [k]} c_j \cdot \delta(t_j) \\ &= \sum_{j \in [k]} (\alpha_j - \lambda(a_j + b_j)) \cdot \delta(t_j) \\ &= \sum_{j \in [k]} \alpha_j \cdot \delta(t_j) - \lambda \sum_{j \in [k]} (a_j + b_j) \cdot \delta(t_j) \\ &= \delta(\pi) - \lambda \sum_{j \in [k]} (a_j + b_j) \cdot \delta(t_j) \\ &= \delta(\pi) - \lambda \sum_{j \in [k]} \frac{\Psi(\pi'_{\text{fwd}})(t_j) + \Psi(\pi'_{\text{bwd}})(t_j)}{\Psi(\text{path}(\pi))(t_j)} \cdot \delta(t_j) \\ &= \delta(\pi) - \lambda \sum_{t \in [\pi]} (\Psi(\pi'_{\text{fwd}})(t) + \Psi(\pi'_{\text{bwd}})(t)) \cdot \delta(t) \\ &= \delta(\pi) - \lambda \delta(\pi'_{\text{fwd}}) - \lambda \delta(\pi'_{\text{bwd}}) \\ &= \mathbf{v} - \mathbf{u} - \lambda \delta(\pi'_{\text{fwd}}) - \lambda \delta(\pi'_{\text{bwd}}) \\ &= (\mathbf{v} - \lambda \delta(\pi'_{\text{bwd}})) - (\mathbf{u} + \lambda \delta(\pi'_{\text{fwd}})) \\ &= \mathbf{y}' - \mathbf{x}'. \end{aligned}$$

Therefore, $q(\mathbf{x}') \xrightarrow{\sigma}_{\mathbb{Q}} q(\mathbf{y}')$. By definition of \mathbf{x}' and \mathbf{y}' , and since $0 < \lambda \leq 1$, we have $[\mathbf{x}] \subseteq [\mathbf{x}']$ and $[\mathbf{y}] \subseteq [\mathbf{y}']$. Thus,

$$\begin{aligned} [\delta(t)] &\subseteq [\mathbf{x}'] \text{ for every } t \in [\pi'_{\text{fwd}}], \\ [\delta(t)] &\subseteq [\mathbf{y}'] \text{ for every } t \in [\pi'_{\text{bwd}}], \text{ and} \end{aligned}$$

hence σ satisfies all of the conditions of Lemma 4.3. Thus, there exists $\sigma' \in T^\star$ such that $\Psi(\sigma') = \Psi(\sigma)$ and $q(\mathbf{x}') \xrightarrow{\sigma'}_{\mathbb{Q}_+} q(\mathbf{y}')$. We obtain

$$q(\mathbf{u}) \xrightarrow{\lambda \pi'_{\text{fwd}}}_{\mathbb{Q}_+} q(\mathbf{x}') \xrightarrow{\sigma'}_{\mathbb{Q}_+} q(\mathbf{y}') \xrightarrow{\lambda \pi'_{\text{bwd}}}_{\mathbb{Q}_+} q(\mathbf{v}).$$

Therefore, it only remains to show that $\Psi(\pi') = \Psi(\pi)$ where $\pi' \stackrel{\text{def}}{=} \lambda \pi'_{\text{fwd}} \cdot \sigma' \cdot \lambda \pi'_{\text{bwd}}$. Let $S \stackrel{\text{def}}{=} [\mathbf{u}]$. Note that

$$[\pi] = [\pi'] = [\sigma] = [\sigma'] = S.$$

Let $t \in S$, we have

$$\begin{aligned}
\Psi(\sigma')(t) &= \Psi(\sigma) \\
&= \sum_{\substack{j \in [k] \\ t_j = t}} c_j \\
&= \sum_{\substack{j \in [k] \\ t_j = t}} \alpha_j - \lambda(a_j + b_j) \\
&= \sum_{\substack{j \in [k] \\ t_j = t}} \alpha_j - \lambda(\Psi(\pi'_{\text{fwd}})(t_j) + \Psi(\pi'_{\text{bwd}})(t_j)) / \Psi(\text{path}(\pi))(t_j) \\
&= \Psi(\pi)(t) + \lambda \cdot \sum_{\substack{j \in [k] \\ t_j = t}} (\Psi(\pi'_{\text{fwd}})(t) + \Psi(\pi'_{\text{bwd}})(t)) / \Psi(\text{path}(\pi))(t) \\
&= \Psi(\pi)(t) + \lambda \cdot (\Psi(\pi'_{\text{fwd}})(t) + \Psi(\pi'_{\text{bwd}})(t)).
\end{aligned}$$

Now, we are done since

$$\begin{aligned}
\Psi(\pi')(t) &= \Psi(\lambda\pi'_{\text{fwd}} \cdot \sigma' \cdot \lambda\pi'_{\text{bwd}})(t) \\
&= \lambda \cdot \Psi(\pi'_{\text{fwd}})(t) + \Psi(\sigma')(t) + \lambda \cdot \Psi(\pi'_{\text{bwd}})(t) \\
&= \Psi(\sigma')(t) + \lambda(\Psi(\pi'_{\text{fwd}})(t) + \Psi(\pi'_{\text{bwd}})(t)) \\
&= \Psi(\pi)(t) + \lambda \cdot (\Psi(\pi'_{\text{fwd}})(t) + \Psi(\pi'_{\text{bwd}})(t)) - \lambda(\Psi(\pi'_{\text{fwd}})(t) + \Psi(\pi'_{\text{bwd}})(t)) \\
&= \Psi(\pi)(t).
\end{aligned}$$

□

B. Missing proofs from Section IV-B

Lemma 4.6: Let $q \in Q$, $\mathbf{u}, \mathbf{v} \in \mathbb{Q}^d$ and $\mathbf{w} \in \mathbb{Q}_+^T$. There exists $\pi \in T^\star$ such that $q(\mathbf{u}) \xrightarrow{\pi} \mathbb{Q} q(\mathbf{v})$ and $\Psi(\pi) = \mathbf{w}$ if and only if there exist $\mathbf{g} \in \mathbb{Q}_+^T$ such that

- (a) $\mathbf{v} = \mathbf{u} + \sum_{t \in T} \mathbf{w}(t) \cdot \delta(t)$,
- (b) $\llbracket \mathbf{g} \rrbracket = \llbracket \mathbf{w} \rrbracket$ and $\mathbf{g}(t) \geq \mathbf{w}(t)$ for every $t \in S$,
- (c) $G_{\mathcal{V}}[\llbracket \mathbf{g} \rrbracket]$ is strongly connected and contains q , and
- (d) $\sum_{t \in \text{in}(p)} \mathbf{g}(t) = \sum_{t \in \text{out}(p)} \mathbf{g}(t)$ for every $p \in Q$.

Proof: \Rightarrow) Let $\pi = \alpha_1 t_1 \alpha_2 t_2 \cdots \alpha_k t_k$ be such that $q(\mathbf{u}) \xrightarrow{\pi} \mathbb{Q} q(\mathbf{v})$ and $\Psi(\pi) = \mathbf{w}$. Let $\mathbf{g}(t) \stackrel{\text{def}}{=} \Psi(\text{path}(\pi))(t)$ for every $t \in T$. Note that (a) immediately holds, and that (b) holds since $0 < \alpha_i \leq 1$ for every $i \in [k]$. Let $S \stackrel{\text{def}}{=} \llbracket \mathbf{g} \rrbracket$. Since $\text{path}(\pi)$ is a path from q to q , $G_{\mathcal{V}}[S]$ is strongly connected and contains q . It remains to show (d). Let H be the directed multigraph obtained from $G_{\mathcal{V}}[S]$ by replacing each edge $t \in S$ by $\Psi(\text{path}(\pi))(t)$ many arcs from $\text{in}(t)$ to $\text{out}(t)$. It is readily seen that $\text{path}(\pi)$ yields an Eulerian cycle in H . Therefore, $|\text{in}(H, p)| = |\text{out}(H, p)|$ for every $p \in Q$, which implies (d). Indeed, for every $p \in Q$, we have

$$\begin{aligned}
&\sum_{t \in \text{in}(G_{\mathcal{V}}, p)} \mathbf{g}(t) \\
&= \sum_{t \in \text{in}(G_{\mathcal{V}}, p)} \Psi(\text{path}(\pi))(t) \\
&= |\text{in}(H, p)| = |\text{out}(H, p)| \\
&= \sum_{t \in \text{out}(G_{\mathcal{V}}, p)} \Psi(\text{path}(\pi))(t) \\
&= \sum_{t \in \text{out}(G_{\mathcal{V}}, p)} \mathbf{g}(t).
\end{aligned}$$

\Leftrightarrow Let $S \stackrel{\text{def}}{=} \llbracket \mathbf{g} \rrbracket$. For every $t \in S$, let $a_t \in \mathbb{N}$ and $b_t \in \mathbb{N}_{>0}$ be such that $\mathbf{g}(t) = a_t/b_t$. Let $\lambda \stackrel{\text{def}}{=} \text{lcm}\{b_t : t \in S\}$, and let H be the directed multigraph obtained from $G_V[S]$ by replacing each arc $t \in S$ by $\lambda \cdot \mathbf{g}(t)$ arcs from $\text{in}(t)$ to $\text{out}(t)$. By (c), $G_V[S]$ is strongly connected and contains q . Thus, H is strongly connected and contains q . Moreover, for every $p \in Q$,

$$\begin{aligned} \text{in}(H, p) &= \sum_{t \in \text{in}(G_V, p)} \lambda \cdot \mathbf{g}(t) \\ &= \lambda \cdot \sum_{t \in \text{in}(G_V, p)} \mathbf{g}(t) \\ &= \lambda \cdot \sum_{t \in \text{out}(G_V, p)} \mathbf{g}(t) && \text{(by (d))} \\ &= \sum_{t \in \text{out}(G_V, p)} \lambda \cdot \mathbf{g}(t) = \text{out}(H, p). \end{aligned}$$

Therefore, H has an Eulerian cycle σ from q to q . Let $\sigma' = t_1 t_2 \cdots t_k$ be the path of G_V obtained from σ . Let $i \in [k]$, and $\alpha_i \stackrel{\text{def}}{=} \mathbf{w}(t_i) / \Psi(\sigma')(t_i)$. Note that $\Psi(\sigma')(t_i) = \lambda \cdot \mathbf{g}(t_i)$. Moreover, $0 < \mathbf{w}(t_i) \leq \mathbf{g}(t_i)$ by (b). Thus, $0 < \alpha_i \leq 1$. Let $\pi \stackrel{\text{def}}{=} \alpha_1 t_1 \alpha_2 t_2 \cdots \alpha_k t_k$. We have

$$\begin{aligned} \delta(\pi) &= \sum_{i \in [k]} \alpha_i \cdot \delta(t_i) \\ &= \sum_{i \in [k]} (\mathbf{w}(t_i) / \Psi(\sigma')(t_i)) \cdot \delta(t_i) \\ &= \sum_{t \in \llbracket \pi \rrbracket} \Psi(\sigma')(t) \cdot (\mathbf{w}(t) / \Psi(\sigma')(t)) \cdot \delta(t) \\ &= \sum_{t \in \llbracket \pi \rrbracket} \mathbf{w}(t) \cdot \delta(t) \\ &= \mathbf{v} - \mathbf{u} && \text{(by (a)).} \end{aligned}$$

Therefore, we obtain $q(\mathbf{u}) \xrightarrow{\pi} q(\mathbf{v})$. Moreover, for every $t \in S$,

$$\begin{aligned} \Psi(\pi)(t) &= \sum_{\substack{i \in [k] \\ t_i = t}} \alpha_i \\ &= \sum_{\substack{i \in [k] \\ t_i = t}} \mathbf{w}(t_i) / \Psi(\sigma')(t_i) \\ &= \Psi(\sigma')(t) \cdot (\mathbf{w}(t) / \Psi(\sigma')(t)) \\ &= \mathbf{w}(t). \end{aligned} \quad \square$$

Note that in the proof of the Lemma 4.6, λ is defined as $\text{lcm}\{b_t : t \in S\}$ where each b_t is the denominator of $\mathbf{g}(t)$. If we were considering CVASS over \mathbb{R} , then $\mathbf{g}(t)$ could be irrational for some $t \in T$. We argue that this does not matter. First note \mathbf{g} is obtained as a solution of the following system of linear inequalities:

$$(\mathbf{v} = \mathbf{u} + \sum_{t \in T} \mathbf{w}(t) \cdot \delta(t)) \wedge \bigwedge_{t \in T} \mathbf{g}(t) \geq \mathbf{w}(t) > 0 \wedge \bigwedge_{p \in Q} \sum_{t \in \text{in}(p)} \mathbf{g}(t) = \sum_{t \in \text{out}(p)} \mathbf{g}(t).$$

If the system has a real solution, then it also has a rational solution (see, e.g. [28]), which can then be used for the argument.

Lemma B.1: Let $p(\mathbf{u}), q(\mathbf{v}) \in Q \times \mathbb{Q}^d$ and $\pi = \alpha_1 t_1 \alpha_2 t_2 \cdots \alpha_k t_k \in T^\star$ be such that $p(\mathbf{u}) \xrightarrow{\pi} q(\mathbf{v})$. If $\mathbf{u} \geq \mathbf{0}$ and

$$\llbracket \delta(t_i) \rrbracket^- \subseteq \llbracket \mathbf{u} \rrbracket \cup \bigcup_{1 \leq j < i} \llbracket \delta(t_j) \rrbracket^+,$$

for every $i \in [k]$ then there exists $\pi' \in T^\star$ such that $\text{path}(\pi') = \text{path}(\pi)$ and π' is q -admissible from $p(\mathbf{u})$.

Proof: Let $\pi = \alpha_1 t_1 \alpha_2 t_2 \cdots \alpha_k t_k$. We prove the claim by induction on $0 \leq i \leq k$. For $i = 0$, we trivially have $p(\mathbf{u}) \xrightarrow{\varepsilon} p(\mathbf{u})$. Let $0 < i \leq k$ and assume the claim holds for $i - 1$. By induction hypothesis, we have

$$p(\mathbf{u}) \xrightarrow{\alpha'_1 t_1 \alpha'_2 t_2 \cdots \alpha'_{i-1} t_{i-1}} r(\mathbf{v}')$$

for some $\alpha'_1, \alpha'_2, \dots, \alpha'_{i-1} \in (0, 1]$ and $r(\mathbf{v}') \in Q \times \mathbb{Q}_+^d$. By Lemma 4.4, there exist $\beta_1, \beta_2, \dots, \beta_{i-1} \in (0, 1]$ and $\mathbf{w} \in \mathbb{Q}_+^d$ such that

$$p(\mathbf{u}) \xrightarrow{\beta_1 t_1 \beta_2 t_2 \cdots \beta_{i-1} t_{i-1}}_{\mathbb{Q}_+} r(\mathbf{w}), \quad (13)$$

and

$$\bigcup_{1 \leq j < i} \llbracket \delta(t_j) \rrbracket \subseteq \llbracket \mathbf{w} \rrbracket. \quad (14)$$

We argue that $\llbracket \mathbf{u} \rrbracket \subseteq \llbracket \mathbf{w} \rrbracket$. For the sake of contradiction, assume there exists $\ell \in [d]$ such that $\mathbf{u}(\ell) > 0$ and $\mathbf{w}(\ell) = 0$. Altogether with (13), there exists some $1 \leq j < i$ such that $\delta(t_j)(\ell) < 0$. By (14), we have $\ell \in \llbracket \mathbf{w} \rrbracket$ which is a contradiction.

Therefore

$$\llbracket \mathbf{u} \rrbracket \cup \bigcup_{1 \leq j < i} \llbracket \delta(t_j) \rrbracket \subseteq \llbracket \mathbf{w} \rrbracket. \quad (15)$$

Let

$$\beta_i \stackrel{\text{def}}{=} \min(1, \{ \lfloor -\mathbf{w}(j) / \delta(t_i)(j) \rfloor : j \in \llbracket \delta(t_i) \rrbracket^- \}).$$

By hypothesis on π , we have

$$\llbracket \delta(t_i) \rrbracket^- \subseteq \llbracket \mathbf{u} \rrbracket \cup \bigcup_{1 \leq j < i} \llbracket \delta(t_j) \rrbracket^+.$$

By (15), this implies that $\beta_i \neq 0$ and hence that $0 < \beta_i \leq 1$. We claim that $\mathbf{w} + \beta_i \cdot \delta(t_i) \geq \mathbf{0}$, in which case we are done by extending (13) with $\beta_i t_i$.

Let us prove the claim. Let $j \in [d]$. If $j \notin \llbracket \delta(t_i) \rrbracket^-$, then $\mathbf{w}(j) + \beta_i \cdot \delta(t_i)(j) \geq 0$ is immediate. If $j \in \llbracket \delta(t_i) \rrbracket^-$, we have

$$\begin{aligned} \mathbf{w}(j) + \beta_i \cdot \delta(t_i)(j) &\geq \mathbf{w}(j) + \left\lfloor \frac{-\mathbf{w}(j)}{\delta(t_i)(j)} \right\rfloor \cdot \delta(t_i)(j) \\ &\geq \mathbf{w}(j) + \frac{-\mathbf{w}(j)}{\delta(t_i)(j)} \cdot \delta(t_i)(j) \\ &\geq 0. \end{aligned} \quad \square$$

Lemma 4.7: Let $p(\mathbf{u}) \in Q \times \mathbb{Q}_+^d$, $q \in Q$ and $S \subseteq T$. There exists $\pi \in T^\star$ such that

- π is q -admissible from $p(\mathbf{u})$, and
- $\llbracket \pi \rrbracket = S$.

if and only if there exists an injection $f: S \rightarrow \mathbb{Q}_{>0}$ such that for every $t, t' \in S$,

- (a) if $f(t) = \min\{f(s) : s \in S\}$, then $\text{in}(t) = p$,
- (b) if $f(t) < f(t')$, then there exists a path π in G_V from $\text{out}(t)$ to $\text{in}(t')$ such that $\llbracket \pi \rrbracket \subseteq \{s \in S : f(s) < f(t')\}$,
- (c) if $f(t) = \max\{f(s) : s \in S\}$, then there exists a path π in G_V from $\text{out}(t)$ to q such that $\llbracket \pi \rrbracket \subseteq S$, and
- (d) $\llbracket \delta(t) \rrbracket^- \subseteq \llbracket \mathbf{u} \rrbracket \cup \bigcup_{\substack{s \in S \\ f(s) < f(t)}} \llbracket \delta(s) \rrbracket^+$.

Proof: \Rightarrow) Let $\pi \in T^\star$ be such that $\llbracket \pi \rrbracket = S$ and π is q -admissible from $p(\mathbf{u})$. Let $\pi = \alpha_1 t_1 \alpha_2 t_2 \cdots \alpha_k t_k$. We build f from the order in which transitions first occur. More formally, for every $t \in S$, let $i_t \stackrel{\text{def}}{=} \min\{j \in [k] : t_j = t\}$. We define f as follows:

$$f(t) \stackrel{\text{def}}{=} \begin{cases} 1 & \text{if } i_t = 1, \\ \max\{f(s) : s \in S, i_s < i_t\} + 1 & \text{otherwise.} \end{cases}$$

It is readily seen that (a), (b) and (c) hold. We prove (d). Let $t \in S$ and $i \in [d]$ be such that $\delta(t)(i) < 0$. For the sake of contradiction, suppose that $\mathbf{u}(i) = 0$ and $\delta(s)(i) \leq 0$ for every $s \in S$ such that $f(s) < f(t)$. By construction of f , $f(t_j) < f(t_{i_t})$ for every $j < i_t$, hence

$$\begin{aligned} \mathbf{u}(i) + \delta(\alpha_1 t_1 \alpha_2 t_2 \cdots \alpha_{i_t} t_{i_t})(i) &\leq \delta(\alpha_{i_t} t_{i_t})(i) \\ &= \alpha_{i_t} \cdot \delta(t)(i) \\ &< 0, \end{aligned}$$

which contradicts the fact that π is q -admissible from $p(\mathbf{u})$.

\Leftarrow) We order transitions of S with respect to the total order induced by f . More formally, let $\ell_{\min} \stackrel{\text{def}}{=} \min\{f(s) : s \in S\}$, let $k \stackrel{\text{def}}{=} |S|$ and let $g : S \rightarrow [k]$ be such that

$$g(t) \stackrel{\text{def}}{=} \begin{cases} 1 & \text{if } f(t) = \ell_{\min}, \\ \max\{g(s) : s \in S, f(s) < f(t)\} + 1 & \text{otherwise.} \end{cases}$$

Note that g is a bijection since f is injective. We define $t_i \stackrel{\text{def}}{=} g^{-1}(i)$ for every $i \in [k]$. Let σ_0 be the path from p to $\text{in}(t_1)$ given by (a), and let σ_k be the path from $\text{out}(t_k)$ to q given by (c). For every $1 \leq i < k$, let σ_i be the path from $\text{out}(t_i)$ to $\text{in}(t_{i+1})$ given by (b). Let

$$\pi \stackrel{\text{def}}{=} t_1 \sigma_1 t_2 \sigma_2 \cdots t_k \sigma_k.$$

By construction, π is a path from p to q , hence $p(\mathbf{u}) \xrightarrow{\pi} \mathbb{Q} q(\mathbf{v})$ for some $\mathbf{v} \in \mathbb{Q}^d$. Since $\mathbf{u} \geq \mathbf{0}$ and (d) holds, we may apply Lemma B.1. Therefore, there exists $\pi' \in T^\star$ such that $\text{path}(\pi') = \text{path}(\pi)$ and π' is q -admissible from $p(\mathbf{u})$. \square

In order to show that CVASS cyclic reachability can be translated into existential $\text{FO}(\mathbb{Q}, +, <)$, we will need some intermediary formulas for graph reachability, strong connectivity and CVASS cyclic \mathbb{Q} -reachability. Even though defining these formulas in existential $\text{FO}(\mathbb{Q}, +, <)$ would suffice for this purpose, we show that they can be defined in $\text{CSP}(\Gamma_c)$, which is useful for other applications.

Lemma B.2: Let $G = (V, E)$ be a directed graph. There exists a logarithmic-space computable formula $\varphi_G^{\text{paths}}(\mathbf{x}, \mathbf{r})$ of $\text{CSP}(\Gamma_c)$, where $\mathbf{x} \in \mathbb{Q}_+^E$ and $\mathbf{r} \in \mathbb{Q}_+^{V \times V}$, such that φ_G is satisfied by (\mathbf{x}, \mathbf{r}) if and only if $\mathbf{r}_{u,v} > 0 \iff v$ is reachable from u in $G[[\mathbf{x}]]$ for every nodes $u, v \in V$.

Proof: We define the formula as follows:

$$\begin{aligned} \varphi_G^{\text{paths}}(\mathbf{x}, \mathbf{r}) \stackrel{\text{def}}{=} & \exists \mathbf{y} \in \mathbb{Q}_+^{V \times V \times [0, |E| - 1]} \bigwedge_{u \in V} \left(\mathbf{y}_{u,u,0} > 0 \wedge \bigwedge_{v \in V \setminus \{u\}} \mathbf{y}_{u,v,0} = 0 \right) \\ & \wedge \bigwedge_{u,v \in V} \bigwedge_{1 \leq i < |E|} \left(\mathbf{y}_{u,v,i} > 0 \rightarrow \bigvee_{e \in \text{in}(v)} (\mathbf{x}_e > 0 \wedge \mathbf{y}_{u, \text{in}(e), i-1} > 0) \right) \\ & \wedge \bigwedge_{u,v \in V} \mathbf{r}_{u,v} > 0 \iff \bigvee_{0 \leq i < |E|} \mathbf{y}_{u,v,i} > 0. \end{aligned}$$

A simple induction shows that φ_G^{paths} is as required. \square

Lemma B.3: Let $G = (V, E)$ be a directed graph and let $u \in V$. There exists a logarithmic-space computable formula $\varphi_{G,u}^{\text{strong-conn}}(\mathbf{x})$ of $\text{CSP}(\Gamma_c)$, where $\mathbf{y} \in \mathbb{Q}_+^E$, such that $\varphi_{G,u}$ is satisfied by \mathbf{x} if and only if $G[[\mathbf{x}]]$ is strongly connected and contains u .

Proof: Let φ_G^{paths} be the formula obtained from Lemma B.2 for graph reachability. We define the formula as follows, where \mathbf{g}_v represents whether node v occurs in $G[[\mathbf{x}]]$:

$$\begin{aligned} \varphi_{G,u}^{\text{strong-conn}}(\mathbf{x}) \stackrel{\text{def}}{=} & \exists \mathbf{r} \in \mathbb{Q}_+^{V \times V}, \mathbf{g} \in \mathbb{Q}_+^V \\ & \varphi_G^{\text{paths}}(\mathbf{x}, \mathbf{r}) \wedge \bigwedge_{v \in V} \mathbf{g}_v > 0 \iff \bigvee_{e \in \text{edges}(v)} \mathbf{x}_e > 0 \\ & \wedge \mathbf{g}_u > 0 \wedge \bigwedge_{v \in V} \mathbf{g}_v > 0 \rightarrow (\mathbf{r}_{u,v} > 0 \wedge \mathbf{r}_{v,u} > 0). \end{aligned}$$

By Lemma B.2, it follows that $\varphi_{G,u}^{\text{strong-conn}}$ is as required. \square

Lemma B.4: Let $q \in \mathcal{Q}$. There exists a logarithmic-space computable formula $\varphi_{\mathcal{V},q}^{\mathbb{Q}\text{-reach}}(\mathbf{x}, \mathbf{y}, \mathbf{z})$ of $\text{CSP}(\Gamma_c)$, where $\mathbf{x}, \mathbf{y} \in \mathbb{Q}_+^d$ and $\mathbf{z} \in \mathbb{Q}_+^T$, such that $\varphi_{G,u}$ is satisfied by $(\mathbf{u}, \mathbf{v}, \mathbf{w})$ if and only if there exists $\pi \in T^\star$ such that $q(\mathbf{u}) \xrightarrow{\pi} \mathbb{Q} q(\mathbf{v})$ in \mathcal{V} and $\Psi(\pi) = \mathbf{w}$.

Proof: Let $\varphi_{G_{\mathcal{V},q}}^{\text{strong-conn}}$ be the formula obtained from Lemma B.3 for strong connectivity. We build $\varphi_{\mathcal{V},q}^{\mathbb{Q}\text{-reach}}$ from the characterization of Lemma 4.6:

$$\begin{aligned} \varphi_{\mathcal{V},q}^{\mathbb{Q}\text{-reach}}(\mathbf{x}, \mathbf{y}, \mathbf{z}) &\stackrel{\text{def}}{=} \exists \mathbf{g} \in \mathbb{Q}_+^T \\ &\quad \mathbf{x} + \sum_{t \in T} \mathbf{z}_t \cdot \delta(t) = \mathbf{y} \\ &\quad \wedge \bigwedge_{t \in T} (\mathbf{g}_t > 0 \leftrightarrow \mathbf{z}_t > 0) \\ &\quad \wedge \bigwedge_{t \in T} \mathbf{g}_t \geq \mathbf{z}_t \\ &\quad \wedge \Phi_{G_{\mathcal{V},q}}^{\text{strong-conn}}(\mathbf{g}) \\ &\quad \wedge \bigwedge_{p \in Q} \left(\sum_{t \in \text{in}(p)} \mathbf{g}_t = \sum_{t \in \text{out}(p)} \mathbf{g}_t \right). \end{aligned}$$

By Lemma B.3 and Lemma 4.6, the formula is as required. \square

We may now translate CVASS cyclic reachability into existential $\text{FO}(\mathbb{Q}, +, <)$.

Lemma 4.8: There exists a logarithmic-space computable existential $\text{FO}(\mathbb{Q}, +, <)$ formula $\varphi_{\mathcal{V},q}(\mathbf{x}, \mathbf{y}, \mathbf{z})$ such that $\varphi_{\mathcal{V},q}(\mathbf{u}, \mathbf{v}, \mathbf{w})$ holds if and only if there exists $\pi \in T^\star$ such that $q(\mathbf{u}) \xrightarrow{\pi}_{\mathbb{Q}_+} q(\mathbf{v})$ and $\Psi(\pi) = \mathbf{w}$.

Proof: We build a formula for the characterization obtained in Proposition 4.5. First, we build a formula $\varphi_{\mathcal{V},q}^{\text{adm}}$ for Proposition 4.5(b), i.e. for the existence of some π such that π is q -admissible from $q(\mathbf{u})$ and $\llbracket \pi \rrbracket = \llbracket \mathbf{z} \rrbracket$.

Let φ_G^{paths} be the formula for graph reachability obtained from Lemma B.2. Formula $\varphi_{\mathcal{V},q}^{\text{adm}}$ is derived from the characterization of obtained in Lemma 4.7:

$$\begin{aligned} \varphi_{\mathcal{V},q}^{\text{adm}}(\mathbf{x}, \mathbf{z}) &\stackrel{\text{def}}{=} \exists \mathbf{f} \in \mathbb{Q}_+^T \\ &\quad \bigwedge_{t \in T} (\mathbf{f}_t > 0 \leftrightarrow \mathbf{z}_t > 0) \\ &\quad \wedge \bigwedge_{\substack{t, t' \in T \\ t \neq t'}} [(\mathbf{f}_t > 0 \wedge \mathbf{f}_{t'} > 0) \rightarrow \mathbf{f}_t \neq \mathbf{f}_{t'}] \\ &\quad \wedge \bigwedge_{t \in T} \varphi_t^{\text{min}}(\mathbf{f}) \rightarrow \text{in}(t) = q \\ &\quad \wedge \bigwedge_{t, t' \in T} (0 < \mathbf{f}_t < \mathbf{f}_{t'}) \rightarrow \exists \mathbf{y} \in \mathbb{Q}_+^T, \mathbf{r} \in \mathbb{Q}_+^{\mathbb{Q} \times \mathbb{Q}} \bigwedge_{s \in T} [\mathbf{y}_s > 0 \leftrightarrow (0 < \mathbf{f}_s < \mathbf{f}_{t'})] \wedge \varphi_{G_{\mathcal{V}}}^{\text{paths}}(\mathbf{y}, \mathbf{r}) \wedge \mathbf{r}_{\text{out}(t), \text{in}(t')} > 0 \\ &\quad \wedge \bigwedge_{t \in T} \varphi_t^{\text{max}}(\mathbf{f}) \rightarrow \exists \mathbf{r} \in \mathbb{Q}_+^{\mathbb{Q} \times \mathbb{Q}} \varphi_{G_{\mathcal{V}}}^{\text{paths}}(\mathbf{f}, \mathbf{r}) \wedge \mathbf{r}_{\text{out}(t), q} > 0 \\ &\quad \wedge \bigwedge_{t \in T} \mathbf{f}_t > 0 \rightarrow \left(\bigwedge_{i \in [d]} \delta(t)(i) < 0 \rightarrow \left[\mathbf{x}_i > 0 \vee \bigvee_{s \in T} (0 < \mathbf{f}_s < \mathbf{f}_t \wedge \delta(t)(i) > 0) \right] \right) \end{aligned}$$

where

$$\begin{aligned} \varphi_t^{\text{min}}(\mathbf{f}) &\stackrel{\text{def}}{=} \mathbf{f}_t > 0 \wedge \bigwedge_{t' \in T \setminus \{t\}} (\mathbf{f}_{t'} > 0 \rightarrow \mathbf{f}_t < \mathbf{f}_{t'}) \\ \varphi_t^{\text{max}}(\mathbf{f}) &\stackrel{\text{def}}{=} \bigwedge_{t' \in T \setminus \{t\}} \mathbf{f}_t > \mathbf{f}_{t'}. \end{aligned}$$

Let $\varphi_{\mathcal{V},q}^{\mathbb{Q}\text{-reach}}$ be the formula for CVASS cyclic \mathbb{Q} -reachability obtained from Lemma B.4. Altogether, we obtain the following formula whose validity follows from Proposition 4.5, Lemma 4.6 and Lemma 4.7:

$$\varphi_{\mathcal{V},q}(\mathbf{x}, \mathbf{y}, \mathbf{z}) \stackrel{\text{def}}{=} \exists \varphi_{\mathcal{V},q}^{\mathbb{Q}\text{-reach}}(\mathbf{x}, \mathbf{y}, \mathbf{z}) \wedge \varphi_{\mathcal{V},q}^{\text{adm}}(\mathbf{x}, \mathbf{z}) \wedge \varphi_{\mathcal{V},q}^{\text{adm}}(\mathbf{y}, \mathbf{z}). \quad \square$$

We may now translate CVASS reachability into existential $\text{FO}(\mathbb{Q}, +, <)$:

Theorem 4.9: Let $p, q \in Q$. There exists a logarithmic-space computable existential $\text{FO}(\mathbb{Q}, +, <)$ formula $\varphi_{\mathcal{V}, p, q}(\mathbf{x}, \mathbf{y}, \mathbf{z})$ such that $\varphi_{\mathcal{V}, p, q}(\mathbf{u}, \mathbf{v}, \mathbf{w})$ holds if and only if there exists $\pi \in T^\star$ such that $p(\mathbf{u}) \xrightarrow{\pi}_{\mathbb{Q}_+} q(\mathbf{v})$ and $\Psi(\pi) = \mathbf{w}$.

Proof: We first give a formula $\varphi_{p, q, \ell}^{\text{run}}(\mathbf{x}, \mathbf{y}, \mathbf{z})$ for “length ℓ reachability”, i.e. for the existence of some $\pi \in T^\star$ such that $|\pi| = \ell$ and $p(\mathbf{x}) \xrightarrow{\pi}_{\mathbb{Q}_+} q(\mathbf{y})$. If $\ell = 0$, we simply build

$$\varphi_{p, q, 0}^{\text{run}}(\mathbf{x}, \mathbf{y}, \mathbf{z}) \stackrel{\text{def}}{=} (p = q) \wedge \bigwedge_{i \in [d]} \mathbf{x}_i = \mathbf{y}_i.$$

If $\ell > 0$, we build a formula that guesses a path with its $(0, 1]$ -coefficients, checks whether counter values do not drop below zero, and whether \mathbf{y} is reached:

$$\begin{aligned} \varphi_{p, q, \ell}^{\text{run}}(\mathbf{x}, \mathbf{y}, \mathbf{z}) \stackrel{\text{def}}{=} & \exists \mathbf{a}_1, \mathbf{a}_2, \dots, \mathbf{a}_\ell \in \mathbb{Q}_+^T, \mathbf{p}_1, \mathbf{p}_2, \dots, \mathbf{p}_\ell \in \mathbb{Q}^d \\ & \bigvee_{t \in \text{out}(p)} 0 < \mathbf{a}_{1, t} \leq 1 \\ & \wedge \bigwedge_{1 \leq i < \ell} \bigwedge_{t \in T} \mathbf{a}_{i, t} > 0 \rightarrow \bigvee_{\substack{t' \in T \\ \text{out}(t) = \text{in}(t')}} 0 < \mathbf{a}_{i+1, t'} \leq 1 \\ & \wedge \bigvee_{t \in \text{in}(q)} 0 < \mathbf{a}_{\ell, t} \leq 1 \\ & \wedge \bigwedge_{i \in [\ell]} \bigwedge_{t \in T} (\mathbf{a}_{i, t} > 0 \rightarrow \bigwedge_{t' \in T \setminus \{t\}} \mathbf{a}_{i, t'} = 0) \\ & \wedge \bigwedge_{i \in [\ell]} \bigwedge_{t \in T} \left(\mathbf{a}_{i, t} > 0 \rightarrow \bigwedge_{j \in [d]} \mathbf{p}_{i, j} = \mathbf{a}_{i, t} \cdot \delta(t)(j) \right) \\ & \wedge \bigwedge_{\ell \in [\ell]} \bigwedge_{j \in [d]} \mathbf{x}(j) + \sum_{1 \leq i \leq \ell} \mathbf{p}_{i, j} \geq 0 \\ & \wedge \bigwedge_{j \in [d]} \mathbf{x}(j) + \sum_{1 \leq i \leq \ell} \mathbf{p}_{i, j} = \mathbf{y} \\ & \wedge \bigwedge_{t \in T} \mathbf{z}_t = \sum_{i \in [\ell]} \mathbf{a}_{i, t} \end{aligned}$$

For every $r \in Q$, let $\varphi_{\mathcal{V}, r}$ be the formula for cyclic reachability obtained from Lemma 4.8. We give a formula that checks whether (a–d) from the main text holds for some fixed k :

$$\begin{aligned} \varphi_{p, q, k}^{\text{decomp}}(\mathbf{x}, \mathbf{y}, \mathbf{z}) \stackrel{\text{def}}{=} & \exists \mathbf{u}_0, \mathbf{u}_1, \dots, \mathbf{u}_{k+1}, \mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_{k+1} \in \mathbb{Q}_+^d \exists \mathbf{q}_0, \mathbf{q}_1, \dots, \mathbf{q}_{k+1} \in \mathbb{Q}_+^Q, \\ & \exists \mathbf{s}_0, \mathbf{s}_1, \dots, \mathbf{s}_k, \mathbf{t}_1, \mathbf{t}_2, \dots, \mathbf{t}_k \in \mathbb{Q}_+^T \\ & \bigwedge_{1 \leq i \leq k} \left(\bigvee_{q \in Q} \mathbf{q}_{i, q} = 1 \wedge \sum_{q \in Q} \mathbf{q}_{i, q} = 1 \right) \\ & \wedge \mathbf{q}_{0, p} = 1 \wedge \mathbf{q}_{k+1, q} = 1 \\ & \wedge \bigwedge_{j \in [d]} \mathbf{u}_0(j) = \mathbf{x}(j) \wedge \mathbf{v}_{k+1}(j) = \mathbf{y}(j) \\ & \wedge \bigwedge_{0 \leq i \leq k} \bigwedge_{r, r' \in Q} (x_{i, r} = 1 \wedge x_{i+1, r'} = 1) \rightarrow \bigvee_{0 \leq \ell < |Q|} \varphi_{r, r', \ell}^{\text{run}}(\mathbf{u}_i, \mathbf{v}_{i+1}, \mathbf{s}_i) \\ & \wedge \bigwedge_{0 < i \leq k} \bigwedge_{r \in Q} x_{i, r} = 1 \rightarrow \varphi_{\mathcal{V}, r}(\mathbf{v}_i, \mathbf{u}_i, \mathbf{t}_i) \\ & \wedge \bigwedge_{t \in T} \mathbf{z}_t = \sum_{0 \leq i \leq k} \mathbf{s}_{i, t} + \sum_{0 < i \leq k} \mathbf{t}_{i, t}. \end{aligned}$$

Altogether, we obtain:

$$\varphi_{\mathcal{V}, p, q}(\mathbf{x}, \mathbf{y}, \mathbf{z}) \stackrel{\text{def}}{=} \bigvee_{0 \leq k \leq |Q|} \varphi_{p, q, k}^{\text{decomp}}(\mathbf{x}, \mathbf{y}, \mathbf{z}). \quad \square$$

C. Missing proofs from Section IV-C

For every d -CVASS $\mathcal{V} = (Q, T)$, let $\|T\| \stackrel{\text{def}}{=} \max\{|\mathbf{z}(i)| : (p, \mathbf{z}, q) \in T, i \in [d]\}$. For every formula ψ , let $\|\psi\|$ denote the largest absolute value of a coefficient occurring in ψ . We give a detailed version of Lemma 4.11:

Lemma B.5: Let $\psi(\mathbf{x})$ be an atomic formula and $n = |\mathbf{x}|$. There exists a logarithmic-space computable $(2n+3)$ -CVASS $\mathcal{V} = (Q, T)$ and states $p, r \in Q$ such that

- $|Q| = n + 2$, $|T| = 3(n + 1)$, $\|T\| \leq \max(1, \|\psi\|^{n+2})$, and
- $\psi(\mathbf{u})$ holds if and only if $p(\text{enc}(\mathbf{u}), 0, 0, 1) \xrightarrow{*}_{\mathbb{Q}_+} r(\mathbf{0})$.

Proof: Let $\psi(\mathbf{x}) = \mathbf{a}^\top \cdot \mathbf{x} \sim c$. We can assume that coefficients appearing in \mathbf{a} and c have integer values, since $(a_1/d_1)x_1 + \dots + (a_n/d_n)x_n \sim c$ has the same solutions as $\lambda(a_1/d_1)x_1 + \dots + \lambda(a_n/d_n)x_n \sim \lambda \cdot c$ where $\lambda \stackrel{\text{def}}{=} d_1 \dots d_n$. Note that λ can be computed in logarithmic space.

We give a construction which is very similar to the one given in the proof of Lemma 3.7. Let \mathcal{V} be the CVASS illustrated in Figure 7. The first $2 \cdot n$ counters encode \mathbf{x} and the two next counters encode the “current value” $\text{val}(\mathcal{V})$. States p_1 to p_n compute $\mathbf{a}^\top \cdot \mathbf{x}$. The transition from p_n to q subtracts c from $\text{val}(\mathcal{V})$, using the last counter to ensure that c is fully subtracted. The self-loop at q allows to decrease $\text{val}(\mathcal{V})$ by any non-negative value. The transition from q to r forces $\text{val}(\mathcal{V})$ to be decreased by a positive value if \sim is strict. Once $r(\mathbf{0}, y, z, 0)$ is reached, $\text{val}(\mathcal{V}) = y - z = \mathbf{a}^\top \cdot \mathbf{x} - c - \lambda$ for some $\lambda \geq 0$ such that $\lambda > 0$ if \sim is strict. The self-loop on r “normalizes” the encoding of $\text{val}(\mathcal{V})$ to test whether $\text{val}(\mathcal{V}) = 0$, and equivalently whether $\mathbf{a}^\top \cdot \mathbf{x} = c + \lambda$.

More formally, let $b \in \{0, 1\}$ be such that $b = 1$ if \sim is “ $>$ ” and $b = 0$ if \sim is “ \geq ”. We define $\mathcal{V} \stackrel{\text{def}}{=} (Q, T)$ as the continuous $(2n+3)$ -VASS such that $Q \stackrel{\text{def}}{=} \{p_i : 1 \leq i \leq n\} \cup \{q, r\}$ and

$$\begin{aligned} T \stackrel{\text{def}}{=} & \{(p_i, (-\mathbf{e}_{2i-1}, \text{enc}(a_i), 0), p_i) : 1 \leq i \leq n\} \\ & \cup \{(p_i, (-\mathbf{e}_{2i}, \text{enc}(-a_i), 0), p_i) : 1 \leq i \leq n\} \\ & \cup \{(p_i, (\mathbf{0}, \text{enc}(0), 0), p_{i+1}) : 1 \leq i < n\} \\ & \cup \{(p_n, (\mathbf{0}, \text{enc}(-c), -1), q)\} \\ & \cup \{(q, (\mathbf{0}, \text{enc}(-1), 0), q)\} \\ & \cup \{(q, (\mathbf{0}, \text{enc}(-b), 0), r)\} \\ & \cup \{(r, (\mathbf{0}, -1, -1, 0), r)\}. \end{aligned}$$

Since we (might) have taken the product of the denominators of ψ , we obtain $\|T\| \leq \max(1, \|\psi\|^{n+2})$. Note that the binary encoding of T is still of polynomial size. \square




Fig. 7. Continuous VASS for $\psi(\mathbf{x}) = \mathbf{a}^\top \cdot \mathbf{x} > c$. Colours highlight similarities with the continuous Petri net of Figure 1.

We give a detailed version of Theorem 4.12:

Theorem B.6: Let $\varphi(\mathbf{x})$ be a formula of existential $\text{FO}(\mathbb{Q}, +, <)$ and $n = |\mathbf{x}|$. There exist $k > 0$, a $(n+k)$ -CVASS $\mathcal{V} = (Q, T)$, $p, q \in Q$ and $\mathbf{v} \in \{0, 1\}^k$, computable in logarithmic space, such that

- $|Q|, |T|, k \in O(|\varphi|^2)$, $\|T\| \leq \max(1, \|\varphi\|^{n+2})$, and
- $\varphi(\mathbf{u})$ holds if and only if $p(\text{enc}(\mathbf{u}), \mathbf{v}) \xrightarrow{*}_{\mathbb{Q}_+} q(\mathbf{0})$.

Proof: We may assume that $\varphi(\mathbf{x}) = \exists \mathbf{y} \in \mathbb{Q}^{n'} \Phi(\mathbf{x}, \mathbf{y})$ for some $n' \in \mathbb{N}$ and some quantifier-free formula Φ made of conjunctions, disjunctions and atomic formulas over \mathbf{x} and \mathbf{y} . Indeed, if this is not the case, then quantified variables can be quantified at the beginning of φ , and negations can be removed in polynomial time by successively applying De Morgan’s law and the rules

$$\begin{aligned} \neg(a_1x_1 + \dots + a_nx_n \geq c) &= -a_1x_1 - \dots - a_nx_n > -c, \\ \neg(a_1x_1 + \dots + a_nx_n > c) &= -a_1x_1 - \dots - a_nx_n \geq -c. \end{aligned}$$

Let $m \stackrel{\text{def}}{=} n + n'$ and $z \stackrel{\text{def}}{=} [\mathbf{x}, \mathbf{y}]$. The CVASS for $\Phi(z)$ is constructed by recursively composing the $(2m+3)$ -CVASS gadgets obtained from Lemma B.5 for atomic formulas. Let ψ and ψ' be quantifier-free formulas over \mathbf{x} . Suppose that for both $f \in \{\psi, \psi'\}$, we have a $(2m+k_f)$ -CVASS $\mathcal{V}_f = (Q_f, T_f)$, some states $p_f, q_f \in Q_f$ and a vector $\mathbf{v}_f \in \{0, 1\}^{k_f}$ such that

$$f(\mathbf{x}) \text{ holds} \iff (\text{enc}(\mathbf{x}), \mathbf{v}_f) \xrightarrow{*}_{\mathbb{Q}_+} (\mathbf{0}, \mathbf{0}).$$

Let $\circ \in \{\vee, \wedge\}$. We can compose \mathcal{V}_ψ and $\mathcal{V}_{\psi'}$ to obtain a $(2m+k_{\psi \circ \psi'})$ -CVASS $\mathcal{V}_{\psi \circ \psi'}$ for $\psi(\mathbf{x}) \circ \psi'(\mathbf{x})$ as illustrated in Figure 8.




Fig. 8. Gadgets for $\psi(\mathbf{x}) \vee \psi'(\mathbf{x})$ and $\psi(\mathbf{x}) \wedge \psi'(\mathbf{x})$. The CVASS for ψ and ψ' , and their counters, are respectively coloured in ■ and ■.

More precisely, $\mathcal{V}_{\psi \vee \psi'}$ starts with $[\text{enc}(\mathbf{x}), \mathbf{0}, \mathbf{0}, 1]$ and non deterministically simulates either \mathcal{V}_ψ on $[\text{enc}(\mathbf{x}), \mathbf{v}_\psi]$, or $\mathcal{V}_{\psi'}$ on $[\text{enc}(\mathbf{x}), \mathbf{v}_{\psi'}]$. For conjunction, $\mathcal{V}_{\psi \wedge \psi'}$ starts with $[\text{enc}(z), \mathbf{0}, \mathbf{0}, 1, \mathbf{0}, \mathbf{0}, 1]$. The initial state consumes $\text{enc}(\mathbf{x})$ and duplicates it twice. Then, $\mathcal{V}_{\psi \wedge \psi'}$ simulates sequentially \mathcal{V}_ψ on $[\text{enc}(z), \mathbf{v}_\psi]$, and $\mathcal{V}_{\psi'}$ on $[\text{enc}(z), \mathbf{v}_{\psi'}]$.

Let V_Φ be the $(2m+k_\Phi)$ -CVASS obtained for Φ by performing the above compositions recursively. To obtain a CVASS \mathcal{V}_φ for φ , it remains to simulate “ $\exists \mathbf{y}$ ”. To do so, \mathcal{V}_φ starts with $[\mathbf{x}, \mathbf{0}, \mathbf{v}_\Phi]$, guesses \mathbf{y} and moves to the initial state of V_Φ as illustrated in Figure 9. To see that \mathcal{V}_φ is of the required size, observe that disjunction and conjunction increase the number of




Fig. 9. CVASS gadget for $\exists \mathbf{y} \Phi(\mathbf{x}, \mathbf{y})$.

states, transitions and counters linearly. More precisely,

$$\begin{aligned} |Q_{\psi \circ \psi'}| &= |Q_\psi| + |Q_{\psi'}| + 2 \text{ for both } \circ \in \{\vee, \wedge\}, \\ |T_{\psi \vee \psi'}| &= |T_\psi| + |T_{\psi'}| + 4, \\ |T_{\psi \wedge \psi'}| &= |T_\psi| + |T_{\psi'}| + 2m + 3, \\ k_{\psi \vee \psi'} &= k_\psi + k_{\psi'} + 1, \text{ and} \\ k_{\psi \wedge \psi'} &= k_\psi + k_{\psi'} + 4m + 2. \end{aligned}$$

Altogether with the bounds of Lemma B.5, we obtain

$$\begin{aligned} |Q_\varphi| &\leq (2m + 2) \cdot (|\varphi| + 1), \\ |T_\varphi| &\leq (3m + 4) \cdot (|\varphi| + 1), \\ k_\varphi &\leq (4m + 3) \cdot |\varphi|. \end{aligned}$$

Finally, note that each number occurring in \mathcal{V}_φ is either 0, -1 , 1 or a coefficient from the gadgets obtained from Lemma B.5. \square

D. Missing proofs from Section IV-D

We detail the NP-hardness reduction of Theorem 4.16 in the next lemma.

Lemma B.7: Coverability, q -admissibility and unboundedness in CVASS are NP-hard.

Proof: We give a single reduction from 3-SAT to the three problems. Let φ be an instance of 3-SAT. As explain in the proof of Lemma 4.13(b), we may build a CVASS $\mathcal{V} = (Q, T)$ and states $p, q \in Q$ such that $G_{\mathcal{V}}$ is acyclic, $\text{out}(q) = \emptyset$, and φ is satisfiable if and only if $p(\mathbf{0})$ is q -admissible. By adding a self-loop labeled by $\mathbf{1}$ on q , we obtain:

$$\begin{aligned} \varphi \text{ is satisfiable} &\iff p(\mathbf{0}) \text{ is } q\text{-admissible} \\ &\iff q(\mathbf{0}) \text{ is coverable from } p(\mathbf{0}) \\ &\iff \mathcal{V} \text{ is unbounded from } p(\mathbf{0}). \end{aligned} \quad \square$$

We now detail the proof sketch of Theorem 4.17.

Theorem 4.17: Reachability in CVASS with zero-tests is undecidable.

Proof: Let $\mathcal{V} = (Q, T)$ be a d -VASS with zero-tests on k counters. We show how to simulate \mathcal{V} by a $(d+2)$ -CVASS $\mathcal{V}' = (Q', T')$ with zero-tests on $k+2$ counters. First, we add two new counters to \mathcal{V} to ensure that each transition of \mathcal{V} is fully fired in \mathcal{V}' , i.e. by $\alpha = 1$ and not by some $0 < \alpha < 1$. This can be achieved by replacing every transition of \mathcal{V} with the gadget illustrated in Figure 10. We then add a transition from q to a new “trap state” r which allows to decrement the two




Fig. 10. Gadget to transform a VASS into a CVASS with zero-tests with the same reachability set.

new counters. Formally, we define \mathcal{V}' as follows:

$$\begin{aligned} Q' &\stackrel{\text{def}}{=} Q \cup \{q_t, q'_t : t \in T\} \cup \{r\}, \\ T' &\stackrel{\text{def}}{=} \{(p, \text{zero}_{i+2}?, q) : (p, \text{zero}_i?, q) \in T\} \cup \\ &\quad \{(p, (-1, 1, z), q_t), (q_t, \text{zero}_1?, q) : t = (p, z, q) \in T\} \cup \\ &\quad \{(p, (1, -1, z), q'_t), (q'_t, \text{zero}_2?, q) : t = (p, z, q) \in T\} \cup \\ &\quad \{(q, (0, 0, \mathbf{0}), r), (r, (-1, 0, \mathbf{0}), r), (r, (0, -1, \mathbf{0}), r)\}. \end{aligned}$$

It can be shown that $p(\mathbf{u}) \xrightarrow{*}_{\mathbb{N}} q(\mathbf{v})$ in \mathcal{V} if and only if $p(1, 0, \mathbf{u}) \xrightarrow{*}_{\mathbb{Q}_+} r(0, 0, \mathbf{v})$ in \mathcal{V}' . Since reachability is undecidable for 2-VASS with zero-tests on both counters [25], we conclude that it is also undecidable for CVASS with zero-tests on (at least) four counters. \square