

The Complexity of Synthesis of b -Bounded Petri Nets

Ronny Tredup*

Institut für Informatik, Theoretische Informatik

Universität Rostock

Albert-Einstein-Straße 22, 18059, Rostock, Germany

ronny.tredup@uni-rostock.de

Abstract. For a fixed type of Petri nets τ , τ -SYNTHESIS is the task of finding for a given transition system A a Petri net N of type τ (τ -net, for short) whose reachability graph is isomorphic to A if there is one. The decision version of this search problem is called τ -SOLVABILITY. If an input A allows a positive decision, then it is called τ -solvable and a sought net N τ -solves A . As a well known fact, A is τ -solvable if and only if it has the so-called τ -event state separation property (τ -ESSP, for short) and the τ -state separation property (τ -SSP, for short). The question whether A has the τ -ESSP or the τ -SSP defines also decision problems. In this paper, for all $b \in \mathbb{N}$, we completely characterize the computational complexity of τ -SOLVABILITY, τ -ESSP and τ -SSP for the types of pure b -bounded Place/Transition-nets, the b -bounded Place/Transition-nets and their corresponding \mathbb{Z}_{b+1} -extensions.

Keywords: Petri nets, synthesis, b -bounded, SSP, ESSP, solvability

1. Introduction

The task of system *analysis* is to examine the behavior of a system and to derive its behavioral properties. Its counterpart, *synthesis*, is the task of automatically finding an implementing system for a given behavioral specification. A valid synthesis procedure then computes a system that is correct by design if it exists.

*Address for correspondence: Universität Rostock, Institut für Informatik, Theoretische Informatik, Albert-Einstein-Straße 22, 18059, Rostock, Germany

In this paper we investigate a certain instance of synthesis: For a fixed type of Petri nets τ , τ -SYNTHESIS is the task to find, for a given directed labeled graph A , called transition system (TS, for short), a Petri net N of type τ (τ -net, for short) whose state graph is isomorphic to A if such a net exists. The decision version of τ -SYNTHESIS is called τ -SOLVABILITY.

Synthesis for Petri nets has been investigated and applied for many years and in various fields: It is used to extract concurrency and distributability data from sequential specifications like transition systems or languages [1]. Synthesis has applications in the field of process discovery to reconstruct a model from its execution traces [2]. In [3], it is employed in supervisory control for discrete event systems. It is also used for the synthesis of speed-independent circuits [4]. In this paper, we investigate the computational complexity of synthesis for certain types of *bounded* Petri nets, that is, Petri nets for which there is a positive integer b that restricts the number of tokens on every place in every reachable marking.

In [5, 6], synthesis has been shown to be solvable in polynomial time for bounded and pure bounded *Place/Transition*-nets (P/T-nets, for short). The approach provided in [5, 6] guarantees a (pure) bounded P/T-net to be output if such a net exists. Unfortunately, it does not work for preselected bounds. In fact, in [7] it has been shown that solvability is NP-complete for 1-bounded P/T-nets (there referred to as *elementary net systems*), that is, if the bound $b = 1$ is chosen in advance. In [8], the type of pure 1-bounded P/T-nets is extended by the additive group \mathbb{Z}_2 of integers modulo 2 (there referred to as *flip-flop nets*). Transitions of these nets can simulate the addition of integers modulo 2. The result of [8] shows that this suffices to bring the complexity of synthesis down to polynomial time. In [9, 10], we progressed the approach of examining the effects of the presence and absence of different interactions on the complexity of synthesis for the broader class of *Boolean* Petri nets that enable independence between places and transitions. This class also contains the type of 1-bounded P/T-nets and its \mathbb{Z}_2 -extension. Although [9, 10] show that synthesis remains hard for 75 of the 128 possible Boolean types (allowing independence), [9] also discovers 36 types for which synthesis is doable in polynomial time. The latter applies in particular for the \mathbb{Z}_2 extension of 1-bounded P/T-nets. As another aspect that possibly might influence the complexity of synthesis of (pure) 1-bounded P/T-nets, the grade g of a TS A as has been introduced in [11]: A TS A is g -grade if every state of A has at most g incoming and at most g outgoing labeled edges. There we showed that synthesis of pure 1-bounded P/T-nets remains NP-complete even for acyclic 1-grade TS. In [12], for any fixed $g \in \mathbb{N}$, we completely characterize the computational complexity of synthesis from g -grade TS for all Boolean Petri net types that enable independence. Surprisingly enough, for many other Boolean types, synthesis remains hard for all $g \geq 1$. For example, this applies to the type of *inhibitor nets* and the type of *contextual nets*, which have originally been introduced in [13] and [14] and are referred to as $\{\text{nop, inp, out, free}\}$ and $\{\text{nop, inp, out, used, free}\}$ in [12], respectively. However, there are several types for which the complexity changes when g becomes small enough. This applies in particular to the Boolean type of *trace nets* that has originally been introduced in [15] and is referred to as $\{\text{nop, inp, out, res, set, used, free}\}$ in [12]. Synthesis for this type is hard if $g \geq 2$, but polynomial for $g < 2$. The same is true for the type of *set nets* that has originally been introduced in [16] and is referred to as $\{\text{nop, inp, set, used}\}$ in [12].

However, some questions in the area of synthesis for Petri nets are still open. Recently, the complexity status of synthesis for (pure) b -bounded P/T-nets, where $b \geq 2$, has been reported as un-

known [17]. Furthermore, it has not yet been analyzed whether extending (pure) b -bounded P/T-nets by the group \mathbb{Z}_{b+1} provides also a tractable superclass if $b \geq 2$.

Let $b \in \mathbb{N}^+$. In this paper, we show that solvability for (pure) b -bounded P/T-nets is NP-complete even if the input is an acyclic 1-grade TS. Moreover, for $b \geq 2$, we introduce (pure) \mathbb{Z}_{b+1} -extended b -bounded P/T-nets. This type originates from (pure) b -bounded P/T-nets by adding interactions between places and transitions simulating the addition of integers modulo $b + 1$. This extension is a natural generalization of Schmitt's approach that does this for $b = 1$ [8]. In contrast to Schmitt's result [8], in this paper, we show that solvability for (pure) \mathbb{Z}_{b+1} -extended b -bounded P/T-nets remains NP-complete for all $b \geq 2$ even if the input is restricted to g -grade TS where $g \geq 2$. In particular, this makes the synthesis of all of these b -bounded Petri net types NP-hard. The question arises whether there are also types of b -bounded P/T-nets for which synthesis is tractable if $b \geq 2$. We affirm this question and propose the type of restricted \mathbb{Z}_{b+1} -extended b -bounded P/T-nets. This paper shows, that synthesis is solvable in polynomial time for this type.

To prove the NP-completeness of solvability we use its well known close connection to the so-called *event state separation property* (ESSP, for short) and *state separation property* (SSP, for short). In fact, a TS A is solvable with respect to a Petri net type if and only if it has the type related ESSP and SSP [6]. The question of whether a TS A has the ESSP or the SSP also defines decision problems. The possibility to efficiently decide if A has at least one of both properties serves as quick-fail pre-processing mechanisms for solvability. Moreover, if A has the ESSP then synthesizing Petri nets up to language equivalence is possible [6]. This makes the decision problems ESSP and SSP worth to study. In [18], both problems have been shown to be NP-complete for pure 1-bounded P/T-nets. This has been confirmed for almost trivial inputs in [11, 19].

In this paper, for all $b \in \mathbb{N}^+$, we show that ESSP and SSP are NP-complete for (pure) b -bounded P/T-nets even if the input is an acyclic 1-grade TS. Moreover, for all $b \geq 2$, the ESSP is shown to remain NP-complete for (pure) \mathbb{Z}_{b+1} -extended b -bounded P/T-nets for g -grade TS where $g \geq 2$. By way of contrast, in this paper, we show that SSP is decidable in polynomial time for the type of (pure) \mathbb{Z}_{b+1} -extended b -bounded P/T-nets, for all $b \in \mathbb{N}$. To the best of our knowledge, so far, this is the first net family where the provable computational complexity of SSP is different to solvability and ESSP.

All presented NP-completeness proofs base on a reduction from the monotone one-in-three 3-SAT problem that is known to be NP-complete [20]. Every reduction starts from a given boolean input expression φ and results in an accordingly restricted g -grade TS A . The expression φ belongs to monotone one-in-three 3-SAT if and only if A has the ESSP or the SSP or the solvability, depending on which of the properties is queried.

The proofs of the announced polynomial time results base on a generalization of Schmitt's approach [8] that reduces ESSP and SSP to systems of linear equations modulo $b + 1$. It exploits that the solvability of such systems is decidable in polynomial time.

This paper is organized as follows: Section 2 introduces necessary definitions and provides them with illustrating examples. Moreover, it also presents some basic results that are used throughout the paper. Section 3 introduces the concept of unions applied by the proofs of our hardness results. Section 4 provides the NP-completeness results and presents the corresponding reductions that prove their validity. Section 5 provide the announced tractability results. Finally, Section 6 closes the paper. This paper is an extended version of [21, 22].

2. Preliminaries

In this section, we introduce necessary notions and provide some basic results that we use throughout the paper as well as some examples.

Definition 2.1. (Transition System)

A (deterministic) *transition system* (TS, for short) $A = (S, E, \delta)$ is a directed labeled graph with states S , events E and partial *transition function* $\delta : S \times E \rightarrow S$, where $\delta(s, e) = s'$ is interpreted as the *edge* $s \xrightarrow{e} s'$. For $s \xrightarrow{e} s'$ we say s is a source and s' is a target of e , respectively. An event e *occurs* at a state s , denoted by $s \xrightarrow{e}$, if $\delta(s, e)$ is defined. A word $w = e_0 \dots e_n \in E^*$ *occurs* at a state s , denoted by $s \xrightarrow{w}$, if it is the empty word ε or there are states q_0, \dots, q_n such that $s = q_0$ and $\delta(q_i, e_{i+1}) = q_{i+1}$ is defined for all $i \in \{0, \dots, n-1\}$. An *initialized* TS $A = (S, E, \delta, \iota)$ is a TS with a distinct state $\iota \in S$ such that every state $s \in S$ is *reachable* from ι by a directed labeled path. The language of A is the set $L(A) = \{w \in E^* \mid \iota \xrightarrow{w}\}$.

In the remainder of this paper, if not explicitly stated otherwise, we assume all TS to be initialized and if a TS A is not explicitly defined, then we refer to its components consistently by $S(A)$ (states) and $E(A)$ (events) and δ_A (transition function) and ι_A (initial states).

Definition 2.2. (g -grade, linear)

Let $g \in \mathbb{N}$. A TS $A = (S, E, \delta, \iota)$ is g -*grade* if, for every state $s \in S$, the number of incoming and outgoing labeled edges at s is at most g : $|\{e \in E \mid \xrightarrow{e} s\}| \leq g$ and $|\{e \in E \mid s \xrightarrow{e}\}| \leq g$. If a TS is 1-grade and cycle free, that is, there are pairwise distinct states s_0, \dots, s_m such that $A = s_0 \xrightarrow{e_1} \dots \xrightarrow{e_m} s_m$, then we say A is *linear*; we call s_m the *terminal state* of A and, for all $i < j \in \{1, \dots, m\}$, we say e_j and s_j occur after e_i .

In this paper, we deal with (different kinds of Petri) nets. Nets have places, transitions, a flow and an initial marking. Places can contain *tokens*. A global marking of a net defines for every place p how many tokens it contains initially. The firing of a transition can change locally the content of some places and thus globally the marking of the net. The flow defines the relations between places and transitions: how many token must a place contain to allow the firing of a transition and in which way changes the firing of a transition the content of a place. Nets are classified by the number of tokens that a place can maximally contain (markings) and according to how places and transitions may influence each other (flow). This way to classify nets leads to infinite many different classes of nets. In order to deal with these classes in a uniform way, the notion of types of nets has been developed in [6]:

Definition 2.3. (Type of nets)

A type of nets τ is a (non-initialized) TS $\tau = (S_\tau, E_\tau, \delta_\tau)$ with $S_\tau \subseteq \mathbb{N}$.

Based on this notion, we are now able to define τ -nets, where the states S_τ of $\tau = (S_\tau, E_\tau, \delta_\tau)$ correspond to possible contents of places, the events E_τ correspond to possible relations between places and transitions and the partial transition function δ_τ describe how the contents of places can be changed by the firing of a transition and, moreover, which contents can inhibit such a firing:

Definition 2.4. (τ -Nets)

Let $\tau = (S_\tau, E_\tau, \delta_\tau)$ be a type of nets. A Petri net $N = (P, T, M_0, f)$ of type τ , (τ -net, for short) is given by finite and disjoint sets P of places and T of transitions, an initial marking $M_0 : P \rightarrow S_\tau$, and a (total) flow function $f : P \times T \rightarrow E_\tau$. A τ -net realizes a certain behavior by firing sequences of transitions: A transition $t \in T$ can fire in a marking $M : P \rightarrow S_\tau$ if $\delta_\tau(M(p), f(p, t))$ is defined for all $p \in P$. By firing, t produces the next marking $M' : P \rightarrow S_\tau$ where $M'(p) = \delta_\tau(M(p), f(p, t))$ for all $p \in P$. This is denoted by $M \xrightarrow{t} M'$. Given a τ -net $N = (P, T, M_0, f)$, its behavior is captured by a transition system A_N , called the reachability graph of N . The state set of A_N is the reachability set $RS(N)$, that is, the set of all markings that, starting from initial state M_0 , are reachable by firing a sequence of transitions. For every reachable marking M and transition $t \in T$ with $M \xrightarrow{t} M'$ the state transition function δ_{A_N} of A_N is defined by $\delta_{A_N}(M, t) = M'$.

Let $b \in \mathbb{N}^+$ be arbitrary but fixed. In this paper, the following types of (b -bounded Petri) nets are the subject of our investigations:

Definition 2.5. (τ_{PT}^b)

The type of b -bounded P/T -nets $\tau_{PT}^b = (S_{\tau_{PT}^b}, E_{\tau_{PT}^b}, \delta_{\tau_{PT}^b})$ has the state set $S_{\tau_{PT}^b} = \{0, \dots, b\}$ and the event set $E_{\tau_{PT}^b} = \{0, \dots, b\}^2$ and, for all $s \in S_{\tau_{PT}^b}$ and all $(m, n) \in E_{\tau_{PT}^b}$, the transition function is defined by $\delta_{\tau_{PT}^b}(s, (m, n)) = s - m + n$ if $s \geq m$ and $s - m + n \leq b$, and undefined otherwise.

Definition 2.6. (τ_{PPT}^b)

The type $\tau_{PPT}^b = (S_{\tau_{PPT}^b}, E_{\tau_{PPT}^b}, \delta_{\tau_{PPT}^b})$ of pure b -bounded P/T -nets is a restriction of τ_{PT}^b that discards all events (m, n) from $E_{\tau_{PT}^b}$ where both m and n are positive. To be exact, $S_{\tau_{PPT}^b} = S_{\tau_{PT}^b}$ and $E_{\tau_{PPT}^b} = E_{\tau_{PT}^b} \setminus \{(m, n) \mid 1 \leq m, n \leq b\}$ and, for all $s \in S_{\tau_{PPT}^b}$ and all $e \in E_{\tau_{PPT}^b}$, we have $\delta_{\tau_{PPT}^b}(s, e) = \delta_{\tau_{PT}^b}(s, e)$.

Definition 2.7. ($\tau_{\mathbb{Z}PT}^b$)

The type $\tau_{\mathbb{Z}PT}^b = (S_{\tau_{\mathbb{Z}PT}^b}, E_{\tau_{\mathbb{Z}PT}^b}, \delta_{\tau_{\mathbb{Z}PT}^b})$ of \mathbb{Z}_{b+1} -extended b -bounded P/T -nets originates from τ_{PT}^b by extending the event set $E_{\tau_{PT}^b}$ with the elements $0, \dots, b$. The transition function additionally simulates the addition modulo $(b+1)$. More exactly, $S_{\tau_{\mathbb{Z}PT}^b} = S_{\tau_{PT}^b}$ and $E_{\tau_{\mathbb{Z}PT}^b} = (E_{\tau_{PT}^b} \setminus \{(0, 0)\}) \cup \{0, \dots, b\}$ and, for all $s \in S_{\tau_{\mathbb{Z}PT}^b}$ and all $e \in E_{\tau_{\mathbb{Z}PT}^b}$ we have that $\delta_{\tau_{\mathbb{Z}PT}^b}(s, e) = \delta_{\tau_{PT}^b}(s, e)$ if $e \in E_{\tau_{PT}^b}$, else $\delta_{\tau_{\mathbb{Z}PT}^b}(s, e) = (s + e) \bmod (b + 1)$.

Definition 2.8. ($\tau_{\mathbb{Z}PPT}^b$)

The type $\tau_{\mathbb{Z}PPT}^b = (S_{\tau_{\mathbb{Z}PPT}^b}, E_{\tau_{\mathbb{Z}PPT}^b}, \delta_{\tau_{\mathbb{Z}PPT}^b})$ of pure \mathbb{Z}_{b+1} -extended b -bounded P/T -nets is a restriction of $\tau_{\mathbb{Z}PT}^b$ such that $S_{\tau_{\mathbb{Z}PPT}^b} = S_{\tau_{\mathbb{Z}PT}^b}$ and $E_{\tau_{\mathbb{Z}PPT}^b} = E_{\tau_{\mathbb{Z}PT}^b} \setminus \{(m, n) \mid 1 \leq m, n \leq b\}$ and, for all $s \in S_{\tau_{\mathbb{Z}PPT}^b}$ and all $e \in E_{\tau_{\mathbb{Z}PPT}^b}$, the transition function is defined by $\delta_{\tau_{\mathbb{Z}PPT}^b}(s, e) = \delta_{\tau_{\mathbb{Z}PT}^b}(s, e)$.

Definition 2.9. ($\tau_{\mathbb{RZ}PT}^b$)

The type of restricted \mathbb{Z}_{b+1} -extended b -bounded P/T -nets $\tau_{\mathbb{RZ}PT}^b = (S_{\tau_{\mathbb{RZ}PT}^b}, E_{\tau_{\mathbb{RZ}PT}^b}, \delta_{\tau_{\mathbb{RZ}PT}^b})$ has the same state set $S_{\tau_{\mathbb{RZ}PT}^b} = S_{\tau_{\mathbb{Z}PT}^b}$ and the same event set $E_{\tau_{\mathbb{RZ}PT}^b} = E_{\tau_{\mathbb{Z}PT}^b}$ as $\tau_{\mathbb{Z}PT}^b$, but a restricted

transition function $\delta_{\tau_{RZPT}^b}$. In particular, $\delta_{\tau_{RZPT}^b}$ restricts $\delta_{\tau_{ZPT}^b}$ in such a way that for $s \in S_{\tau_{RZPT}^b}$ and $(m, n) \in E_{\tau_{RZPT}^b}$ we have that $\delta_{\tau_{RZPT}^b}(s, (m, n)) = \delta_{\tau_{ZPT}^b}(s, (m, n))$ if $s = m$; otherwise if $s \neq m$, then $\delta_{\tau_{RZPT}^b}(s, (m, n))$ remains undefined. Hence, every $(m, n) \in E_{\tau_{RZPT}^b}$ occurs exactly once in τ_{RZPT}^b . Furthermore, if $(s, e) \in \{0, \dots, b\}^2$ then $\delta_{\tau_{RZPT}^b}(s, e) = \delta_{\tau_{ZPT}^b}(s, e)$.

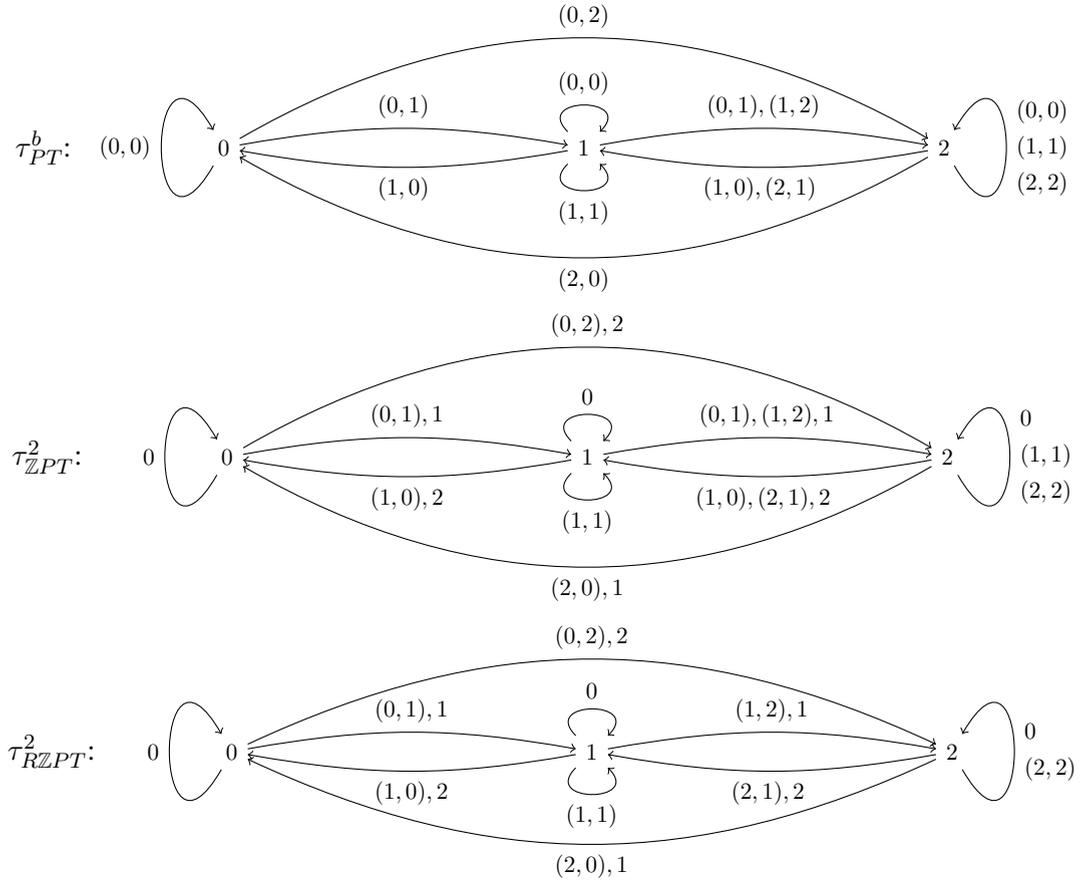


Figure 1: The 2-bounded types τ_{PT}^2 (top) and τ_{ZPT}^2 (middle) and τ_{RZPT}^2 (bottom), respectively, with states 0 and 1 and 2. Multiple edges with the same source and target (but different events) are represented as one edge with multiple events. For example, τ_{PT}^b has three edges from 2 to 2: one labeled (0, 0), another (1, 1), and a third labeled (2, 2). Similarly, τ_{ZPT}^2 has two edges from 2 to 0: one labeled (2, 0) and another labeled 1.

Example 2.10. Figure 1 sketches τ_{PT}^2 (top) and τ_{ZPT}^2 (middle). Events separated by commas label different edges. Omitting the events (1, 1), (1, 2), (2, 1) and (2, 2) and the corresponding edges yields τ_{PT}^2 and τ_{ZPT}^2 , respectively. Moreover, Figure 1 sketches τ_{RZPT}^2 (bottom).

Example 2.11. Figure 2 sketches the τ_{PPT}^2 -net N_1 and its reachability graph A_{N_1} . N_1 has places R_1 and R_2 , transitions a and b and flow $f(R_1, a) = (1, 0)$, $f(R_2, b) = (1, 0)$ and $f(R_1, b) = f(R_2, a) = (0, 0)$ and initial marking $M_0(R_1) = M_0(R_2) = 1$. The $(0, 0)$ -labeled edges are omitted.

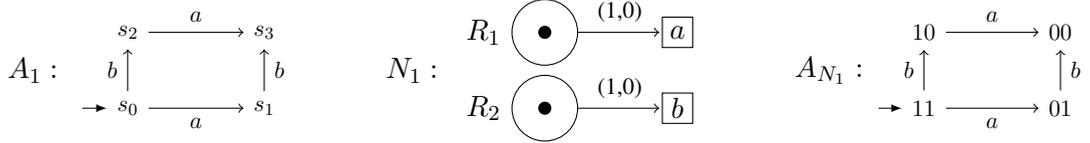


Figure 2: The TS A_1 , the τ_{PPT}^1 -Net N_1 and the reachability graph A_{N_1} of N_1 .

According to Definition 2.4, for every τ -net N , there is always a TS A_N , that reflects the global behavior of N , namely the corresponding reachability graph. Moreover, by firing all possible sequences of transitions, the reachability graph A_N can be computed effectively. Naturally, this raises the question whether a given TS A corresponds to the behavior of a τ -net N . Furthermore, in case of a positive decision, N should be constructed. This is the subject of the following search problem:

τ -SYNTHESIS

Input: A TS $A = (S, E, \delta, \iota)$.

Task: Find a τ -net N whose reachability graph is isomorphic to A if it exists.

If an input $A = (S, E, \delta, \iota)$ of τ -Synthesis allows a positive decision, then we want to construct a corresponding τ -net N purely from A . Since A and the reachability graph A_N of N shall be isomorphic, the events E of A become transitions of N . The places, the flow function and the initial marking of N originate from so-called τ -regions of A .

Definition 2.12. (τ -Regions)

Let $\tau \in \{\tau_0^b, \tau_1^b, \tau_2^b, \tau_3^b, \tau_4^b\}$ and $A = (S, E, \delta, \iota)$ be a TS. A τ -region of A is a pair (sup, sig) of support $sup : S \rightarrow S_\tau$ and signature $sig : E \rightarrow E_\tau$ such that for every edge $s \xrightarrow{e} s'$ of A the image $sup(s) \xrightarrow{sig(e)} sup(s')$ is present in τ . If $sig(e) = (m, n)$, then we define $sig^-(e) = m$ and $sig^+(e) = n$ and $|sig(e)| = 0$, and if $sig(e) \in \{0, \dots, b\}$, then we define $sig^-(e) = sig^+(e) = 0$ and $|sig(e)| = sig(e)$.

A region (sup, sig) models a place p and its initial marking $M_0(p)$ as well as the corresponding part of the flow function $f(p, \cdot)$ of a sought τ -net if it exist. In particular, $sig(e)$ models $f(p, e)$ and $sup(\iota)$ models the number of tokens that p contains initially and, more generally, $sup(s)$ models the number of tokens $M(p)$ in the marking M that corresponds to the state s according to the isomorphism φ that justifies $A \cong A_N$.

Definition 2.13. (Synthesized net)

Every set \mathcal{R} of τ -regions of $A = (S, E, \delta, \iota)$ defines the synthesized τ -net $N_A^{\mathcal{R}} = (\mathcal{R}, E, f, M_0)$ with set of places \mathcal{R} , set of transitions E , flow function $f((sup, sig), e) = sig(e)$ and initial marking $M_0((sup, sig)) = sup(\iota)$ for all $(sup, sig) \in \mathcal{R}$ and all $e \in E$.

To make sure that a synthesized net N realizes the behavior of a TS exactly, distinct states s and s' of A must correspond to different markings M and M' of the net. Moreover, the firing of a transition e needs to be inhibited at a marking M , when the event e does not occur at the state s that corresponds to M by the isomorphism φ . This is stated by so-called separation atoms and separation properties.

Definition 2.14. (τ -State Separation)

Let τ be a type of nets and $A = (S, E, \delta, \iota)$ a TS. A pair (s, s') of distinct states of A defines a *state separation atom* (SSA, for short). A τ -region $R = (sup, sig)$ solves (s, s') if $sup(s) \neq sup(s')$. The meaning of R is to ensure that $N_A^{\mathcal{R}}$ contains at least one place R such that $M(R) \neq M'(R)$ for the markings M and M' corresponding to s and s' , respectively. If $s \in S$ is a state of A and, for all states $s' \in S$ such that $s' \neq s$, there is a τ -region that solves (s, s') then s is called τ -solvable. If every state of A or, equivalently, every SSA of A is τ -solvable, then A has the τ -state separation property (τ -SSP, for short).

Definition 2.15. (τ -Event State Separation)

Let τ be a type of nets and $A = (S, E, \delta, \iota)$ a TS. A pair (e, s) of event $e \in E$ and state $s \in S$ where e does not occur at s , that is $\neg s \xrightarrow{e}$, defines an *event state separation atom* (ESSA atom, for short). A τ -region $R = (sup, sig)$ solves (e, s) if $sig(e)$ is not defined at $sup(s)$ in τ , that is, $\neg \delta_{\tau}(sup(s), sig(e))$. The meaning of R is to ensure that there is at least one place R in $N_A^{\mathcal{R}}$ such that $\delta_{\tau}(M(R), f(R, e))$ is not defined for the marking M that corresponds to s via the isomorphism, that is, e cannot fire in M . If, for all $s \in S$ such that $\neg s \xrightarrow{e}$, there is a τ -region that solves (e, s) , then e is called τ -solvable. If every event of A or, equivalently, every ESSA of A is τ -solvable, then A has the τ -event state separation property (τ -ESSP, for short).

Definition 2.16. (Witness, τ -admissible set)

A set \mathcal{R} of τ -region is a (τ -) *witness* of the τ -(E)SSP of A if it contains for every (E)SSA a τ -region that solves it. If A has the τ -SSP and the τ -ESSP, then A is called τ -solvable. A set \mathcal{R} that is a witness of both the τ -SSP and the τ -ESSP of A is called τ -admissible.

The following lemma, borrowed from [6, p.163], summarizes the already implied connection between the existence of τ -admissible sets of A and (the solvability of) τ -synthesis:

Lemma 2.17. ([6])

Let A be a TS and τ a type of nets. The reachability graph A_N of a τ -net N is isomorphic to A if and only if there is a τ -admissible set \mathcal{R} of A such that $N = N_A^{\mathcal{R}}$.

Example 2.18. Let $\tau \in \{\tau_{PT}^b, \tau_{PT}^b \mid b \in \mathbb{N}^+\}$. The TS A_1 of Figure 2 has the τ -ESSP and the τ -SSP: The region $R_1 = (sup_1, sig_1)$, which is defined by $sup_1(s_0) = sup_1(s_2) = 1$ and $sup_1(s_1) = sup_1(s_3) = 0$ and $sig_1(a) = (1, 0)$ and $sig_1(b) = (0, 0)$, solves the ESSA (a, s_1) and (a, s_3) as well as the SSA (s_0, s_1) and (s_0, s_3) and (s_2, s_1) and (s_2, s_3) . Moreover, the region $R_2 = (sup_2, sig_2)$, which is defined by $sup_2(s_0) = sup_2(s_1) = 1$, $sup_2(s_2) = sup_2(s_3) = 0$, $sig_2(a) = (0, 0)$ and $sig_2(b) = (1, 0)$ solves the remaining ESSA (b, s_2) and (b, s_3) as well as the SSA (s_0, s_2) and (s_1, s_3) of A_1 . Since R_1 and R_2 solve all SSA and ESSA, $\mathcal{R} = \{R_1, R_2\}$ is a τ -admissible set. Figure 2 sketches the synthesized net $N_1 = N_{A_1}^{\mathcal{R}}$, where $(0, 0)$ -labeled flow edges are omitted, and its reachability graph

A_{N_1} . The isomorphism φ between A_1 and A_{N_1} is given by $\varphi(s_0) = 11$, $\varphi(s_1) = 01$, $\varphi(s_2) = 10$ and $\varphi(s_3) = 00$.

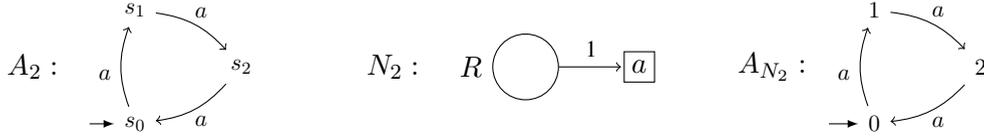


Figure 3: The TS A_2 , the $\tau_{\mathbb{Z}PPT}^2$ -Net N_2 and the reachability graph A_{N_2} of N_2 .

Example 2.19. The TS A_2 of Figure 3 has no ESSA, since the only event a occurs at every state of A_2 . Consequently, A_2 has the τ -ESSP for all types of nets. However, A_2 has the SSA (s_0, s_1) , (s_0, s_2) and (s_1, s_2) . If $\tau \in \{\tau_{PPT}^b, \tau_{PT}^b \mid b \in \mathbb{N}^+\}$, then neither of these atoms is τ -solvable, since every τ -region $R = (sup, sig)$ of A_2 satisfies $sup(s_0) = sup(s_0) - 2sig^-(a) + 2sig^+(a)$, which implies $sig(a) = (0, 0)$ and thus $sup(s_0) = sup(s_1) = sup(s_2)$. Nevertheless, if $\tau \in \{\tau_{\mathbb{Z}PPT}^b, \tau_{\mathbb{Z}PT}^b \mid b \geq 2\}$, then A_2 has the τ -SSP, since the following τ -Region $R = (sup, sig)$ solves all SSA in one blow: $sup(s_0) = 0$, $sup(s_1) = 1$ and $sup(s_2) = 2$ and $sig(a) = 1$. Since A_2 has also the τ -ESSP, $\mathcal{R} = \{R\}$ is a τ -admissible set of A_2 .

Figure 3 sketches the synthesized net $N_2 = N_{A_2}^{\mathcal{R}}$ and its reachability graph A_{N_2} . This example also shows that the group-extended types $\tau_{\mathbb{Z}PPT}$ and $\tau_{\mathbb{Z}PT}$ are strictly more powerful than the types τ_{PPT} and τ_{PT} .

A purpose of this paper is to characterize the computational complexity of τ -SYNTHESIS for all introduced b -bounded types of nets completely. Since the corresponding complexity classes are defined for decision problems, we restrict our investigations to the decision version of τ -SYNTHESIS that is called τ -SOLVABILITY. By Lemma 2.17, there is a τ -admissible set \mathcal{R} of A if and only if there is a τ -net N whose reachability graph is isomorphic to A . This allows us to formulate the solvability problem for τ -nets as follows:

τ -SOLVABILITY

Input: A TS $A = (S, E, \delta, \iota)$.

Question: Does there exist a τ -admissible set \mathcal{R} of A ?

Although we are mainly interested in synthesis, the τ -SSP and the τ -ESSP are also interesting on their own. This is because, for example, an algorithm that decides in polynomial time whether A has the τ -SSP or the τ -ESSP could serve as a pre-synthesis method, which rejects inputs that does not have the property in question. This leads to the following decision problems:

τ -SSP

Input: A TS $A = (S, E, \delta, \iota)$.

Question: Does there exist a witness \mathcal{R} for the τ -SSP of A ?

τ -ESSP

Input: A TS $A = (S, E, \delta, \iota)$.

Question: Does there exist a witness \mathcal{R} for the τ -ESSP of A ?

In [7], it was originally shown that τ_{PPT}^1 -SOLVABILITY (there referred to as *elementary net synthesis*) is NP-complete. In [11, 19], we have shown that this remains true even for strongly restricted inputs and applies also to τ_{PPT}^1 -SSP and τ_{PPT}^1 -ESSP. Moreover, the type τ_{ZPPT}^1 coincides with Schmitt's type (*flip-flop nets*) for which the considered decision problems are tractable [8]. In [9, p. 619], this characterization was found to be true for τ_{ZPT}^1 (there referred to as the Boolean type of nets $\tau = \{\text{nop, inp, out, used, swap}\}$) as well. In this paper, we complete the complexity characterization of τ -SOLVABILITY, τ -SSP and τ -ESSP for all introduced b -bounded types of nets and all $b \in \mathbb{N}^+$. (Observe that the problems are trivial if $b = 0$.) Figure 4 provides an overview over our findings and shows, depending on τ and b , which of the problems are NP-complete (NPC) and which are solvable in polynomial time (P).

Bound	Problem	τ_{PPT}^b	τ_{PT}^b	τ_{ZPPT}^b	τ_{ZPT}^b	τ_{RZPT}^b
$b = 1$	SSP	NPC	NPC	P	P	P
	ESSP	NPC	NPC	P	P	P
	SOLVABILITY	NPC	NPC	P	P	P
$b \geq 2$	SSP	NPC	NPC	P	P	P
	ESSP	NPC	NPC	NPC	NPC	P
	SOLVABILITY	NPC	NPC	NPC	NPC	P

Figure 4: Overview of the computational complexity of τ -SOLVABILITY, τ -SSP and τ -ESSP for all $\tau \in \{\tau_{PPT}^b, \tau_{PT}^b, \tau_{ZPPT}^b, \tau_{ZPT}^b, \tau_{RZPT}^b\}$ and all $b \in \mathbb{N}^+$.

In the following, if not explicitly stated otherwise, for all $\tau \in \{\tau_{PT}^b, \tau_{PPT}^b\}$, we let $b \in \mathbb{N}^+$ and, for all $\tau \in \{\tau_{ZPT}^b, \tau_{ZPPT}^b\}$, we let $2 \leq b \in \mathbb{N}$ be arbitrary but fixed, since the case $b = 1$ is already solved for the latter. The observations of the next lemma are used to simplify our proofs:

Lemma 2.20. Let $\tau \in \{\tau_{PPT}^b, \tau_{PT}^b, \tau_{ZPT}^b, \tau_{ZPPT}^b\}$ and $A = (S, E, \delta, \iota)$ be a TS.

- Two mappings $sup : S \rightarrow S_\tau$ and $sig : E \rightarrow E_\tau$ define a τ -region of A if and only if for every directed labeled path $q_0 \xrightarrow{e_1} \dots \xrightarrow{e_m} q_m$ of A holds $sup(q_i) = sup(q_{i-1}) - sig^-(e_i) + sig^+(e_i) + |sig(e_i)|$ for all $i \in \{1, \dots, \ell\}$, where this equation is to consider modulo $(b + 1)$. In particular, every region (sup, sig) is implicitly completely defined by sig and $sup(\iota)$.
- If $s_0, s_1, \dots, s_b \in S$, $e \in E$ and $s_0 \xrightarrow{e} \dots \xrightarrow{e} s_b$ then a τ -region (sup, sig) of A satisfies $sig(e) = (m, n)$ with $m \neq n$ if and only if $(m, n) \in \{(1, 0), (0, 1)\}$. If $sig(e) = (0, 1)$ then $sup(s_0) = 0$ and $sup(s_b) = b$. If $sig(e) = (1, 0)$ then $sup(s_0) = b$ and $sup(s_b) = 0$.

Proof:

(1): The first claim follows directly from the definitions of τ and τ -regions. For the second claim, we observe that every state $s \in S$ is reachable by a directed labeled path $q_0 \xrightarrow{e_1} \dots \xrightarrow{e_m} q_m$, where $q_0 = \iota$ and $q_m = s$. Thus, if $\text{sup}(\iota)$ and a valid signature sig are given, then, by the first claim, we get $\text{sup}(s)$ by $\text{sup}(s) = \text{sup}(\iota) + \sum_{i=1}^m (-\text{sig}^-(e_i) + \text{sig}^+(e_i) + |\text{sig}(e_i)|)$.

(2): The *If*-direction is trivial. For the *Only-if*-direction we show that the assumption $(m, n) \notin \{(1, 0), (0, 1)\}$ yields a contradiction.

By (1), we have that $\text{sup}(s_b) = \text{sup}(s_0) + b \cdot (n - m)$. If $|n - m| > 1$, then either $b \cdot (n - m) < -b$ or $b \cdot (n - m) > b$; since $0 \leq \text{sup}(s_0) \leq b$, the first case contradicts $\text{sup}(s_b) \geq 0$, and the latter case contradicts $\text{sup}(s_b) \leq b$, respectively. Hence, if $n \neq m$ then $|n - m| = 1$. For a start, we show that $m > n$ implies $m = 1$ and $n = 0$. By $n \leq m - 1$ and $\text{sup}(s_0) \leq b$ we obtain the estimation

$$\text{sup}(s_{b-1}) = \text{sup}(s_0) + (b-1)(n-m) \leq b + (b-1)(m-1-m) = 1$$

By $n < m \leq \text{sup}(s_{b-1}) \leq 1$ we have $(m, n) = (1, 0)$. Similarly, we obtain that $(m, n) = (0, 1)$ if $m < n$. Hence, if $\text{sig}(e) = (m, n)$ and $n \neq m$ then $\text{sig}(e) \in \{(1, 0), (0, 1)\}$.

The second statement follows directly from (1). \square

The following lemma shows that if A is a linear TS and \mathcal{R} is a witness of the τ -ESSP of A , then \mathcal{R} witnesses also the τ -SSP of A . In particular, this implies that a linear TS A is τ -solvable if and only if it has the τ -ESSP. Notice that Lemma 2.21 provides a very general result, since its statement is independent from the actual choice of τ .

Lemma 2.21. (ESSP implies SSP for Linear TS)

Let τ be a type of nets and let $A = s_0 \xrightarrow{e_1} \dots \xrightarrow{e_n} s_n$ be a linear TS and let \mathcal{R} be a set of τ -regions of A . If \mathcal{R} is a witness of the τ -ESSP of A , then \mathcal{R} witnesses also the τ -SSP of A .

Proof:

Let \mathcal{R} be a witness of the τ -ESSP of A . Assume that there is an SSA that can not be solved by a region of \mathcal{R} . Then there is an SSA $\alpha = (z_{i_j}, z_{i_k})$ of A , where $i_j, i_k \in \{0, \dots, n\}$, such that the state z_{i_k} has the maximum index among all states of A that participate at SSA of A that can not be solved by a region of \mathcal{R} : if (z_{i_ℓ}, z_{i_m}) is an SSA of A that can not be solved by regions of \mathcal{R} , then $i_\ell \leq i_k$ and $i_m \leq i_k$. In particular, this implies $i_j < i_k$. Since $i_j < i_k$, there is the edge $z_{i_j} \xrightarrow{e_{i_j+1}} z_{i_j+1}$ in A . Since α is not solvable by regions of \mathcal{R} , we have $\text{sup}(z_{i_j}) = \text{sup}(z_{i_k})$ for all $(\text{sup}, \text{sig}) \in \mathcal{R}$. This implies, that the event e_{i_j+1} occurs at z_{i_k} , since the ESSA (e_{i_j+1}, z_{i_k}) would not be solvable otherwise:

$\text{sup}(z_{i_j}) \xrightarrow{\text{sig}(e_{i_j+1})}$ and $\neg \text{sup}(z_{i_k}) \xrightarrow{\text{sig}(e_{i_j+1})}$, implies the contradiction $\text{sup}(z_{i_j}) \neq \text{sup}(z_{i_k})$. Hence, $z_{i_k} \xrightarrow{e_{i_j+1}} z_{i_k+1}$ is an edge in A . Since $\text{sup}(z_{i_j}) = \text{sup}(z_{i_k})$ for all $(\text{sup}, \text{sig}) \in \mathcal{R}$ and since δ_τ is a function, by $z_{i_j} \xrightarrow{e_{i_j+1}} z_{i_j+1}$ and $z_{i_k} \xrightarrow{e_{i_j+1}} z_{i_k+1}$, we get $\text{sup}(z_{i_j+1}) = \text{sup}(z_{i_k+1})$ for all $(\text{sup}, \text{sig}) \in \mathcal{R}$. In particular, the SSA (z_{i_j+1}, z_{i_k+1}) is not solvable by regions of \mathcal{R} . Since $i_k < i_k + 1$, this contradicts the choice of α . Consequently, α does not exist and \mathcal{R} witnesses the τ -SSP of A . \square

3. The concept of unions

For our reductions, we use the technique of *component design* [23]. Every implemented constituent is a TS (in the context of the reduction also referred to as gadget) that locally ensures the satisfaction of some constraints. Commonly, all constituents are finally joined together in a target instance (TS) such that all required constraints are properly globally translated. However, the concept of unions saves us the need to actually create the target instance:

Definition 3.1. (Union)

If $A_0 = (S_0, E_0, \delta_0, \iota_0), \dots, A_n = (S_n, E_n, \delta_n, \iota_n)$ are TS with pairwise disjoint states (but not necessarily disjoint events) then we call $U(A_0, \dots, A_n)$ their *union* with set of states $S(U) = \bigcup_{i=0}^n S_i$ and set of events $E(U) = \bigcup_{i=0}^n E_i$.

Let $\tau = (Z_\tau, E_\tau, \delta_\tau)$ be a type of nets and $U = U(A_0, \dots, A_n)$ a union, where $A_i = (S_i, E_i, \delta_i, \iota_i)$ for all $i \in \{0, \dots, n\}$. The concepts of SSA, ESSA, τ -regions, τ -SSP, and τ -ESSP as defined in the preliminaries are transferred to U as follows:

Definition 3.2. (Region of a Union)

A pair (sup, sig) of mappings $sup : S(U) \rightarrow S_\tau$ and $sig : E(U) \rightarrow E_\tau$ is called a τ -region (of U), if $s \xrightarrow{e} s' \in A_i$ implies $sup(s) \xrightarrow{sig(e)} sup(s') \in \tau$ for all $i \in \{0, \dots, n\}$.

Definition 3.3. (τ -State Separation in Unions)

A pair (s, s') of distinct states $s, s' \in S(U)$ of the *same* TS A_i , where $i \in \{0, \dots, n\}$, defines an SSA of U . A τ -region (sup, sig) of U solves (s, s') , if $sup(s) \neq sup(s')$. U has the τ -SSP, if all of its SSA are τ -solvable.

Definition 3.4. (τ -Event State Separation in Unions)

A pair (e, s) of event $e \in E(U)$ and state $s \in S(U)$ such that $\neg s \xrightarrow{e}$, defines an ESSA of U . A τ -region (sup, sig) of U solves it, if $\neg sup(s) \xrightarrow{e}$. U has the τ -ESSP if all of its ESSA are τ -solvable.

In the same way, the notion of *witness* and τ -admissible set and τ -solvable are transferred to unions. From the perspective of τ -SSP and τ -ESSP, unions are intended to treat a lot of unjoined TS as if they were joined to a TS. To be able to do so, in the following, we introduce the *linear joining* $LJ(U)$ and the *joining* $J(U)$ of a union U and argue that $LJ(U)$ or $J(U)$ has the τ -(E)SSP if and only if U has the τ -(E)SSP.

Definition 3.5. (Linear Joining)

Let $U = U(A_0, \dots, A_n)$ be a union such that, for all $i \in \{0, \dots, n\}$, the TS $A_i = (S_i, E_i, \delta_i, \iota_i)$ is linear and its terminal state is t_i and let $Q = \{q_1, \dots, q_n\}$ be a set of states, which is disjoint with $S(U)$, and $W = \{w_1, \dots, w_n\}$ and $Y = \{y_1, \dots, y_n\}$ be sets of events which are disjoint with $E(U)$. The *linear joining* of U is the linear TS $LJ(U) = (S(U) \cup Q, E(U) \cup W \cup Y, \delta, \iota_0)$ with transition

function δ that is, for all $e \in E(U) \cup W \cup Y$ and all $s \in S(U) \cup Q$, defined as follows:

$$\delta(s, e) = \begin{cases} \delta_i(s, e), & \text{if } s \in S_i \text{ and } e \in E_i \text{ and } i \in \{0, \dots, n\} \\ q_{i+1}, & \text{if } s = t_i \text{ and } e = w_{i+1} \text{ and } i \in \{0, \dots, n-1\} \\ \iota_i, & \text{if } s = q_i \text{ and } e = y_i \text{ and } i \in \{1, \dots, n\} \\ \text{undefined,} & \text{otherwise} \end{cases}$$

Remark 3.6. The linear joining $LJ(U)$ of U can be sketched as follows:

$$LJ(U) = A_0 \xrightarrow{w_1} q_1 \xrightarrow{y_1} A_1 \xrightarrow{w_2} q_2 \xrightarrow{y_2} \dots \xrightarrow{w_n} q_n \xrightarrow{y_n} A_n$$

Definition 3.7. (Joining)

Let $U = U(A_0, \dots, A_n)$ be a union of TS $A_i = (S_i, E_i, \delta_i, \iota_i)$ for all $i \in \{0, \dots, n\}$, and let $Q = \{q_0, \dots, q_n\}$ be a set of states, which is disjoint with $S(U)$, and $W = \{w_1, \dots, w_n\}$ and $Y = \{y_0, \dots, y_n\}$ be sets of events, which are disjoint with $E(U)$. The *joining* of U is the TS $J(U) = (S(U) \cup Q, E(U) \cup W \cup Y, \delta, q_0)$ with transition function δ that is, for all $e \in E(U) \cup W \cup Y$ and all $s \in S(U) \cup Q$, defined as follows:

$$\delta(s, e) = \begin{cases} \delta_i(s, e), & \text{if } s \in S_i \text{ and } e \in E_i \text{ and } i \in \{0, \dots, n\} \\ q_{i+1}, & \text{if } s = q_i \text{ and } e = w_{i+1} \text{ and } i \in \{0, \dots, n-1\} \\ \iota_i, & \text{if } s = q_i \text{ and } e = y_i \text{ and } i \in \{0, \dots, n\} \\ \text{undefined,} & \text{otherwise} \end{cases}$$

Remark 3.8. The joining $J(U)$ of U can be sketched as follows:

$$\begin{array}{ccccccc} & & q_0 & \xrightarrow{w_1} & q_1 & \xrightarrow{w_2} & \dots & \xrightarrow{w_n} & q_n \\ & & \downarrow y_0 & & \downarrow y_1 & & & & \downarrow y_n \\ & & A_0 & & A_1 & & & & A_n \end{array}$$

The following lemma proves the announced functionality of unions. For technical reasons, we restrict ourselves to unions U where for every event $e \in E(U)$ there is at least one ESSA (e, s) to solve. The unions of our reductions satisfy this property, which is used to ensure that if U has the τ -ESSP, then $LJ(U)$ and $J(U)$ have the τ -ESSP, too. Moreover, our reductions ensure that if $\tau \in \{\tau_{PT}^b, \tau_{PPT}^b\}$, then only the linear joining $LJ(U)$ is to consider, and if $\tau \in \{\tau_{ZPT}^b, \tau_{ZPPT}^b\}$, then only the joining $J(U)$ is used. Thus, for the sake of simplicity, the lemma is formulated accordingly.

Lemma 3.9.

1. Let $U = U(A_0, \dots, A_n)$ be a union of linear TS such that t_i is the terminal state of $A_i = (S_i, E_i, \delta_i, \iota_i)$ for all $i \in \{0, \dots, n\}$, and for every event $e \in E(U)$ there is a state $s \in S(U)$ with $\neg s \xrightarrow{e}$. If $\tau \in \{\tau_{PT}^b, \tau_{PPT}^b\}$, then U has the τ -ESSP, respectively the τ -SSP, if and only if $LJ(U)$ has the τ -ESSP, respectively the τ -SSP.

2. Let $U = U(A_0, \dots, A_n)$ be a union such that $A_i = (S_i, E_i, \delta_i, \iota_i)$ for all $i \in \{0, \dots, n\}$, and for every event $e \in E(U)$ there is a state $s \in S(U)$ with $\neg s \xrightarrow{e}$. If $\tau \in \{\tau_{\mathbb{Z}PT}^b, \tau_{\mathbb{Z}PPT}^b\}$, then U has the τ -ESSP, respectively the τ -SSP, if and only if $J(U)$ has the τ -ESSP, respectively the τ -SSP.

Proof:

(1): The *if*-direction is trivial.

Only-if: Let $R = (sup, sig)$ be a τ -region of U , which solves an ESSA (a, z) or an SSA (z, z') of U . We can extend R to a τ -region $R' = (sup', sig')$ of $LJ(U)$ that also solves these atoms, by defining R' for all $s \in S(U) \cup Q$ and all $e \in E(U) \cup W \cup Y$ as follows:

$$sup'(s) = \begin{cases} sup(s), & \text{if } s \in S(U), \\ sup(z), & \text{if } s \in Q \end{cases}$$

$$sig'(e) = \begin{cases} sig(e), & \text{if } e \in E(U), \\ (sup(t_i) - sup(z), 0) & \text{if } e = w_{i+1} \text{ and } sup(t_i) > sup(z) \text{ and } i \in \{0, \dots, n-1\} \\ (0, sup(z) - sup(t_i)) & \text{if } e = w_{i+1} \text{ and } sup(t_i) \leq sup(z) \text{ and } i \in \{0, \dots, n-1\} \\ (0, sup(\iota_i) - sup(z)) & \text{if } e = y_i \text{ and } sup(\iota_i) > sup(z) \text{ and } i \in \{1, \dots, n\} \\ (sup(z) - sup(\iota_i), 0) & \text{if } e = y_i \text{ and } sup(\iota_i) \leq sup(z) \text{ and } i \in \{1, \dots, n\} \end{cases}$$

Notice that this extension also τ -solves (e, q_i) for all $i \in \{1, \dots, n\}$, since $sup(q_i) = sup(z)$. Since, for every event $e \in E(U)$ there is an atom (e, s) to solve, this implies that all events of U are τ -solvable in $LJ(U)$. Moreover, it is easy to see that the connector states q_1, \dots, q_n and the connector events y_1, \dots, y_n are τ -solvable: If $i \in \{1, \dots, n\}$ is arbitrary but fixed then the following region $R = (sup, sig)$ (by Lemma 2.20, completely defined) τ -solves q_i and y_i : $sup(\iota_0) = b$; for all $e \in E(LJ(U))$, if $e = y_i$, then $sig(e) = (0, b)$; if $e = w_i$, then $sig(e) = (b, 0)$; otherwise $sig(e) = (0, 0)$.

So far, we have already proven that if U has the τ -SSP, then $LJ(U)$ has the τ -SSP, too. Thus, to prove that the τ -ESSP of U implies τ -ESSP of $LJ(U)$, it remains to show that w_1, \dots, w_n are solvable if U has the τ -ESSP. Let $i \in \{1, \dots, n\}$ be arbitrary but fixed. The following region $R = (sup, sig)$ solves (w_i, s) for all $s \in S(LJ(U)) \setminus S_{i-1}$: if $i = 1$, then $sup(\iota_0) = 0$, otherwise $sup(\iota_0) = b$; for all $e \in E(LJ(U))$, if $i \neq 1$ and $e = y_{i-1}$, then $sig(e) = (b, 0)$; if $e = w_i$, then $sig(w_i) = (0, b)$; otherwise, $sig(e) = (0, 0)$.

It remains to argue that (w_i, s) is τ -solvable for all $s \in S_{i-1}$. Since U has the τ -ESSP, there is a set \mathcal{R} of regions that witnesses the τ -ESSP. In particular, for every ESSA (e, s) of A_{i-1} there is a region of \mathcal{R} that solves it. Restricting the corresponding regions to A_{i-1} yields a set of regions that witnesses the τ -ESSP of A_{i-1} . Since A_i is linear, by Lemma 2.21, these regions witness also the τ -SSP of A_{i-1} . Consequently, for every state $s \in S_{i-1} \setminus \{t_{i-1}\}$, there is a region $(sup, sig) \in \mathcal{R}$ such that $sup(s) \neq sup(t_{i-1})$. We extend this region to a region of $LJ(U)$ that solves (w_i, s) as follows: Besides of w_i, q_i and y_i , the extension of (sup, sig) is defined as R' above; if $sup(s) > sup(t_{i-1})$, then $sup(q_i) = b$, otherwise $sup(q_i) = 0$; if $sup(q_i) = b$, then $sig(w_i) = (0, b - sup(t_{i-1}))$; otherwise $sig(w_i) = (sup(t_{i-1}), 0)$; finally, if $sup(q_i) = b$, then $sig(y_i) = (b - sup(\iota_i), 0)$; otherwise $sig(y_i) = (0, sup(\iota_i))$.

(2): *If*: Again, the *if*-direction is trivial.

Only-if: Let $R = (sup, sig)$ be a τ -region of U , which solves an ESSA (a, z) or an SSA (z, z') of U . We can extend R to a τ -region $R' = (sup', sig')$ of $J(U)$ that also solves these atoms, by defining R' for all $s \in S(U) \cup Q$ and all $e \in E(U) \cup W \cup Y$ as follows:

$$sup'(s) = \begin{cases} sup(s), & \text{if } s \in S(U), \\ sup(z), & \text{if } s \in Q \end{cases}$$

$$sig'(e) = \begin{cases} sig(e), & \text{if } e \in E(U), \\ 0, & \text{if } e \in W, \\ (sup(z) - sup(\iota_i), 0) & \text{if } e = y_i \text{ and } sup(\iota_i) < sup(z) \text{ and } i \in \{0, \dots, n\} \\ (0, sup(\iota_i) - sup(z)) & \text{if } e = y_i \text{ and } sup(\iota_i) \geq sup(z) \text{ and } i \in \{0, \dots, n\} \end{cases}$$

Notice that R' also solves (a, q_i) for all $i \in \{0, \dots, n\}$ as $sup(q_i) = sup(z)$. Consequently, since there is at least one state $s \in S(U)$ for every event $e \in E(U)$ such that (e, s) is an ESSA of U , the atom (e, q_i) is solvable for every $e \in E(U)$ and every $i \in \{0, \dots, n\}$. As a result, to prove the τ -(E)SSP for $J(U)$ it remains to argue that the remaining SSA and ESSA at which the states of Q and the events of $W \cup Y$ participate are solvable in $J(U)$. If $i \in \{0, \dots, n\}$ and $s \in S(J(U))$ and $e \in E(J(U))$ then the following region (sup, sig) simultaneously solves every valid atom (y_i, \cdot) , (q_i, \cdot) and (w_{i+1}, \cdot) in $J(U)$ (if the latter exists):

$$sup(s) = \begin{cases} 0, & \text{if } s = q_i \\ b, & \text{otherwise} \end{cases} \quad sig(e) = \begin{cases} (0, b), & \text{if } e = y_i \text{ or } (i < n \text{ and } e = w_{i+1}) \\ (b, 0), & \text{if } 1 \leq i \text{ and } e = w_{i-1} \\ 0, & \text{otherwise} \end{cases} \quad \square$$

4. NP-completeness results

The following theorem is the main contribution of this section:

Theorem 4.1. 1. If $\tau \in \{\tau_{PT}^b, \tau_{PPT}^b\}$, then τ -SOLVABILITY and τ -ESSP and τ -SSP are NP-complete, even when restricted to linear TS.

2. Let $\tau \in \{\tau_{ZPT}^b, \tau_{ZPPT}^b\}$. For any fixed $g \geq 2$, τ -SOLVABILITY and τ -ESSP are NP-complete, even when restricted to g -grade TS.

For the proof of Theorem 4.1, on the one hand, we have to argue that τ -SOLVABILITY, τ -ESSP and τ -SSP are in NP. This can be seen as follows. By Definition 2.14 and Definition 2.15, a TS $A = (S, E, \delta, \iota)$ has at most $|S|^2$ SSA and at most $|S| \cdot |E|$ ESSA, respectively. This implies that if a TS A is τ -solvable or has the τ -SSP or the τ -ESSP, then there is a set of τ -regions \mathcal{R} of A of size at most $|S|^2 + |S| \cdot |E|$ that witnesses the corresponding property of A . Consequently, there is a non-deterministic Turing machine that guesses \mathcal{R} in a non-deterministic computation and verifies the validity of \mathcal{R} in (deterministic) polynomial time.

On the other hand, we have to argue that the decision problems are NP-hard for accordingly restricted input TS. The NP-hardness proofs base on polynomial-time reductions of the following decision problem, which has been shown to be NP-complete in [20]:

CUBIC MONOTONE ONE-IN-THREE 3-SAT (CM1IN33SAT)	
<i>Input:</i>	A boolean expression $\varphi = \{\zeta_0, \dots, \zeta_{m-1}\}$ of 3-clauses such that, for all $i \in \{0, \dots, m-1\}$, the clause $\zeta_i = \{X_{i_0}, X_{i_1}, X_{i_2}\}$ contains 3 distinct non-negated variables, where $i_0 < i_1 < i_2$; every variable $X \in V(\varphi)$ occurs in exactly three distinct clauses, where $V(\varphi) = \bigcup_{i=0}^{m-1} \zeta_i$ denotes the set of all variables of φ .
<i>Question:</i>	Does there exist a one-in-three model of φ , that is, a subset $M \subseteq V(\varphi)$ such that $ M \cap \zeta_i = 1$ for all $i \in \{0, \dots, m-1\}$?

Notice that the characterization of the input φ implies $|V(\varphi)| = m$. The following example provides –up to renaming– the smallest instance of CM1IN33SAT that allows a positive decision:

Example 4.2. The boolean expression $\varphi = \{\zeta_0, \dots, \zeta_5\}$ with clauses $\zeta_0 = \{X_0, X_1, X_2\}$, $\zeta_1 = \{X_0, X_2, X_3\}$, $\zeta_2 = \{X_0, X_1, X_3\}$, $\zeta_3 = \{X_2, X_4, X_5\}$, $\zeta_4 = \{X_1, X_4, X_5\}$, $\zeta_5 = \{X_3, X_4, X_5\}$ is a well-defined input of CM1IN33SAT and has the one-in-three model $M = \{X_0, X_4\}$.

General reduction approach. Let $\tau \in \{\tau_{PT}^b, \tau_{PPT}^b, \tau_{ZPT}^b, \tau_{ZPPT}^b\}$. For the proof of the NP-hardness of τ -SOLVABILITY and τ -ESSP we reduce φ to a union U_τ of gadget TS. The index τ emphasizes that the actual peculiarity of the union depends on τ . In particular, if $\tau \in \{\tau_{PT}^b, \tau_{PPT}^b\}$, then all these TS are linear, and $JL(U_\tau)$ is a well defined linear TS. Otherwise, if $\tau \in \{\tau_{ZPT}^b, \tau_{ZPPT}^b\}$, then these gadgets are 2-grade TS where no initial state has an incoming edge, which implies that $J(U_\tau)$ is a 2-grade TS.

In U_τ , the variables of φ are represented by events and the clauses of φ are represented by paths on which the variables of the clauses occur as events. More exactly, for every $i \in \{0, \dots, m-1\}$ and clause $\zeta_i = \{X_{i_0}, X_{i_1}, X_{i_2}\}$, the union U_τ contains (a gadget with) a directed labeled path $P_i = \dots \xrightarrow{X_{i_0}} \dots \xrightarrow{X_{i_1}} \dots \xrightarrow{X_{i_2}} \dots$ on which the variables X_{i_0}, X_{i_1} and X_{i_2} of ζ_i occur as events. Moreover, by construction, the union U_τ provides an ESSA α whose τ -solvability is connected with the existence of a one-in-three model of φ . In particular, we build the union U_τ in a way such that there is a subset $\mathfrak{E} \subseteq E_\tau$ of events of τ so that the following properties are satisfied: If $R = (sup, sig)$ is a τ -region of U_τ that solves α , then the variable events whose signature belongs to \mathfrak{E} define a one-in-three model of φ , that is, the set $M = \{X \in V(\varphi) \mid sig(X) \in \mathfrak{E}\}$ satisfies $|M \cap \zeta_i| = 1$ for all $i \in \{0, \dots, m-1\}$. Hence, if U_τ has the τ -ESSP, then α is τ -solvable and φ allows a positive decision. Moreover, the construction of U_τ ensures that if φ has a one-in-three model, then α as well as all the other ESSA and SSA of U_τ are τ -solvable. Thus, U_τ has the τ -ESSP if and only if φ is one-in-three satisfiable if and only if U_τ has both the τ -ESSP and the τ -SSP. Since Lemma 3.9 lifts these implications to the linear joining $LJ(U_\tau)$, if $\tau \in \{\tau_{PT}^b, \tau_{PPT}^b\}$, and to the joining $J(U_\tau)$, if $\tau \in \{\tau_{ZPT}^b, \tau_{ZPPT}^b\}$, this proves the NP-hardness of the τ -ESSP and τ -SOLVABILITY for accordingly restricted TS.

Let $\tau \in \{\tau_{PT}^b, \tau_{PPT}^b\}$. For the proof of the NP-hardness of τ -SSP we reduce φ to a union U of linear TS. Since this union is the same for both τ_{PT}^b and τ_{PPT}^b , U needs no index. Using essentially the same approach as just sketched, the union U provides an SSA α that is τ -solvable if and only if φ has a one-in-three model. Moreover, if α is τ -solvable, then U has the τ -SSP. Consequently, again by Lemma 3.9, this implies that $LJ(U)$ has the τ -SSP if and only if φ has a one-in-three model. This proves the NP-hardness of τ -SSP for linear inputs.

4.1. NP-hardness of τ_{PPT}^b -SOLVABILITY and τ_{PPT}^b -ESSP

In the remainder of this section, unless explicitly stated otherwise, let $\tau = \tau_{PPT}^b$. In the following, we first introduce the gadgets (TS) of the union U_τ and the atom α . Figure 5 presents a concrete example of $U_{\tau_{PPT}^b}$, where φ corresponds to Example 4.2. Secondly, we argue that these gadgets collaborate in a way such that if α is τ -solvable, then φ has a one-in-three model. Finally, we show that if φ is one-in-three satisfiable, then U_τ is τ -solvable.

The union U_τ has the following gadget H_1 that provides the announced ESSA $\alpha = (k, h_{1,2b+4})$:

$$\begin{array}{cccccccccccccccc}
 h_{1,0} & \xrightarrow{k} & \cdots & \xrightarrow{k} & h_{1,b} & \xrightarrow{y_0} & h_{1,b+1} & \xrightarrow{o_0} & h_{1,b+2} & \xrightarrow{k} & \cdots & \xrightarrow{k} & h_{1,2b+2} & \xrightarrow{y_1} & h_{1,2b+3} & \xrightarrow{y_0} & h_{1,2b+4} \\
 & & & & & & & & & & & & & & & & & \downarrow o_1 \\
 & & & & & & & & & & & & & & & & & h_{1,3b+5} \xleftarrow{k} \cdots \xleftarrow{k} h_{1,2b+5}
 \end{array}$$

For all $j \in \{0, 1, 2, 3\}$, the union U_τ has the following gadget D_j that provides the event k_j :

$$D_{j,1} = d_{j,0} \xrightarrow{o_0} d_{j,1} \xrightarrow{k_j} d_{j,2} \xrightarrow{o_1} d_{j,3}$$

For all $j \in \{0, \dots, 2m - 1\}$, the union U_τ has the following gadgets F_j and G_j that provide the event z_j :

$$F_j = f_{j,0} \xrightarrow{k_0} f_{j,1} \xrightarrow{z_j} f_{j,2} \qquad G_j = g_{j,0} \xrightarrow{z_j} g_{j,1} \xrightarrow{o_0} g_{j,2}$$

For all $i \in \{0, \dots, m - 1\}$, the union U_τ has the following gadget M_i , that uses the variable X_i as event:

$$M_i = m_{i,0} \xrightarrow{k_1} m_{i,1} \xrightarrow{X_i} \cdots \xrightarrow{X_i} m_{i,b+1}$$

For all $i \in \{0, \dots, m - 1\}$, the union U_τ has the following gadget T_i that uses the elements of $\zeta_i = \{X_{i_0}, X_{i_1}, X_{i_2}\}$ as events:

$$\begin{array}{cccccccccccccccc}
 T_i = & t_{i,0} & \xrightarrow{k_2} & t_{i,1} & \xrightarrow{X_{i_0}} & \cdots & \xrightarrow{X_{i_0}} & t_{i,b+1} & \xrightarrow{z_{2i}} & t_{i,b+2} & \xrightarrow{X_{i_1}} & \cdots & \xrightarrow{X_{i_1}} & t_{i,2b+2} \\
 & & & & & & & & & & & & & & \downarrow z_{2i+1} \\
 & & & & & & & & & & & & & & t_{i,3b+4} \xleftarrow{k_3} t_{i,3b+3} \xleftarrow{X_{i_2}} \cdots \xleftarrow{X_{i_2}} t_{i,2b+3}
 \end{array}$$

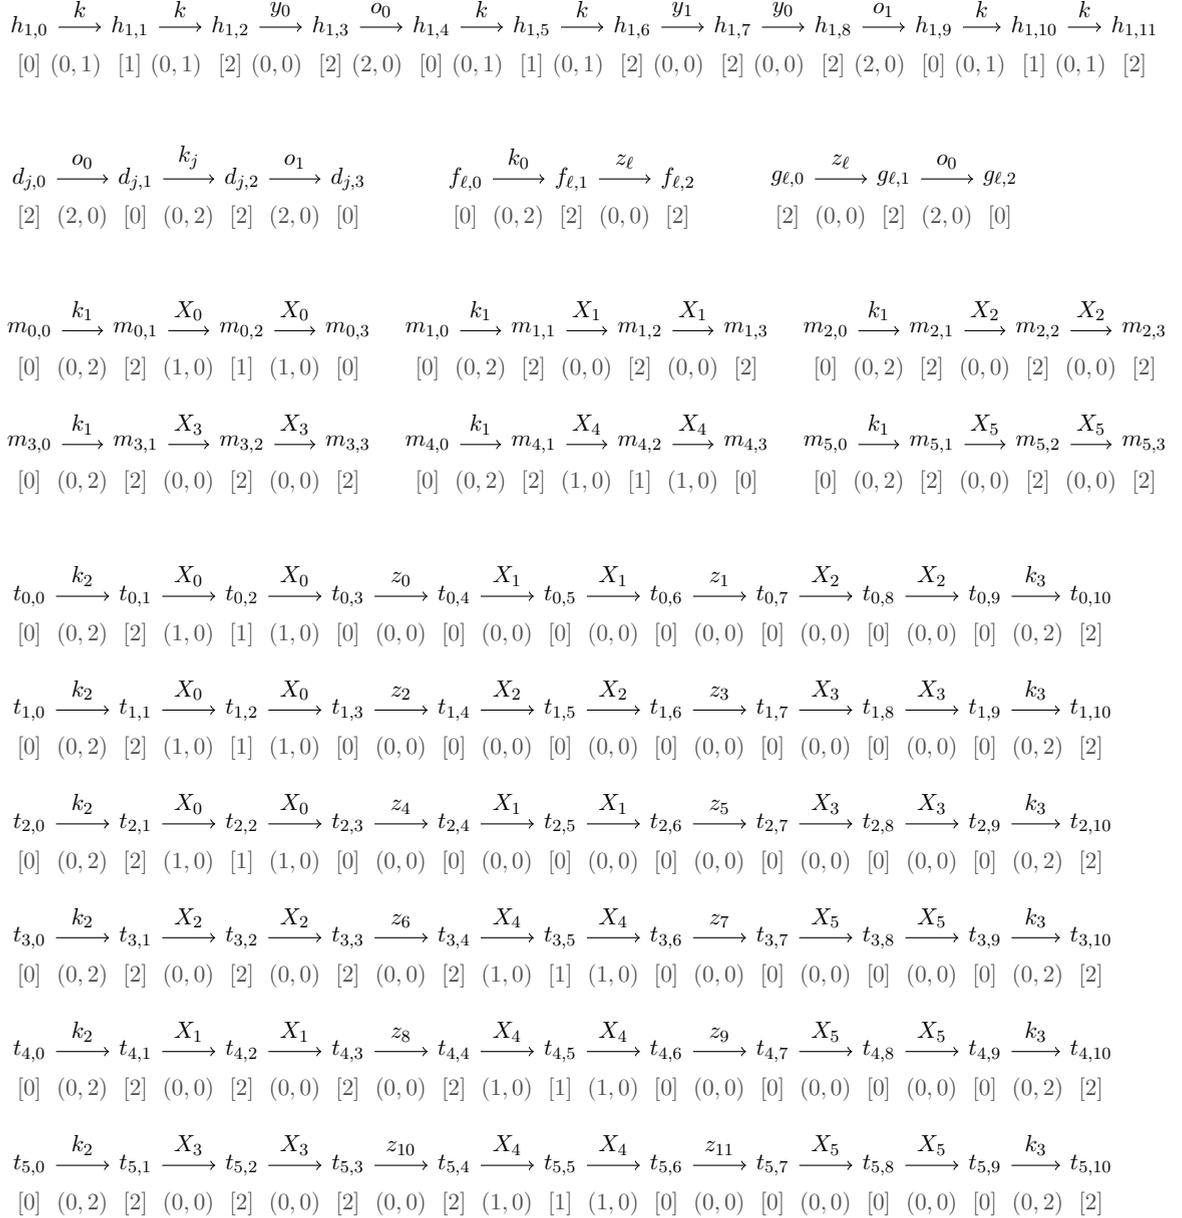


Figure 5: The gadgets of the union $U_{\tau_{PPT}^2}$ that originates from the input φ of Example 4.2; we assume $j \in \{0, 1, 2, 3\}$ and $\ell \in \{0, \dots, 11\}$. A number $[i]$ and a pair (k, ℓ) below a state s and event e define the support $sup(s) = i$ and the signature $sig(e) = (k, \ell)$ in correspondence to the region $R = (sup, sig)$, which is defined to prove the τ_{PPT}^2 -solvability of $\alpha = (k, h_{1,8})$.

Altogether,

$$U_\tau = U(H_1, D_0, \dots, D_3, F_0, \dots, F_{2m-1}, G_0, \dots, G_{2m-1}, M_0, \dots, M_{m-1}, T_0, \dots, T_{m-1}).$$

Lemma 4.3. If U_τ has the τ -ESSP, then φ has a one-in-three model.

Proof:

Since U_τ has the τ -ESSP, there is a τ -region that solves α . Let $R = (\text{sup}, \text{sig})$ be such a region. In the following we argue, that the set $\{X \in V(\varphi) \mid \text{sig}(X) = (0, 1)\}$ or the set $\{X \in V(\varphi) \mid \text{sig}(X) = (1, 0)\}$ is a one-in-three model of φ . Since R solves α , we have that $\text{sig}(k)$ does not occur at $\text{sup}(h_{1,2b+4})$. This implies $\text{sig}(k) \neq (0, 0)$. By Lemma 2.20, we get $\text{sig}(k) \in \{(1, 0), (0, 1)\}$. In what follows, we let $\text{sig}(k) = (0, 1)$ and show that $M = \{X \in V(\varphi) \mid \text{sig}(X) = (1, 0)\}$ defines a one-in-three model of φ . The arguments for the case $\text{sig}(k) = (1, 0)$ are quite similar and lead to the fact that $\{X \in V(\varphi) \mid \text{sig}(X) = (0, 1)\}$ defines a searched model.

Let $\text{sig}(k) = (0, 1)$ and $\neg \text{sup}(h_{1,2b+4}) \xrightarrow{\text{sig}(k)}$. We argue that this implies $\text{sig}(o_0) = \text{sig}(o_1) = (b, 0)$: For all $s \in \{0, \dots, b-1\}$, the event $(0, 1)$ occurs at s in τ . Since $\text{sig}(k)$ does not occur at $\text{sup}(h_{1,2b+4})$, this implies $\text{sup}(h_{1,2b+4}) = b$. Moreover, by $\text{sig}(k) = (0, 1)$ and Lemma 2.20, we get $\text{sup}(h_{1,b}) = b$ and $\text{sup}(h_{1,b+2}) = \text{sup}(h_{1,2b+5}) = 0$. By $\text{sup}(h_{1,2b+4}) = b$ and $\text{sup}(h_{1,2b+5}) = 0$, we obtain $\text{sig}(o_1) = (b, 0)$. Moreover, $\text{sup}(h_{1,b}) = b$ and $h_{1,b} \xrightarrow{y_0}$ imply $\text{sig}^+(y_0) = 0$, and by $\text{sup}(h_{1,2b+4}) = b$ and $h_{1,2b+4} \xrightarrow{y_0}$ imply $\text{sig}^-(y_0) = 0$. (Recall that R is pure.) Hence, $\text{sig}(y_0) = (0, 0)$, which implies $\text{sup}(h_{1,b+1}) = b$. Thus, by $\text{sup}(h_{1,b+1}) = b$ and $\text{sup}(h_{1,b+2}) = 0$, we obtain $\text{sig}(o_0) = (b, 0)$.

The gadgets D_0, \dots, D_3 use the signatures of o_0 and o_1 to determine the signatures of k_0, \dots, k_3 . More exactly, $\text{sig}(o_0) = \text{sig}(o_1) = (b, 0)$ implies $\text{sup}(d_{j,1}) = 0$ and $\text{sup}(d_{j,2}) = b$ for all $j \in \{0, 1, 2, 3\}$. Consequently, this implies $\text{sig}(k_0) = \dots = \text{sig}(k_3) = (0, b)$.

Let $j \in \{0, \dots, 2m-1\}$ be arbitrary but fixed. The gadgets F_j and G_j ensure that $\text{sig}(z_j) = (0, 0)$: By $\text{sig}(o_0) = (b, 0)$ and $\text{sig}(k_0) = (0, b)$, we get $\text{sup}(f_{j,1}) = b$ and $\text{sup}(g_{j,1}) = b$. Since R is pure, that is $\text{sig}^+(z_j) = 0$ or $\text{sig}^-(z_j) = 0$, by $f_{j,2} \xrightarrow{z_j}$, we get $\text{sig}^-(z_j) \geq \text{sig}^+(z_j)$. Similarly, by $g_{j,1} \xrightarrow{z_j}$, we get $\text{sig}^+(z_j) \geq \text{sig}^-(z_j)$. Consequently, $\text{sig}^-(z_j) = \text{sig}^+(z_j)$, which implies $\text{sig}(z_j) = (0, 0)$, since R is pure.

Let $i \in \{0, \dots, m-1\}$ be arbitrary but fixed. The gadget M_i ensures for X_i that $\text{sig}(X_i) \in \{(1, 0), (0, 0)\}$: By $\text{sig}(k_1) = (0, b)$, we have $\text{sup}(m_{i,1}) = b$, which implies $\text{sig}^-(X_i) \geq \text{sig}^+(X_i)$. Since X_i occurs b times in a row at $m_{i,1}$, by Lemma 2.20, this implies $\text{sig}(X_i) \in \{(1, 0), (0, 0)\}$.

The gadget T_i ensures that there is exactly one event $X \in \{X_{i_0}, X_{i_1}, X_{i_2}\}$ such that $\text{sig}(X) = (1, 0)$: By $\text{sig}(k_2) = \text{sig}(k_3) = (0, b)$, we have that $\text{sup}(t_{i,1}) = b$ and $\text{sup}(t_{i,3b+3}) = 0$. Consequently, the image of the sub-path $t_{i,1} \xrightarrow{X_{i_1}} \dots \xrightarrow{X_{i_2}} t_{i,3b+3}$ under (sup, sig) is a path of τ that starts at b and terminates at 0 . Hence, there is an event e on this path that satisfies $\text{sig}^-(e) > \text{sig}^+(e)$. Since $\text{sig}(z_{2i}) = \text{sig}(z_{2i+1}) = (0, 0)$, we obtain that $e \in \{X_{i_0}, X_{i_1}, X_{i_2}\}$. Moreover, since each of X_{i_0}, X_{i_1} and X_{i_2} occurs b times in a row, if $\text{sig}^-(e) > \text{sig}^+(e)$, then $\text{sig}(e) = (1, 0)$. In the following, we argue that if $e \in \{X_{i_0}, X_{i_1}, X_{i_2}\}$ such that $\text{sig}(e) = (1, 0)$, then $\text{sig}(e') \neq (1, 0)$ for all $e' \in \{X_{i_0}, X_{i_1}, X_{i_2}\} \setminus \{e\}$.

If $\text{sig}(X_{i_0}) = (1, 0)$, then we get $\text{sup}(t_{i,b+1}) = 0$, by Lemma 2.20. By $\text{sig}(z_{2i}) = (0, 0)$, this implies $\text{sup}(t_{i,b+2}) = 0$ and $\text{sig}^-(X_{i_1}) = 0$. Thus, by $\text{sig}(X_{i_1}) \in \{(1, 0), (0, 0)\}$, we conclude $\text{sig}(X_{i_1}) = (0, 0)$. By $\text{sup}(t_{i,b+2}) = 0$, $\text{sig}(X_{i_1}) = (0, 0)$ and $\text{sig}(z_{2i+1}) = (0, 0)$, we have that

$\text{sup}(t_{i,2b+3}) = 0$. This implies $\text{sig}^-(X_{i_2}) = 0$ and, thus, $\text{sig}(X_{i_2}) = (0, 0)$. In particular, we have $\text{sig}(X_{i_1}) \neq (1, 0)$ and $\text{sig}(X_{i_2}) \neq (1, 0)$.

If $\text{sig}(X_{i_1}) = (1, 0)$, then we get $\text{sup}(t_{i,b+1}) = b$ and $\text{sup}(t_{i,2b+3}) = 0$, by Lemma 2.20 and $\text{sig}(z_{2i}) = \text{sig}(z_{2i+1}) = (0, 0)$. By $\text{sup}(t_{i,b+1}) = b$, we get $\text{sig}(X_{i_0}) \neq (1, 0)$. Moreover, just like before, by $\text{sup}(t_{i,2b+3}) = 0$, we have $\text{sig}(X_{i_2}) \neq (1, 0)$.

If $\text{sig}(X_{i_2}) = (1, 0)$, then we get $\text{sig}(X_{i_0}) \neq (1, 0)$ and $\text{sig}(X_{i_1}) \neq (1, 0)$, since $\text{sig}(X_{i_0}) = (1, 0)$ or $\text{sig}(X_{i_1}) = (1, 0)$ imply $\text{sig}(X_{i_2}) \neq (1, 0)$, as just discussed.

Altogether, we have shown that if $R = (\text{sup}, \text{sig})$ is a τ -region that solves α such that $\text{sig}(k) = (0, 1)$, then, for all $i \in \{0, \dots, m-1\}$, there is exactly one event $e \in \{X_{i_0}, X_{i_1}, X_{i_2}\}$ that satisfies $\text{sig}(e) = (1, 0)$. As a result, the set $\{X \in V(\varphi) \mid \text{sig}(X) = (1, 0)\}$ defines a one-in-three model of φ . It is noteworthy that we use the pureness of τ only for the functionality of H_1 and (by the signature of o_1 , implicitly) for D_0, \dots, D_3 . That is, once we have that $\text{sig}(k_0) = \dots = \text{sig}(k_3) = (0, b)$ and $\text{sig}(o_0) = (b, 0)$, the arguments for the functionality of the remaining gadgets essentially work also for the (impure) b -bounded type τ_{PT}^b . The only difference then is that we can not conclude that $\text{sig}(z_j) = (0, 0)$, since $\text{sig}(z_j) = (m, m)$ would also be possible for τ_{PT}^b . The same is true for $e' \in \{X_{i_0}, X_{i_1}, X_{i_2}\} \setminus \{e\}$ if $e \in \{X_{i_0}, X_{i_1}, X_{i_2}\}$ such that $\text{sig}(e) = (1, 0)$. However, if $\text{sig}(e) = (m, m)$, then $s \xrightarrow{e} s'$ implies also $\text{sup}(s) = \text{sup}(s')$, and that is what actually matters in our arguments. Thus, we will reuse the corresponding gadgets for the type τ_{PT}^b .

If $\text{sig}(k) = (1, 0)$ and $\text{sup}(h_{1,2b+4}) = 0$, then one argues similarly that the set $\{X \in V(\varphi) \mid \text{sig}(X) = (0, 1)\}$ defines a one-in-three model of φ . Altogether, this shows that if U_τ has the τ -ESSP, which implies the τ -solvability of α , then φ has a one-in-three model. \square

For the opposite direction, we have to prove the following lemma:

Lemma 4.4. If φ has a one-in-three model, then U_τ has the τ -ESSP and the τ -SSP.

For the proof of Lemma 4.4 it is sufficient to show that if φ has a one-in-three model M , then U_τ has the τ -ESSP. Since all introduced gadgets are linear TS, by Lemma 2.21, this implies that U_τ has the τ -SSP, too. The brut-force approach of this proof would be to explicitly present for every ESSA of U_τ a τ -region that solves it. In fact, for some atoms of U_τ , we need to explicitly present regions that solve them. In particular, this applies to $(k, h_{1,2b+4})$. On the other hand, the gadgets and the events of U_τ meet some regularities that allow us to solve many events homogeneously. In the following, for the purpose to discover these regularities, we first introduce the notions of consistent and thinly distributed events. After that we present a lemma that uses these notions and exploits a certain structure of U_τ to solve most events uniformly.

Definition 4.5. (c-consistent)

Let $U = U(A_0, \dots, A_n)$ be a union, where $A_i = (S_i, E_i, \delta_i, \iota_i)$ is a linear TS for all $i \in \{0, \dots, n\}$, and let $c \in \mathbb{N}$. We say an event $e \in E(U)$ is c -consistent (in U), if the following condition is satisfied for all $i \in \{0, \dots, n\}$: if $s \xrightarrow{e} s' \in A_i$, then e occurs always exactly c times in a row in A_i , that is, there are states $s, s' \in \{s_0, \dots, s_c\} \subseteq S_i$ such that $s_0 \xrightarrow{e} \dots \xrightarrow{e} s_c$ and $\neg \xrightarrow{e} s_0$ and $\neg s_c \xrightarrow{e}$.

Definition 4.6. (thinly distributed)

Let $U = U(A_0, \dots, A_n)$ be a union, where $A_i = (S_i, E_i, \delta_i, \iota_i)$ is a linear TS for all $i \in \{0, \dots, n\}$, and let $e \in E(U)$ such that e is c -consistent for some $c \in \{1, b\}$. We say e is *thinly distributed* (in U) if the following condition is satisfied for all $i \in \{0, \dots, n\}$: if $e \in E_i$, then there is exactly one path (with pairwise distinct states) $s_0 \xrightarrow{e} \dots \xrightarrow{e} s_c$ in A_i .

Example 4.7. Every event of U_τ is either b -consistent as, for example, k and X_0, \dots, X_{m-1} , or 1-consistent as, for example, o_0 and o_1 . Moreover, the event o_1 occurs once at the edge $h_{1,2b+4} \xrightarrow{o_1} h_{1,2b+5}$ and, for all $j \in \{0, 1, 2, 3\}$, the it occurs once at the edge $d_{j,2} \xrightarrow{o_1} d_{j,3}$. No other gadget of U_τ applies o_1 . Thus, o_1 is thinly distributed in U_τ . Moreover, for all $i \in \{0, \dots, m-1\}$, if X_i occurs in a gadget of U_τ , then it occurs exactly once b -times in a row in this gadget. Hence, X_i is thinly distributed.

Lemma 4.8. Let $\tau \in \{\tau_{PT}^b, \tau_{PPT}^b\}$. Let $U = U(A_0, \dots, A_n)$ be a union, where $A_i = (S_i, E_i, \delta_i, \iota_i)$ is a linear TS for all $i \in \{0, \dots, n\}$, such that every event $e \in E(U)$ is 1-consistent or b -consistent, and let $a \in E(U)$ be a thinly distributed event and $q \in S_i$ a state such that $\neg \xrightarrow{a}$, where $i \in \{0, \dots, n\}$ is arbitrary but fixed. If one of the following conditions is satisfied, then there is a τ -region of U that solves (a, q) :

1. $a \notin E_i$ or $e \in E_i$ and q occurs after a ;
2. $a \in E_i$ and a occurs after q and there is an event $x \in E_i \setminus \{a\}$ such that $\xrightarrow{x} z \xrightarrow{a}$ in A_i and
 - (a) x is thinly distributed and
 - (b) for all $j \in \{0, \dots, n\}$, if $a, x \in E_j$, then x does not occur after a in A_j and
 - (c) if a is b -consistent, then x is 1-consistent.

Proof:

(1): The following τ -region $R = (sup, sig)$ solves (a, q) : For all $j \in \{0, \dots, m-1\}$, if $a \in E_j$, then $sup(\iota_j) = 0$, otherwise $sup(\iota_j) = b$; for all $e \in E(U)$, if $e = a$ and a is b -consistent, then $sig(e) = (0, 1)$; if $e = a$ and e is 1-consistent, then $sig(e) = (0, b)$; otherwise $sig(e) = (0, 0)$.

(2): The following τ -region $R = (sup, sig)$ solves (a, q) : for all $j \in \{0, \dots, n\}$, if $a \in E_j$ and $x \notin E_j$, then $sup(\iota_j) = 0$, otherwise $sup(\iota_j) = b$; for all $e \in E(A)$, if $e = a$, then $sig(e) = (0, 1)$ if a is b -consistent, else $sig(e) = (0, b)$; if $e = x$, then $sig(e) = (1, 0)$ if x is b -consistent, else $sig(e) = (b, 0)$; otherwise $sig(e) = (0, 0)$. \square

Armed with these results, we are now able to provide the proof of Lemma 4.4:

Proof:

[Lemma 4.4] Let M be a one-in-three model of φ . We proceed as follows. First, we apply Lemma 4.8 to solve most of U_τ 's ESSA. After that, we explicitly present τ -regions that solve the remaining atoms and, in particular, solve α . This proves that U_τ has the τ -ESSP and, by Lemma 2.21, the τ -SOLVABILITY, too.

Let $e \in E(U_\tau) \setminus \{k, y_0, y_1\}$, let G be a gadget of U_τ and let $s \in S(G)$ such that $\neg s \xrightarrow{e}$, where all of e, G and s are arbitrary but fixed. For a start, we notice that e is thinly distributed. Moreover, recall

that, for all $i \in \{0, \dots, m-1\}$, the clause $\zeta_i = \{X_{i_0}, X_{i_1}, X_{i_2}\}$ satisfies $i_0 < i_1 < i_2$. Consequently, if $e \notin \{o_0, o_1\}$ or if $e \in \{o_0, o_1\}$ and $G \neq H_1$, then e satisfies Condition 1 or Condition 2 of Lemma 4.8. Thus, by Lemma 4.8, the atom (e, s) is τ -solvable.

It remains to argue that the remaining atoms are also τ -solvable. For convenience, we let $I = \{\iota_G \mid G \text{ is a gadget of } U\}$ be the set of the initial states of the gadgets of U .

For a start, we argue for the solvability of k . The following region $R = (sup, sig)$ solves (k, s) for all relevant $s \in S(H_1)$; in particular, it solves $(k, h_{1,2b+4})$. Figure 5 presents a concrete example of R for the union U_τ that originates from φ of Example 4.2. Let's start with the support of the initial states: if $s \in \{h_{1,0}, f_{0,0}, \dots, f_{2m-1,0}, m_{0,0}, \dots, m_{m-1,0}, t_{0,0}, \dots, t_{m-1,0}\}$, then $sup(s) = 0$; if $s \in \{d_{0,0}, \dots, d_{3,0}, g_{0,0}, \dots, g_{2m-1,0}\}$, then $sup(s) = b$. The signature is defined as follows: for all $e \in E(U_\tau)$, if $e = k$, then $sig(e) = (0, 1)$; if $e \in \{o_0, o_1\}$, then $sig(e) = (b, 0)$; if $e \in \{k_0, \dots, k_3\}$, then $sig(e) = (0, b)$; if $e \in M$, then $sig(e) = (1, 0)$; otherwise $sig(e) = (0, 0)$.

The following region $R = (sup, sig)$ solves (k, s) for all other relevant states of U_τ : $sup(h_{1,0}) = 0$; for all $s \in I \setminus \{h_{1,0}\}$, $sup(s) = b$; for all $e \in E(U_\tau)$, if $e = k$, then $sig(e) = (0, 1)$; if $e = y_0$, then $sig(e) = (b, 0)$; otherwise $sig(e) = (0, 0)$. This proves the solvability of k .

In the following, we argue that (o_0, q) is solvable for all relevant $q \in S(H_1)$: The following region $R = (sup, sig)$ solves (o_0, s) for all $s \in \{h_{1,b+2}, \dots, h_{1,3b+5}\}$: for all gadgets $G \in U_\tau$, we define $sup(\iota_G) = 0$ for G 's the initial state ι_G ; for all $e \in E(U_\tau)$, if $e = o_0$, then $sig(e) = (0, b)$; otherwise $sig(e) = (0, 0)$.

The following region $R = (sup, sig)$ solves (o_0, s) for all $s \in \{h_{1,0}, \dots, h_{1,b}\}$: $sup(h_{1,0}) = b$; for all $s \in I \setminus \{h_{1,0}\}$, we define $sup(s) = 0$; for all $e \in E(U_\tau)$, if $e = o_0$, then $sig(e) = (0, b)$; if $e = y_0$, then $sig(e) = (b, 0)$; otherwise $sig(e) = (0, 0)$.

Similarly, one argues that (o_1, q) is solvable for all relevant $q \in S(H_1)$. So far, we have proven the solvability of all $e \in E(U_\tau) \setminus \{k, y_0, y_1\}$. It remains to argue for the solvability of y_0 and y_1 .

The following region $R = (sup, sig)$ solves (y_0, s) for all $s \in \{h_{1,0}, \dots, h_{1,b-1}\}$: for all $s \in I$, we define $sup(\iota_G) = 0$; for all $e \in E(U_\tau)$, if $e = y_0$, then $sig(e) = (b, 0)$; if $e = k$, then $sig(e) = (0, 1)$; otherwise $sig(e) = (0, 0)$.

The following region $R = (sup, sig)$ solves (y_0, s) for all $s \in S(H_1) \setminus \{h_{1,0}, \dots, h_{1,b-1}\}$: $sup(h_{i,0}) = b$ and for all $s \in I \setminus \{h_{1,0}\}$, we define $sup(s) = 0$; for all $e \in E(U_\tau)$, if $e = y_0$, then $sig(e) = (b, 0)$; if $e = y_1$, then $sig(e) = (0, b)$; otherwise $sig(e) = (0, 0)$.

It is easy to see, that y_1 is solvable. □

Altogether, since the construction of U_τ and thus A_τ is obviously polynomial, by Lemma 3.9, Lemma 4.3 and Lemma 4.4 and the NP-completeness of CM1IN33SAT, we have finally proven that τ_{PPT}^b -ESSP and τ_{PPT}^b -SOLVABILITY are NP-complete for all $b \in \mathbb{N}^+$.

4.2. NP-hardness of τ_{PPT}^b -SOLVABILITY and τ_{PPT}^b -ESSP

In the remainder of section, unless stated explicitly otherwise, we assume that $\tau = \tau_{PPT}^b$. The union U_τ has the following TS H_0 that provides the ESSA $\alpha = (k, h_{0,4b+1})$:

$$\begin{array}{cccccccccccccccc}
 H_0 = h_{0,0} & \xrightarrow{k} & \cdots & \xrightarrow{k} & h_{0,b} & \xrightarrow{z} & \cdots & \xrightarrow{z} & h_{0,2b} & \xrightarrow{o_0} & h_{0,2b+1} & \xrightarrow{k} & \cdots & \xrightarrow{k} & h_{0,3b+1} \\
 & & & & & & & & & & & & & & & \downarrow z \\
 & & & & & & & & & & & & & & & h_{0,6b+1} & \xleftarrow{k} & \cdots & \xleftarrow{k} & h_{0,5b+1} & \xleftarrow{o_1} & \cdots & \xleftarrow{o_1} & h_{0,4b+1} & \xleftarrow{z} & \cdots
 \end{array}$$

For every $j \in \{0, 1, 2, 3\}$, the union U_τ has the following gadget C_j that provides k_j :

$$C_j = c_{j,0} \xrightarrow{o_0} c_{j,1} \xrightarrow{k_j} c_{j,2} \xrightarrow{o_1} \cdots \xrightarrow{o_1} c_{j,b+2}$$

Finally, for all $j \in \{0, \dots, 2m - 1\}$ and for all $i \in \{0, \dots, m - 1\}$, the union U_τ has the gadgets F_j, G_j, M_i and T_i as defined in Section 4.1. Altogether,

$$U_\tau = U(H_0, C_0, \dots, C_3, F_0, \dots, F_{2m-1}, G_0, \dots, G_{2m-1}, M_0, \dots, M_{m-1}, T_0, \dots, T_{m-1}).$$

Lemma 4.9. If U_τ has the τ -ESSP, then φ has a one-in-three model.

Proof:

Since U_τ has the τ -ESSP, there is a τ -region of U_τ that solves α . Let $R = (sup, sig)$ be such a region. In the following, we argue that either $sig(k_0) = \dots = sig(k_3) = (0, b)$ or $sig(k_0) = \dots = sig(k_3) = (b, 0)$. As already argued at the end of the proof of Lemma 4.3, by the functionality of the remaining gadgets, this implies that $\{X \in V(\varphi) \mid sig(X) = (1, 0)\}$ or $\{X \in V(\varphi) \mid sig(X) = (0, 1)\}$ is a one-in-three model of φ .

Let $E_0 = \{(m, m) \mid 0 \leq m \leq b\}$. By definition, if $sig(k) = (m, m) \in E_0$ then $sup(h_{0,3b+1}) \geq m$ and $sup(h_{0,5b+1}) \geq m$. Event (m, m) occurs at every state $s \in S_{\tau_{PT}^b}$ that satisfies $s \geq m$.

Hence, by $\neg h_{0,4b+1} \xrightarrow{(m,m)}$, we get $sup(h_{0,4b+1}) < m$. Since $sup(h_{0,3b+1}) \geq m$ and $sup(h_{0,4b+1}) < m$, we have $sig^-(z) > sig^+(z)$. Observe, that z is b -consistent. Thus, by Lemma 2.20, we have $sup(z) = (1, 0)$. Similarly, we get $sig(o_1) = (0, 1)$. This immediately implies $sup(h_{0,2b}) = 0$ and $sup(h_{0,3b+1}) = b$. Moreover, by $sig(k) = (m, m)$ and $sup(h_{0,3b+1}) = b$ we get $sup(h_{0,2b+1}) = b$. By $sup(h_{0,2b}) = 0$, this implies $sig(o_0) = (0, b)$. Thus, we have $sig(o_0) = (0, b)$ and $sig(o_1) = (0, 1)$.

Otherwise, if $sig(k) \notin E_0$, then Lemma 2.20 ensures $sig(k) \in \{(1, 0), (0, 1)\}$. If $sig(k) = (0, 1)$ then we have $sup(h_{0,4b+1}) = b$, since $s \xrightarrow{(0,1)}$ for every state $s \in \{0, \dots, b - 1\}$ of τ_{PT}^b . Moreover, again by $sig(k) = (0, 1)$ we have $sup(h_{0,b}) = sup(h_{0,3b+1}) = b$ and $sup(h_{0,2b+1}) = sup(h_{0,5b+1}) = 0$. By $sup(h_{0,3b+1}) = sup(h_{0,4b+1}) = b$ we have $sig(z) \in E_0$, which together with $sup(h_{0,b}) = b$ implies $sup(h_{0,2b}) = b$. Thus, by $sup(h_{0,2b}) = b$ and $sup(h_{0,2b+1}) = 0$, it is $sig(o_0) = (b, 0)$. Moreover, by $sup(h_{0,4b+1}) = b$ and $sup(h_{0,5b+1}) = 0$, we conclude $sig(o_1) = (1, 0)$. Hence, we have $sig(o_0) = (b, 0)$ and $sig(o_1) = (1, 0)$. Similar arguments show that $sig(k) = (1, 0)$ implies $sig(o_0) = (0, b)$ and $sig(o_1) = (0, 1)$.

So far we have argued that if (sup, sig) is a τ_{PT}^b -region of U_τ that solves α , then either $sig(o_0) = (0, b)$ and $sig(o_1) = (0, 1)$ or $sig(o_0) = (b, 0)$ and $sig(o_1) = (1, 0)$. One easily finds out that if $sig(o_0) = (0, b)$ and $sig(o_1) = (0, 1)$, then $sup(c_{j,1}) = b$ and $sup(c_{j,2}) = 0$ and thus $sig(k_j) = (b, 0)$ for all $j \in \{0, \dots, 3\}$. Similarly, if $sig(o_0) = (b, 0)$ and $sig(o_1) = (1, 0)$, then $sup(c_{j,1}) = 0$,

$sup(c_{j,2}) = b$ and $sig(k_j) = (0, b)$ for all $j \in \{0, \dots, 3\}$. By the functionality of F_j, G_j this implies $z_j \in E_0$. Moreover, by the functionality of M_i , this implies if $sig(k_1) = (0, b)$, then $sig(X_i) \in \{(1, 0)\} \cup E_0$ and if $sig(k_1) = (b, 0)$, then $sig(X_i) \in \{(0, 1), (0, 0)\}$ for all $i \in \{0, \dots, m-1\}$. Similar to the arguments for τ_{PPT}^b , one argues that the gadgets T_0, \dots, T_{m-1} then ensure that $\{e \in V(\varphi) \mid sig(e) = (0, 1)\}$ or $\{e \in V(\varphi) \mid sig(e) = (1, 0)\}$ defines a sought model of φ . Thus, if U_τ has the τ -ESSP or is τ -solvable, which implies that α is τ -solvable, then φ has a one-in-three model. \square

The following lemma is dedicated to the opposite direction:

Lemma 4.10. If φ has a one-in-three model, then U_τ has the τ -ESSP and the τ -SSP.

Proof:

In the following, we argue that if M is a one-in-three model of φ , then U_τ has the τ -ESSP and thus has also the τ -SSP, since all gadgets are linear. Notice that if e is an event and G is a gadget of U_τ such that e does not occur in G , then (e, s) is τ -solvable for all $s \in S(G)$. A solving region $R = (sup, sig)$ is defined as follows: for all gadgets G' of U_τ , if $e \in E(G')$, then $sup(\iota_{G'}) = b$, otherwise $sup(\iota_{G'}) = 0$; for all events $e' \in E(U_\tau)$, if $e' = e$, then $sig(e') = (b, b)$; otherwise $sig(e') = (0, 0)$. Thus, in the following, we only argue for valid atoms (e, s) where e and s occur in the same gadget.

Let $e \in E(U_\tau) \setminus \{k, z\}$, let G be a gadget of U_τ and let $s \in S(G)$ such that $\neg s \xrightarrow{e}$, where all of e, G and s are arbitrary but fixed. The event e is thinly distributed. Moreover, if $e \notin \{o_0, o_1\}$ or if $e \in \{o_0, o_1\}$ and $G \neq H_0$, then e satisfies Condition 1 or Condition 2 of Lemma 4.8. Thus, by Lemma 4.8, in these cases, the atom (e, s) is τ -solvable.

For convenience, let $I = \{\iota_G \mid G \text{ is a gadget of } U_\tau\}$ be the set of the initial states of the gadgets of U_τ .

To complete the proof for the solvability of o_0 and o_1 , it remains to argue that (o_0, s) and (o_1, s') are solvable for all relevant $s, s' \in S(H_0)$: By Lemma 4.8.1, the atoms (o_0, s) and (o_1, s') are solvable for all $s \in \{h_{0,2b+1}, \dots, h_{0,6b+1}\}$ and for all $s' \in \{h_{0,5b+1}, \dots, h_{0,6b+1}\}$. The following region $R = (sup, sig)$ solves (o_0, s) for all $s \in \{h_{0,0}, \dots, h_{0,2b-1}\}$: $sup(h_{0,0}) = b$; for all $s \in I \setminus \{h_{0,0}\}$, we define $sup(s) = 0$; for all $e \in E(U_\tau)$, if $e = o_0$, then $sig(e) = (0, b)$; if $e = z$, then $sig(e) = (1, 0)$; otherwise $sig(e) = (0, 0)$.

The following region $R = (sup, sig)$ solves (o_1, s) for all $s \in \{h_{0,0}, \dots, h_{0,4b}\} \setminus \{h_{0,2b}\}$ and uses the model M of φ : We start with the support of the initial states: $sup(h_{0,0}) = 0$; if $s \in \{f_{0,0}, \dots, f_{2m-1,0}, m_{0,0}, \dots, m_{m-1,0}, t_{0,0}, \dots, t_{m-1,0}\}$, then $sup(s) = 0$; if $s \in \{c_{0,0}, \dots, c_{3,0}\} \cup \{g_{0,0}, \dots, g_{2m-1,0}\}$, then $sup(s) = b$. The signature is defined as follows: for all $e \in E(U_\tau)$, if $e = o_1$, then $sig(e) = (b, b)$; if $e = z$, then $sig(e) = (0, 1)$; if $e = o_0$, then $sig(e) = (b, 0)$; if $e \in \{k_0, \dots, k_3\}$, then $sig(e) = (0, b)$; if $e \in M$, then $sig(e) = (1, 0)$; otherwise $sig(e) = (0, 0)$.

The following region $R = (sup, sig)$ solves $(o_1, h_{0,2b})$: $sup(s) = 0$ for all $s \in I$; for all $e \in E(U_\tau)$, if $e = o_1$, then $sig(o_1) = (b, b)$; if $e = o_0$, then $sig(e) = (0, b)$; otherwise $sig(e) = (0, 0)$. This proves the solvability of o_1 .

Since z occurs only in H_0 , for the solvability of z , it remains to argue that (z, s) is τ -solvable for all relevant $s \in S(H_0)$. The following region $R = (sup, sig)$ does this for all $s \in S(H_0) \setminus \{h_{0,6b+1}\}$ and uses the model M of φ . Moreover, this region also solves (k, s) for all $s \in S(H_0)$ and, thus,

proves the solvability of k : if $s \in \{h_{0,0}, f_{0,0}, \dots, f_{2m-1,0}, m_{0,0}, \dots, m_{m-1,0}, t_{0,0}, \dots, t_{m-1,0}\}$, then $\text{sup}(s) = 0$; if $s \in \{c_{0,0}, \dots, c_{3,0}, g_{0,0}, \dots, g_{2m-1,0}\}$, then $\text{sup}(s) = b$; for all $e \in E(U_\tau)$, if $e = z$, then $\text{sig}(z) = (b, b)$; if $e = k$, then $\text{sig}(e) = (0, 1)$; if $e \in \{o_0, o_1\}$, then $\text{sig}(e) = (b, 0)$; if $e \in \{k_0, \dots, k_3\}$, then $\text{sig}(e) = (0, b)$; if $e \in M$, then $\text{sig}(e) = (1, 0)$; otherwise $\text{sig}(e) = (0, 0)$.

One easily finds that $(z, h_{0,6b+1})$ is τ -solvable. Altogether, this proves that if φ is one-in-three satisfiable, then U_τ has the τ -ESSP. Since all gadgets are linear, this completes the proof. \square

4.3. NP-hardness of τ_{PPT}^b -SSP and τ_{PT}^b -SSP

In the remainder of this section, unless stated explicitly otherwise, let $\tau \in \{\tau_{PPT}^b, \tau_{PT}^b\}$ be arbitrary but fixed. The union U_τ has the following gadget H_2 that provides the atom $\alpha = (h_{2,0}, h_{2,b})$:

$$H_2 = h_{2,0} \xrightarrow{k} \dots \xrightarrow{k} h_{2,b} \xrightarrow{o_0} h_{2,b+1} \xrightarrow{k} \dots \xrightarrow{k} h_{2,2b+1} \xrightarrow{o_2} h_{2,2b+2} \xrightarrow{k} \dots \xrightarrow{k} h_{2,3b+2}$$

Moreover, the union U_τ has every gadget that has been defined for $U_{\tau_{PPT}^b}$ in Section 4.1 except for H_1 . Altogether, U_τ is defined as follows:

$$U_\tau = U(H_2, D_0, \dots, D_3, F_0, \dots, F_{2m-1}, G_0, \dots, G_{2m-1}, M_0, \dots, M_{m-1}, T_0, \dots, T_{m-1}).$$

Lemma 4.11. If U_τ has the τ -SSP, then φ has a one-in-three model.

Proof:

Since U_τ has the τ -SSP, there is a τ -region that solves α . Let $R = (\text{sup}, \text{sig})$ be such a region. We argue that the signature of the variable events define a sought model of φ : The event k occurs b times in a row at $h_{2,0}$. Thus, by Lemma 2.20, a region (sup, sig) solving $(h_{2,0}, h_{2,b})$ satisfies $\text{sig}(k) \in \{(1, 0), (0, 1)\}$. If $\text{sig}(k) = (1, 0)$, then $\text{sup}(h_{2,b}) = \text{sup}(h_{2,2b+1}) = b$ and $\text{sup}(h_{2,b+1}) = \text{sup}(h_{2,2b+2}) = 0$. This implies $\text{sig}(o_0) = \text{sig}(o_2) = (b, 0)$ and, thus, $\text{sig}(k_j) = (0, b)$ for all $j \in \{0, \dots, 3\}$. Otherwise, if $\text{sig}(k) = (0, 1)$ then $\text{sup}(h_{2,b}) = \text{sup}(h_{2,2b+1}) = 0$ and $\text{sup}(h_{2,b+1}) = \text{sup}(h_{2,2b+2}) = b$. This implies $\text{sig}(o_0) = \text{sig}(o_2) = (0, b)$ and $\text{sig}(k_j) = (b, 0)$ for all $j \in \{0, \dots, 3\}$. Just like before, this proves the one-in-three satisfiability of φ . \square

The following lemma addresses the opposite direction:

Lemma 4.12. If φ has a one-in-three model, then U_τ has the τ -SSP.

Proof:

Let M be a one-in-three model of φ . We briefly argue, that U_τ has the τ -SSP. For start, let $e \in E(U_\tau) \setminus \{k\}$ be arbitrary but fixed. The event e is thinly distributed. Moreover, if $s \in S(U_\tau) \setminus S(H_2)$ and $\neg s \xrightarrow{e}$, then, by Lemma 4.8, (e, s) is τ -solvable. By Lemma 2.21, this implies that if (s, s') is an SSA of U_τ such that $s, s' \notin S(H_2)$, then (s, s') is τ -solvable. Thus, it remains to show that any SSA (s, s') of U_τ where $s, s' \in S(H_2)$ is τ -solvable, too. The corresponding regions can be defined similar to those from Section 4.1 and Section 4.2. In particular, the atom $(h_{2,0}, h_{2,b})$ can be solved by a region

that is defined in accordance to the region $R = (sup, sig)$ of Section 4.1 that solves $(k, h_{1,2b+4})$; one simply has to replace $sup(h_{1,0}) = 0$ by $sup(h_{2,0}) = 0$ and to ignore the events y_0 and y_1 . The resulting region also solves (s, s') if $s \neq s' \in \{h_{2,0}, \dots, h_{2,b}\}$ or $s \neq s' \in \{h_{2,b+1}, \dots, h_{2,2b+1}\}$ or $s \neq s' \in \{h_{2,2b+2}, \dots, h_{2,3b+2}\}$. Finally, it is easy to see that all states of $\{h_{2,0}, \dots, h_{2,b}\}$ are separable from all states of $\{h_{2,b+1}, \dots, h_{2,2b+1}\} \cup \{h_{2,2b+2}, \dots, h_{2,3b+2}\}$, and that all states of $\{h_{2,b+1}, \dots, h_{2,2b+1}\}$ are separable from all states of $\{h_{2,2b+2}, \dots, h_{2,3b+2}\}$. Altogether, this proves that if M has a one-in-three model, then U_τ has the τ -SSP. \square

4.4. NP-hardness of τ -SOLVABILITY and τ -ESSP for $\tau = \tau_{\mathbb{Z}PPT}$ and $\tau = \tau_{\mathbb{Z}PT}$

In the remainder of this section, unless stated explicitly otherwise, let $\tau \in \{\tau_{\mathbb{Z}PT}^b, \tau_{\mathbb{Z}PPT}^b\}$ and let $E_0 = \{(m, m) \mid 1 \leq m \leq b\} \cup \{0\}$. The union U_τ has the following TS H_3 that provides the atom $\alpha = (k, h_{3,1,b-1})$:

$$H_3 = \begin{array}{ccccccc} h_{3,0,0} & \xrightarrow{k} & \dots & \xrightarrow{k} & h_{3,0,b-1} & \xrightarrow{k} & h_{3,0,b} \\ & & & & & \nearrow z & \\ u \downarrow & & & & & & \\ h_{3,1,0} & \xrightarrow{k} & \dots & \xrightarrow{k} & h_{3,1,b-1} & & \end{array}$$

Moreover, for all $j \in \{0, \dots, m-1\}$, the union U_τ has the following gadgets F_j and G_j that use the variable X_j as event:

$$F_j = \begin{array}{ccccccc} f_{j,0,0} & \xrightarrow{k} & \dots & \xrightarrow{k} & f_{j,0,b-1} & \xrightarrow{k} & f_{j,0,b} \\ & & & & & \nearrow X_j & \\ v_j \downarrow & & & & & & \\ f_{j,1,0} & \xrightarrow{k} & \dots & \xrightarrow{k} & f_{j,1,b-1} & & \end{array} \quad G_j = g_{j,0} \xrightarrow{k} \dots \xrightarrow{k} g_{j,b} \xrightarrow{X_j} g_{j,b+1}$$

Finally, for all $i \in \{0, \dots, m-1\}$, the union U_τ has the following gadget T_i that uses the variables of the clause $\zeta_i = \{X_{i_0}, X_{i_1}, X_{i_2}\}$ as events:

$$T_i = t_{i,0} \xrightarrow{k} \dots \xrightarrow{k} t_{i,b} \xrightarrow{X_{i_0}} t_{i,b+1} \xrightarrow{X_{i_1}} t_{i,b+2} \xrightarrow{X_{i_2}} t_{i,b+3} \xrightarrow{z} t_{i,b+4} \xrightarrow{k} \dots \xrightarrow{k} t_{i,2b+4}$$

Altogether,

$$U_\tau = (H_3, F_0, G_0, \dots, F_{m-1}, G_{m-1}, T_0, \dots, T_{m-1}).$$

Lemma 4.13. If U_τ has the τ -ESSP, then φ has a one-in-three model.

Proof:

Since U_τ has the τ -ESSP, there is a τ -region, that solves α . Let $R = (sup, sig)$ be such a region. In the following, we first argue that $sig(k) \in \{(1, 0), (0, 1)\}$ and $sig(z) \in E_0$. Secondly, we show that this implies that $M = \{X \in V(\varphi) \mid sig(X) = 1\}$ is a one-in-three model of φ .

Let (sup, sig) be a τ -region that solves α , that is, $\neg sup(h_{3,1,b-1}) \xrightarrow{sig(k)}$. If $sig(k) \in E_0$, then we inductively obtain $sup(h_{3,1,0}) = sup(h_{3,1,b-1})$. This contradicts $\neg sup(h_{3,1,b-1}) \xrightarrow{sig(k)}$. Moreover, if $e \in \{0, \dots, b\}$, then $s \xrightarrow{e}$ for all $s \in S_\tau$. Consequently, we have $sig(k) \notin E_0 \cup \{1, \dots, b\}$.

The event k occurs b times in a row. Therefore, by Lemma 2.20, we get $\text{sig}(k) \in \{(1, 0), (0, 1)\}$. Moreover, if $\text{sig}(k) = (1, 0)$, then $\text{sup}(h_{3,0,b}) = 0$ and if $\text{sig}(k) = (0, 1)$, then $\text{sup}(h_{3,0,b}) = b$. If $s \in \{0, \dots, b-1\}$ then $s \xrightarrow{(0,1)}$ is true, and if $s \in \{1, \dots, b\}$, then $s \xrightarrow{(1,0)}$ is true. Consequently, by $\neg \text{sup}(h_{3,1,b-1}) \xrightarrow{\text{sig}(k)}$, if $\text{sig}(k) = (0, 1)$, then $\text{sup}(h_{3,1,b-1}) = b$, and if $\text{sig}(k) = (1, 0)$, then $\text{sup}(h_{3,1,b-1}) = 0$. For both cases, this implies $\text{sup}(h_{3,0,b}) = \text{sup}(h_{3,1,b-1})$ and thus $\text{sig}(z) \in E_0$.

We now argue that this makes M a one-in-three model of φ . Let $i \in \{0, \dots, m-1\}$ be arbitrary but fixed. By the definition of τ -regions, if p_i is defined by

$$p_i = \text{sup}(t_{i,b}) \xrightarrow{\text{sig}(X_{i_0})} \text{sup}(t_{i,b+1}) \xrightarrow{\text{sig}(X_{i_1})} \text{sup}(t_{i,b+2}) \xrightarrow{\text{sig}(X_{i_2})} \text{sup}(t_{i,b+3})$$

then p_i is a directed labeled path in τ . By $\text{sig}(z) \in E_0$ and $t_{i,b+3} \xrightarrow{z} t_{i,b+4}$ we obtain that $\text{sup}(t_{i,b+3}) = \text{sup}(t_{i,b+4})$. Moreover, k occurs b times in a row at $t_{i,0}$ and $t_{i,b+4}$. By Lemma 2.20, this implies if $\text{sig}(k) = (0, 1)$, then $\text{sup}(t_{i,b}) = b$ and $\text{sup}(t_{i,b+4}) = 0$. Similarly, and if $\text{sig}(k) = (1, 0)$, then $\text{sup}(t_{i,b}) = 0$ and $\text{sup}(t_{i,b+4}) = b$. Altogether, we obtain that the following conditions are true: If $\text{sig}(z) \in E_0$ and $\text{sig}(k) = (1, 0)$, then the path p_i starts at 0 and terminates at b , and if $\text{sig}(z) \in E_0$ and $\text{sig}(k) = (0, 1)$, then the path p_i starts at b and terminates at 0. In particular, both cases imply that there has to be at least one event $X \in \{X_{i_0}, X_{i_1}, X_{i_2}\}$ whose signature satisfies $\text{sig}(X) \notin E_0$. Via the functionality of the gadgets $F_0, G_0, \dots, F_{m-1}, G_{m-1}$, our reduction ensures that X is unique. More exactly, the aim of $F_0, G_0, \dots, F_{m-1}, G_{m-1}$ is to restrict the possible signatures for the variable events as follows:

- If $\text{sig}(k) = (1, 0)$, then $X \in V(\varphi)$ implies $\text{sig}(X) \in E_0 \cup \{b\}$, and
- if $\text{sig}(k) = (0, 1)$, then $X \in V(\varphi)$ implies $\text{sig}(X) \in E_0 \cup \{1\}$.

Before we argue that $F_0, G_0, \dots, F_{m-1}, G_{m-1}$ satisfy the announced functionality, we first argue that these restrictions of the signature of $X_{i_0}, X_{i_1}, X_{i_2}$ ensure that there is exactly one variable event $X \in \{X_{i_0}, X_{i_1}, X_{i_2}\}$ with $\text{sig}(X) \notin E_0$. Remember that, by definition, if $\text{sig}(X) \in E_0$ then $\text{sig}^-(X) + \text{sig}^+(X) = |\text{sig}(X)| = 0$.

For a start, let $\text{sig}(z) \in E_0$ and $\text{sig}(k) = (1, 0)$, which implies that p_i starts at 0 and terminates at b . Moreover, assume $\text{sig}(X) \in E_0 \cup \{b\}$. By Lemma 2.20, we obtain

$$(|\text{sig}(X_{i_0})| + |\text{sig}(X_{i_1})| + |\text{sig}(X_{i_2})|) \equiv b \pmod{b+1} \quad (1)$$

If $\text{sig}(X_{i_0}), \text{sig}(X_{i_1}), \text{sig}(X_{i_2}) \in E_0$, then $|\text{sig}(X_{i_0})| = |\text{sig}(X_{i_1})| = |\text{sig}(X_{i_2})| = 0$. This contradicts Equation 1. Hence, there has to be at least one variable event $X \in \{X_{i_0}, X_{i_1}, X_{i_2}\}$ such that $\text{sig}(X) = b$. In the following, we argue that X is unique.

Assume, for a contradiction, that there are two different variable events $X, Y \in \{X_{i_0}, X_{i_1}, X_{i_2}\}$ such that $\text{sig}(X) = \text{sig}(Y) = b$ and that $\text{sig}(Z) \in E_0$ for $Z \in \{X_{i_0}, X_{i_1}, X_{i_2}\} \setminus \{X, Y\}$. By symmetry and transitivity, we obtain

$$b \equiv (|sig(X_{i_0})| + |sig(X_{i_1})| + |sig(X_{i_2})|) \bmod (b + 1) \quad |(1) \quad (2)$$

$$(|sig(X_{i_0})| + |sig(X_{i_1})| + |sig(X_{i_2})|) \equiv 2b \bmod (b + 1) \quad |assumpt. \quad (3)$$

$$b \equiv 2b \bmod (b + 1) \quad |(2), (3) \quad (4)$$

$$2b \equiv (b - 1) \bmod (b + 1) \quad |def. \equiv \quad (5)$$

$$b \equiv (b - 1) \bmod (b + 1) \quad |(4), (5) \quad (6)$$

$$\exists m \in \mathbb{Z} : m(b + 1) = 1 \quad |(6) \quad (7)$$

By Equation 7, we get $b = 0$, a contradiction. Similarly, if we assume that $|sig(X_{i_0})| = |sig(X_{i_1})| = |sig(X_{i_2})| = b$, then we obtain

$$(|sig(X_{i_0})| + |sig(X_{i_1})| + |sig(X_{i_2})|) \equiv 3b \bmod (b + 1) \quad |assumpt. \quad (8)$$

$$b \equiv 3b \bmod (b + 1) \quad |(2), (8) \quad (9)$$

$$3b \equiv (b - 2) \bmod (b + 1) \quad |def. \equiv \quad (10)$$

$$b \equiv (b - 2) \bmod (b + 1) \quad |(9), (10) \quad (11)$$

$$\exists m \in \mathbb{Z} : m(b + 1) = 2 \quad |(11) \quad (12)$$

By Equation 12, we have $b \in \{0, 1\}$, which contradicts $b \geq 2$. Consequently, if $sig(z) \in E_0$ and $sig(k) = (1, 0)$ and $sig(X) \in E_0 \cup \{b\}$, then there is exactly one variable event $X \in \{X_{i_0}, X_{i_1}, X_{i_2}\}$ with $sig(X) \notin E_0$.

Otherwise, if $sig(z) \in E_0$, $sig(k) = (0, 1)$, implying that p_i starts at b and terminates at 0, and $sig(X) \in E_0 \cup \{1\}$, then the following equation is true:

$$(b + |sig(X_{i_0})| + |sig(X_{i_1})| + |sig(X_{i_2})|) \equiv 0 \bmod (b + 1) \quad (13)$$

This implies $|sig(X_{i_0})| + |sig(X_{i_1})| + |sig(X_{i_2})| \equiv 1 \bmod (b + 1)$. If there is more than one $X \in \{X_{i_0}, X_{i_1}, X_{i_2}\}$ such that $sig(X) = 1$, then $2 \equiv 1 \bmod (b + 1)$ or $3 \equiv 1 \bmod (b + 1)$ is true. If $2 \equiv 1 \bmod (b + 1)$, then $b = 0$, and if $3 \equiv 1 \bmod (b + 1)$, then $b \in \{0, 1\}$. Since $b \geq 2$, both cases yield a contradiction. Consequently, there is exactly one $X \in \{X_{i_0}, X_{i_1}, X_{i_2}\}$ such that $sig(X) = 1$, and if $Y \in \{X_{i_0}, X_{i_1}, X_{i_2}\} \setminus \{X\}$, then $sig(Y) \in E_0$.

Under the assumption that the gadgets $F_0, G_0, \dots, F_{m-1}, G_{m-1}$ behave as announced, we have shown the following: If (sup, sig) is a τ -region of U_τ such that $sig(k) \in \{(0, 1), (1, 0)\}$ and $sig(z) \in E_0$, then, for every $i \in \{0, \dots, m - 1\}$, there is exactly one variable event $X \in \{X_{i_0}, X_{i_1}, X_{i_2}\}$ such that $sig(X) \notin E_0$. As a result, the set $M = \{X \in V(\varphi) | sig(X) \notin E_0\}$ defines a one-in-three model of φ .

It remains to argue that the gadgets $F_0, G_0, \dots, F_{m-1}, G_{m-1}$ behave as announced. Let $j \in \{0, \dots, m - 1\}$. In the following, we show that if $sig(k) = (1, 0)$, then $sig(X_j) \in E_0 \cup \{b\}$, and if $sig(k) = (0, 1)$, then $sig(X_j) \in E_0 \cup \{1\}$.

To begin with, let $sig(k) = (1, 0)$. The event k occurs b times in a row at $f_{j,0,0}$ and $g_{j,0}$ and $b - 1$ times in a row at $f_{j,1,0}$. By Lemma 2.20 this implies $sup(f_{j,0,b}) = sup(g_{j,b}) = 0$ and $sup(f_{j,1,b-1}) \in \{0, 1\}$. Clearly, if $sup(f_{j,0,b}) = sup(f_{j,1,b-1}) = 0$ then $sig(X_j) \in E_0$. We argue that $sup(f_{j,1,b-1}) = 1$ implies $sig(X_j) = b$.

Assume, for a contradiction, that $\text{sig}(X_j) \neq b$. If $\text{sig}(X_j) = (m, m)$ for some $m \in \{1, \dots, b\}$, then $-\text{sig}^-(X_j) + \text{sig}^+(X_j) = |\text{sig}(X_j)| = 0$. By Lemma 2.20, this contradicts $\text{sup}(f_{j,0,b}) \neq \text{sup}(f_{j,1,b-1})$. If $\text{sig}(X_j) = (m, n)$ with $m \neq n$, then $|\text{sig}(X_j)| = 0$. By Lemma 2.20, we have $\text{sup}(f_{j,0,b}) = \text{sup}(f_{j,1,b-1}) - \text{sig}^-(X_j) + \text{sig}^+(X_j)$, which implies $\text{sig}(X_j) = (1, 0)$. But, this contradicts $\text{sup}(g_{j,b}) \xrightarrow{\text{sig}(X_j)} \text{sup}(g_{j,b}) = 0$ and $-0 \xrightarrow{(1,0)}$, in τ . Finally, if $\text{sig}(X_j) = e \in \{0, \dots, b-1\}$, then we have $1+e \not\equiv 0 \pmod{b+1}$. This contradicts $\text{sup}(f_{j,1,b-1}) \xrightarrow{\text{sig}(X_j)} \text{sup}(f_{j,0,b})$. Hence, we have $\text{sig}(X_j) = b$. Overall, it is proven that if $\text{sig}(k) = (1, 0)$, then $\text{sig}(X_j) \in E_0 \cup \{b\}$.

To continue, let $\text{sig}(k) = (0, 1)$. Similarly to the former case, by Lemma 2.20, we obtain that $\text{sup}(f_{j,0,b}) = \text{sup}(g_{j,b}) = b$ and $\text{sup}(f_{j,1,b-1}) \in \{b-1, b\}$. If $\text{sup}(f_{j,1,b-1}) = b$, then $\text{sig}(X_j) \in E_0$. We show that $\text{sup}(f_{j,1,b-1}) = b-1$ implies $\text{sig}(X_j) = 1$. Assume $\text{sig}(X_j) = (m, n) \in E_\tau$. If $m = n$ or $m > n$, then, by $\text{sup}(f_{j,0,b}) = \text{sup}(f_{j,1,b-1}) - \text{sig}^-(X_j) + \text{sig}^+(X_j)$, we get $\text{sup}(f_{j,0,b}) < b$. This is a contradiction. If $m < n$ then, by $\text{sup}(g_{j,b+1}) = \text{sup}(g_{j,b}) - \text{sig}^-(X_j) + \text{sig}^+(X_j)$, we get the contradiction $\text{sup}(g_{j,b+1}) > b$. Hence, $\text{sig}(X_j) \in \{0, \dots, b\}$ and $(b-1 + |\text{sig}(X_j)|) \equiv b \pmod{b+1}$. This implies that $(b+1)$ divides $(|\text{sig}(X_j)| - 1)$ and thus $|\text{sig}(X_j)| \equiv 1 \pmod{b+1}$. Consequently, we obtain $\text{sig}(X_j) = 1$. This shows that $\text{sig}(k) = (0, 1)$ and $z \in E_0$ implies $\text{sig}(X_j) \in E_0 \cup \{1\}$. \square

Conversely, a one-in-three model of φ implies the τ -ESSP and the τ -SSP for U_τ :

Lemma 4.14. If φ has a one-in-three model, then U_τ has the τ -ESSP and the τ -SSP.

Proof:

Let M be a one-in-three model of φ , and let $I = \{h_{3,0,0}, t_{j,0}, f_{j,0,0}, g_{j,0} \mid 0 \leq j \leq m-1\}$ be the set of the initial states of the gadgets of U_τ .

We start with the solvability of k . The following τ -region $R = (\text{sup}, \text{sig})$ solves $\alpha = (k, h_{3,1,b-1})$ and thus k completely in H_3 : for all $s \in I$, $\text{sup}(s) = 0$; for all $e \in E(U_\tau)$, if $e = k$, then $\text{sig}(e) = (0, 1)$; if $e \in \{z\} \cup (V(\varphi) \setminus M)$, then $\text{sig}(e) = 0$; for all $j \in \{0, \dots, m-1\}$, if $e = v_j$ and $X_j \in M$, then $\text{sig}(e) = 0$; for all $j \in \{0, \dots, m-1\}$, if $e = v_j$ and $X_j \notin M$, then $\text{sig}(e) = 1$; otherwise holds $e \in M \cup \{u\}$, and we define $\text{sig}(e) = 1$.

Notice that this region solves also a lot SSA of U_τ . In particular, if $q_0 \xrightarrow{k} \dots \xrightarrow{k} q_b$, then this region solves (s, s') for all $s \neq s' \in \{1_0, \dots, q_b\}$.

The following region $R = (\text{sup}, \text{sig})$ solves (k, s) for all remaining relevant states of U_τ : for all $s \in I$, $\text{sup}(s) = 0$; for all $e \in E(U_\tau)$, if $e = k$, then $\text{sig}(e) = (0, 1)$; if $e \in \{z\} \cup \{v_0, \dots, v_{m-1}\}$, then $\text{sig}(e) = 1$; otherwise, $\text{sig}(e) = 0$.

We proceed with the solvability of z . Let $i \in \{0, \dots, m-1\}$ be arbitrary but fixed. Let $j, \ell \in \{0, \dots, m-1\} \setminus \{i\}$ such that $j \neq \ell$ and $X_{i_2} \in E(T_j)$ and $X_{i_2} \in E(T_\ell)$. The following region solves (z, s) for all $s \in \{h_{3,0,0}\} \cup S(T_i)$: for all $s \in \{h_{3,0,0}, t_{i,0}, t_{j,0}, t_{\ell,0}\}$, $\text{sup}(s) = b$; for all $s \in \{f_{0,0,0}, g_{j,0} \mid j \in \{0, \dots, m-1\}\}$, $\text{sup}(s) = 1$; for all $s \in \{t_{j,0} \mid j \in \{0, \dots, m-1\}\} \setminus \{t_{i,0}, t_{j,0}, t_{\ell,0}\}$, $\text{sup}(s) = 0$; for all $e \in E(U_\tau)$, if $e = z$, then $\text{sig}(z) = (0, b)$; if $e = X_{i_2}$, then $\text{sig}(e) = 1$; if $n \in \{0, \dots, m-1\}$ and $e = v_n$ and $X_{i_2} \in E(F_n)$, then $\text{sig}(e) = b$; if $e = u$, then $\text{sig}(e) = 1$; otherwise $\text{sig}(e) = 0$. By the arbitrariness of i , this proves also the τ -solvability of (z, s) for all relevant $s \in \bigcup_{j=0}^{m-1} S(T_j)$.

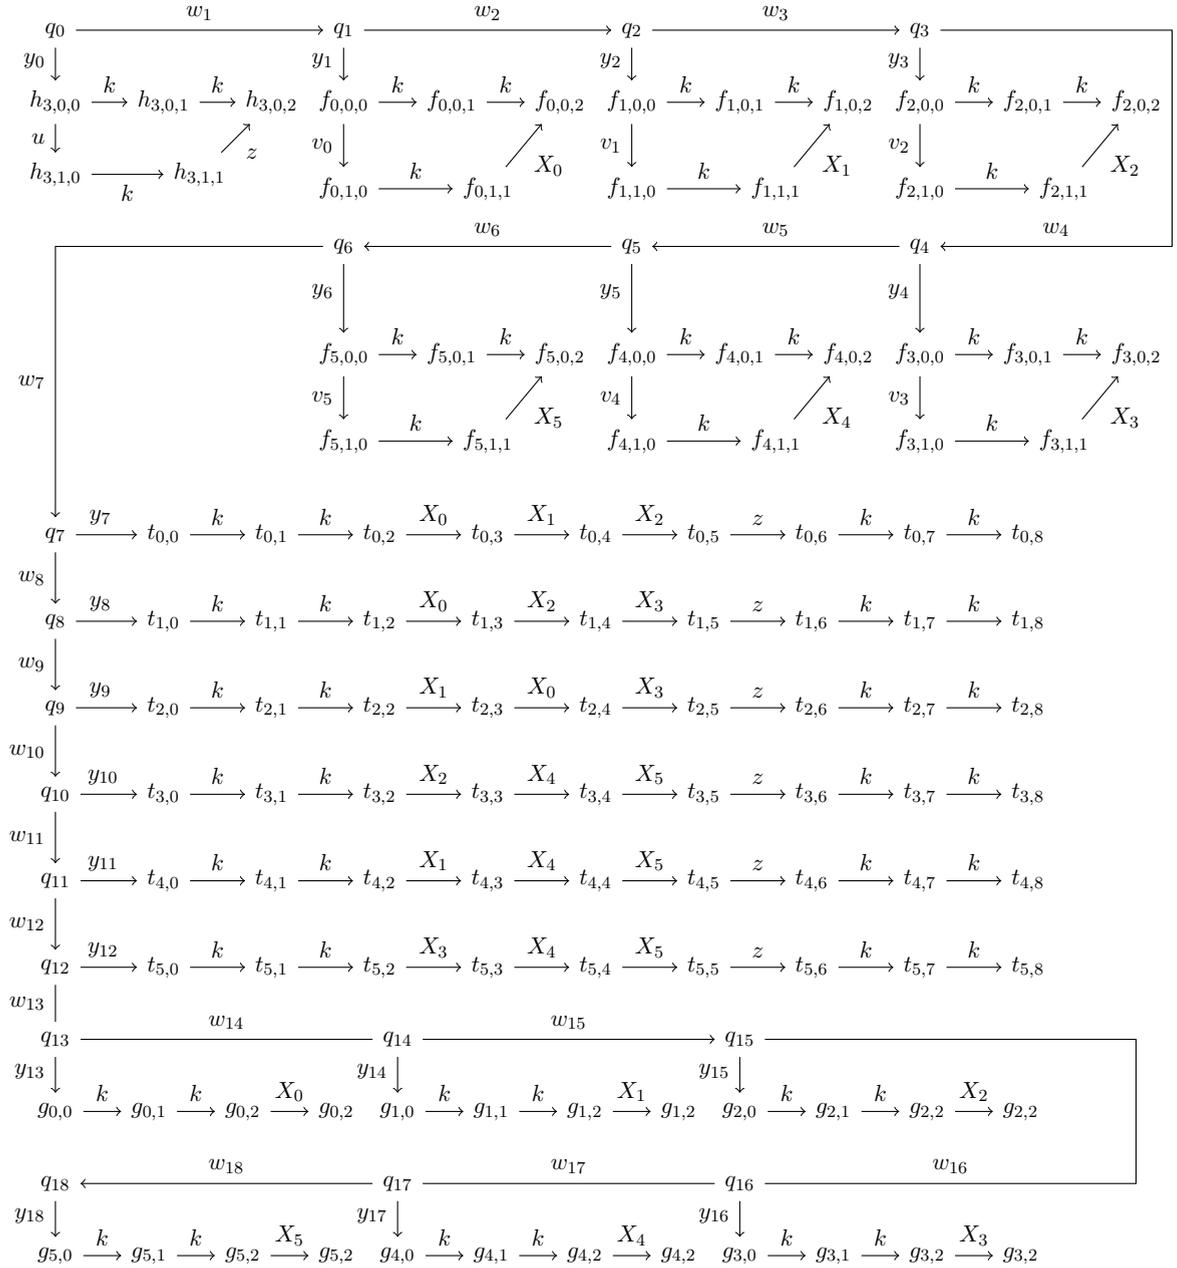


Figure 6: For all $\tau \in \{\tau_{\mathbb{Z}PPT}^2, \tau_{\mathbb{Z}PT}^2\}$, the joining $J(U_\tau)$ where φ corresponds to Example 4.2.

Notice that if $s \in \{h_{3,0,0}, \dots, h_{3,0,b}\}$ and $s' \in \{h_{3,1,0}, \dots, h_{3,1,b-1}\}$ or if $s \in \{f_{i,0,0}, \dots, h_{i,0,b}\}$ and $s' \in \{s_{i,1,0}, \dots, f_{i,1,b-1}\}$, then this region also solves (s, s') . Thus, altogether, we already have proven the solvability of all states of H_3, F_0, \dots, F_{m-1} .

The following region $R = (sup, sig)$ solves (z, s) for all relevant $s \in S(H_3) \setminus \{h_{3,0,0}\}$: for all $s \in I \setminus \{t_{i,0} \mid i \in \{0, \dots, m-1\}\}$, $sup(s) = 0$; for all $s \in \{t_{i,0} \mid i \in \{0, \dots, m-1\}\}$, $sup(s) = 1$; for all $e \in E(U_\tau)$, if $e = z$, then $sig(e) = (0, b)$; if $e = k$, then $sig(k) = 1$; if $e = u$, then $u = 2$; otherwise $sig(e) = 0$.

The following region $R = (sup, sig)$ solves (z, s) for all remaining relevant states: for all $s \in \{h_{3,0,0}\} \cup \{f_{j,0,0}, g_{j,0} \mid j \in \{0, \dots, m-1\}\}$, $sup(s) = b$; for all $s \in \{t_{i,0} \mid i \in \{0, \dots, m-1\}\}$, $sup(s) = 0$; for all $e \in E(U_\tau)$, if $e = z$, then $sig(e) = (0, b)$; if $e = u$, then $sig(e) = 1$; otherwise, $sig(e) = 0$.

We proceed by arguing for the solvability of u . The following region $R = (sup, sig)$ solves (u, s) for all $s \in \{h_{3,0,1}, \dots, h_{3,0,b}\}$: for all $s \in I$, $sup(s) = 0$; for all $e \in E(U_\tau)$, if $e = u$, then $sig(e) = (0, b)$; if $e = z$, then $sig(e) = 2$; if $e = k$, then $sig(e) = 1$; otherwise $sig(e) = 0$.

If $b > 2$, then the following region $R = (sup, sig)$ solves (u, s) for relevant states $s \in S(U_\tau) \setminus \{h_{3,0,1}, \dots, h_{3,0,b}\}$: for all $s \in I$, if $s = h_{3,0,0}$, then $sup(s) = 0$; otherwise, $sup(s) = 1$; for all $e \in E(U_\tau)$, if $e = u$, then $sig(e) = (0, b)$; if $e = z$, then $sig(e) = 1$; otherwise $sig(e) = 0$. If $b = 2$, then we additionally need a slightly modified region that maps $sup(s) = 0$ for all $s \in \{t_{j,0} \mid j \in \{0, \dots, m-1\}\}$. This proves the solvability of u .

We proceed with the solvability of the events v_0, \dots, v_{m-1} . Let $i \in \{0, \dots, m-1\}$ be arbitrary but fixed. The following region $R = (sup, sig)$ solves (u_i, s) for all $s \in \{f_{i,0,1}, \dots, f_{i,0,b}\}$: for all $s \in I$, $sup(s) = 0$; for all $e \in E(U_\tau)$, if $e = u_i$, then $sig(e) = (0, b)$; if $e = X_i$, then $sig(e) = 2$; if $e = k$, then $sig(e) = 1$; otherwise $sig(e) = 0$.

If $b > 2$, then the following region $R = (sup, sig)$ solves (v_i, s) for all remaining relevant states $S(U_\tau) \setminus \{f_{i,0,1}, \dots, f_{i,0,b}\}$: $sup(f_{i,0,0}) = 0$; for all $s \in I \setminus \{f_{i,0,0}\}$, $sup(s) = 1$; for all $e \in E(U_\tau)$, if $e = u_i$, then $sig(e) = (0, b)$; if $e = X_i$, then $sig(e) = 1$; otherwise $sig(e) = 0$. If $b = 2$, then we additionally need a slightly modified region that maps $sup(s) = 0$ for all $s \in \{g_{j,0}, t_{j,0} \mid j \in \{0, \dots, m-1\}\}$. This proves the solvability of v_i . Since i was arbitrary, this proves the solvability of all v_0, \dots, v_{m-1} .

It is easy to see, that the variable events X_0, \dots, X_{m-1} are solvable. Thus, for the sake of simplicity, we refrain from the explicit representation of the corresponding regions. Moreover, one easily verifies that the remaining regions that complete the τ -ESSP of U_τ also solve the remaining SSA of U_τ . Altogether, we have finally proven that if M has a one in three model, then U_τ has the τ -ESSP and the τ -SSP. \square

5. Polynomial time results

The following theorem states the main result of this section:

- Theorem 5.1.**
1. τ_{RZPT}^b -ESSP can be solved in time polynomial in the size of input A .
 2. If $\tau \in \{\tau_{ZPT}^b, \tau_{ZPPT}^b, \tau_{RZPT}^b\}$, then τ -SSP can be solved in time polynomial in the size of input A .

The contribution of Theorem 5.1 is threefold. Firstly, τ -ESSP and τ -SOLVABILITY are NP-complete for all $\tau \in \{\tau_{ZPT}^b, \tau_{ZPPT}^b\}$ by Theorem 4.1. However, Theorem 5.1.2 states that τ -SSP is solvable in

polynomial-time for these types. Hence, to the best of our knowledge, Theorem 5.1 discovers the first Petri net types for which τ -SSP and τ -ESSP as well as τ -SSP and τ -SOLVABILITY provably have a different computational complexity.

Secondly, in [8], Schmitt extended the type τ_{PPT}^1 by the additive group of integers modulo 2, which leads to the tractable (super-) type to τ_{ZPPT}^1 . Moreover, in [9], we argued that Schmitts approach transferred to τ_{PT}^1 yields the tractable type τ_{ZPT}^1 . However, by Theorem 4.1, lifting Schmitts technique to τ_{PPT}^b and τ_{PT}^b does not lead to superclasses with a tractable synthesis problem for all $2 \leq b \in \mathbb{N}$. Hence, Theorem 5.1 proposes the first tractable type of b -bounded Petri nets, where $b \geq 2$, so far. Finally, Theorem 5.1 gives us insight into which of the τ -net properties, where $\tau \in \{\tau_{PPT}^b, \tau_{PT}^b\}$, cause the hardness of τ -SYNTHESIS and the corresponding separation problems. In particular, flow arc relations (events in τ) between places and transitions in a τ -net define conditions when a transition is able to fire. For example, if N is a τ -net with transition t and place p such that $f(p, t) = (1, 0)$ then the firing of t in a marking M requires $M(p) \geq 1$. By Theorem 5.1, the hardness of finding a τ -net N for A originates from the potential possibility of τ -nets to satisfy such conditions by multiple markings $M(p) \in \{1, \dots, b\}$. In fact, the definition of τ_{RZPT}^b implies that $f(p, t) = (m, n)$ requires $M(p) = m$ for the firing of t and prohibits the possibility of multiple choices. By Theorem 5.1, this makes τ_{RZPT}^b -synthesis tractable.

While the question of whether there are superclasses of $\tau_{PT}^b, \tau_{PPT}^b, b \geq 2$, for which synthesis is doable in polynomial time remains unanswered, the following lemma shows that the type τ_{RZPT}^b yields at least a tractable superclasses of Schmitt's type τ_{RZPT}^b [8]. More generally, if $b < b'$ then the class of τ_{RZPT}^b -nets is strictly more comprehensive than the class of $\tau_{RZPT}^{b'}$ -nets:

Lemma 5.2. If $b < b' \in \mathbb{N}^+$ and if \mathcal{T} is the set of τ_{RZPT}^b -solvable TS and \mathcal{T}' the set of $\tau_{RZPT}^{b'}$ -solvable TS then $\mathcal{T} \subset \mathcal{T}'$.

Proof:

We present a TS A that is $\tau_{RZPT}^{b'}$ -solvable but not τ_{RZPT}^b -solvable: Let $A = (\{s_0, \dots, s_{b'}\}, \{a\}, \delta, s_0)$ be the TS with transition function $\delta(s_i, a) = s_{i+1}$ for $i \in \{0, \dots, b' - 1\}$ and $\delta(s_{b'}, a) = s_0$. By other words, A is a directed labeled cycle $s_0 \xrightarrow{a} \dots \xrightarrow{a} s_{b'} \xrightarrow{a} s_0$ where every transition is labeled by a . Notice, that A has no ESSA. Hence, it has the τ -ESSP for every type of nets τ . Consequently, A is τ -solvable if and only if it has the τ -SSP.

Assume, for a contradiction, that A is τ_{RZPT}^b -solvable. By $b < b'$, A provides the SSA (s_0, s_{b+1}) and the τ_{RZPT}^b -solvability of A implies that there is a τ_{RZPT}^b -region (sup, sig) that solves it. If $sig(a) = (m, n)$ then $sup(s_1) = sup(s_0) - m + n \neq sup(s_0)$ and, by definition of τ_{RZPT}^b , $\neg sup(s_1) \xrightarrow{(m, n)}$. This is a contradiction to $s_1 \xrightarrow{a}$. Hence, $sig(a) \in \{1, \dots, b\}$. By induction, $sup(s_{b+1}) = sup(s_0) + (b + 1) \cdot sig(a) = sup(s_0) \pmod{b + 1}$ implying $sup(s_{b+1}) = sup(s_0)$. Thus, (sup, sig) does not solve (s_0, s_{b+1}) , which proves that A is not τ_{RZPT}^b -solvable.

On the contrary, it is easy to see that the $\tau_{RZPT}^{b'}$ -region (sup, sig) , which is defined by $sup(s_0) = 0, sig(a) = 1$ and $sup(s_{i+1}) = sup(s_i) + sig(a)$ for $i \in \{0, \dots, b' - 1\}$, solves every SSA of A . Hence, A is $\tau_{RZPT}^{b'}$ -solvable. \square

5.1. Abstract regions and fundamental cycles

In the remainder of this paper, unless explicitly stated otherwise, we assume that $A = (S, E, \delta, \iota)$ is an arbitrary but fixed (non-trivial) TS with at least two states and event set $E = \{e_1, \dots, e_n\}$. Recall that $\tau \in \{\tau_{\mathbb{Z}PT}^b, \tau_{\mathbb{Z}PPT}^b, \tau_{\mathbb{R}\mathbb{Z}PT}^b\}$ and $b \in \mathbb{N}^+$ are also arbitrary but fixed.

The proof of Theorem 5.1 bases on a generalization of the approach used in [8] that reduces the solvability of ESSA and SSA to the solvability of systems of linear equations modulo $b+1$. It exploits that the solvability of such systems is decidable in polynomial time:

Lemma 5.3. ([24])

Let $M \in \mathbb{Z}_{b+1}^{k \times n}$ and $c \in \mathbb{Z}_{b+1}^k$. There is an algorithm that decides in time $\mathcal{O}(nk \cdot \max\{n, k\})$ whether there is an element $x \in \mathbb{Z}_{b+1}^n$ such that $Mx = c$.

Essentially, our generalization composes for every ESSA and every SSA $\alpha = (x, y)$ of the TS A a system of equations modulo $b+1$ that has a solution if and only if α is τ -solvable. Hence, the TS A has the τ -ESSP, respectively the τ -SSP, if and only if every system, defined by the ESSA of A , respectively by the SSA of A , has a solution.

We proceed by deducing the notion of abstract regions. Our starting point is the goal to obtain τ -regions (sup, sig) of A as solutions of linear equation systems modulo $b+1$. By Definition 2.12 and the definition of τ , (sup, sig) is a τ -region of A if and only if for every transition $s \xrightarrow{e} s'$ it is true that

$$sup(s') = (sup(s) - sig^-(e) + sig^+(e) + |sig(e)|) \bmod (b+1) \quad (14)$$

Hence, installing for every transition $s \xrightarrow{e} s'$ the corresponding Equation 14 yields a linear system of equations whose solutions are regions of A . If (sup, sig) is a solution of this system such that $sig(e) = (m, n) \in E_\tau \setminus \{0, \dots, b\}$ for $e \in E(A)$ then, by definition, for every transition $s \xrightarrow{e} s'$ it has to be true that $m \leq sup(s)$ and $sup(s') - m + n \leq b$. Unfortunately, the conditions $m \leq sup(s)$ and $sup(s') - m + n \leq b$ can not be tested in the group \mathbb{Z}_{b+1} . To cope with this obstacle, we abstract from elements $(m, n) \in E_\tau$ by restricting to regions (solutions) that identify (m, n) with the unique element $x \in \{0, \dots, b\}$ such that $x = (n - m) \bmod (b+1)$. This leads to the notion of *abstract* τ -regions.

Definition 5.4. (Abstract Region)

A τ -region (sup, sig) of $A = (S, E, \delta, \iota)$ is called *abstract* if the codomain of sig is restricted to the elements of \mathbb{Z}_{b+1} , that is, $sig : E \rightarrow \{0, \dots, b\}$. If (sup, sig) is an abstract region, then we call sig an *abstract signature*.

Remark 5.5. (Notation of abstract regions)

For the sake of clarity, we denote abstract signatures by abs instead of sig and abstract regions by (sup, abs) instead of (sup, sig) . For convenience, we also identify $abs = (abs(e_1), \dots, abs(e_n))$.

By definition, two mappings $sup : \{0, \dots, b\} \rightarrow \{0, \dots, b\}$ and $abs : E \rightarrow \{0, \dots, b\}$ define an abstract τ -region if and only if for every transition $s \xrightarrow{e} s'$ of A it is true that

$$sup(s') = (sup(s) + abs(e)) \bmod (b+1) \quad (15)$$

Obviously, for abstract regions, the Equation 14 reduces to Equation 15. Installing for every transition $s \xrightarrow{e} s'$ of A its corresponding Equation 15 yields a system modulo $b + 1$ whose solutions are abstract regions. However, such systems require to deal with sup and abs simultaneously, which is very inconvenient. It is better to first obtain abs independently of sup and then to define sup with the help of abs . The following observations show how to realize this idea.

By induction and Equation 15, one immediately obtains that (sup, abs) is an abstract region if and only if for every directed labeled path $p = \iota \xrightarrow{e'_1} \dots \xrightarrow{e'_m} s_m$ of A from the initial state ι to the state s_m the *path equation* holds:

$$sup(s_m) = (sup(\iota) + abs(e'_1) + \dots + abs(e'_m)) \bmod (b + 1) \quad (16)$$

In order to exploit Equation 16, we first introduce the following notions:

Definition 5.6. (Parikh-vector)

Let $p = z_0 \xrightarrow{a_1} \dots \xrightarrow{a_m} z_m$ be a path of the TS A on pairwise distinct states z_0, \dots, z_m . The *Parikh-vector* of p is the mapping $\psi_p : \{e_1, \dots, e_n\} \rightarrow \{0, \dots, b\}$ such that $\psi_p(e) = |\{i \in \{0, \dots, m-1\} \mid z_i \xrightarrow{e}\}| \bmod (b + 1)$ for every event $e \in \{e_1, \dots, e_n\}$, that is, ψ_p assigns to e the number of its occurrences on p modulo $b + 1$. For convenience, we identify $\psi_p = (\psi_p(e_1), \dots, \psi_p(e_n))$.

Definition 5.7. (Product)

If $x = (x_1, \dots, x_n)$ and $y = (y_1, \dots, y_n)$ are two elements of \mathbb{Z}_{b+1}^n , then we say $x \cdot y = (x_1 \cdot y_1 + \dots + x_n \cdot y_n) \bmod b + 1$ is the *product* of x and y .

Definition 5.6 and Definition 5.7 allow us to reformulate the path equation 16 as follows:

$$sup(s_m) = (sup(\iota) + \psi_p \cdot abs) \bmod (b + 1) \quad (17)$$

Notice that if p, p' are two different paths from ι to s_m , then $\psi_p \cdot abs = \psi_{p'} \cdot abs$. Thus, the support sup is fully determined by $sup(\iota)$ and abs . We obtain sup explicitly by $sup(s) = (sup(\iota) + \psi_p \cdot abs) \bmod (b + 1)$ for all $s \in S$, where p is an arbitrary but fixed path of A that starts at ι and terminates at s . Consequently, every abstract signature abs implies $b + 1$ different abstract τ -regions of A , one for every $sup(\iota) \in \{0, \dots, b\}$. Altogether, we have argued that the challenge of finding abstract regions of A reduces to the task of finding the abstract signatures of A . In the following, we introduce the notion of fundamental cycles, defined by so-called chords of a spanning tree of A , which enables us to find abstract signatures.

Definition 5.8. (Spanning tree, chord)

A *spanning tree* A' of TS A is a sub-transition system $A' = (S, E', \delta_{A'}, \iota)$ of A with the same set of states S , an event set $E' \subseteq E$ and a restricted transition function $\delta_{A'}$ such that first $\delta_{A'}(s, e) = s'$ entails $\delta_A(s, e) = s'$ and, moreover, for every $s \in S$ there is *exactly* one path $p = \iota \xrightarrow{e_1} \dots \xrightarrow{e_m} s$ in A' . Every transition $s \xrightarrow{e} s'$ of A which is not in A' is called a *chord* (of A').

Remark 5.9. (Parikh-vector of a state in the spanning tree)

For every $s \in S$, by ψ_s we denote the Parikh-vector ψ_p of the unique path $p = \iota \xrightarrow{e_1} \dots \xrightarrow{e_m} s$ in A' .

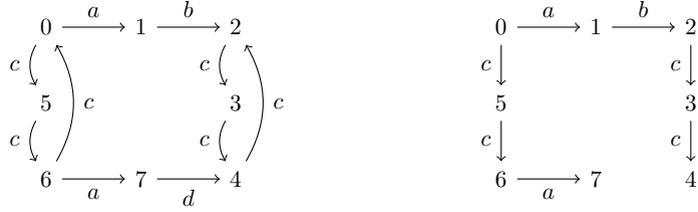


Figure 7: Left: An input TS A . Right: A spanning tree A' of TS A . The unique Parikh vectors ψ_0, \dots, ψ_7 of A' (written as rows) are given by $\psi_0 = (0, 0, 0, 0), \psi_1 = (1, 0, 0, 0), \psi_2 = (1, 1, 0, 0), \psi_3 = (1, 1, 1, 0), \psi_4 = (1, 1, 2, 0), \psi_5 = (0, 0, 1, 0), \psi_6 = (0, 0, 2, 0)$ and $\psi_7 = (1, 0, 2, 0)$. The transitions $\delta_A(7, d) = 4, \delta_A(4, c) = 2$ and $\delta_A(6, c) = 0$ of A define the chords of A' . The corresponding fundamental cycles are given by $\psi_t = \psi_7 + (0, 0, 0, 1) - \psi_4 = (0, 2, 0, 1)$ and $\psi_{t'} = \psi_4 + (0, 0, 1, 0) - \psi_2 = (0, 0, 0, 0)$ and $\psi_{t''} = \psi_6 + (0, 0, 1, 0) - \psi_0 = (0, 0, 0, 0)$. Hence, if $abs = (x_a, x_b, x_c, x_d)$ then $\psi_t \cdot abs = 0 \cdot x_a + 2 \cdot x_b + 0 \cdot x_c + x_d = 2 \cdot x_b + x_d$. By $\psi_{t'} \cdot abs = \psi_{t''} \cdot abs = 0$ for every map abs , only the equation $2 \cdot x_b + x_d = 0$ contributes to the basic part of every upcoming system.

Notice that the underlying undirected graph of A' is a tree in the common graph-theoretical sense. The chords of A' are exactly the edges that induce a cycle in the underlying undirected graph of A' . This gives rise to the following notion of fundamental cycles:

Definition 5.10. (Fundamental cycle)

Let $t = s \xrightarrow{e} s'$ be a chord of A' . The *fundamental cycle* of t is the mapping $\psi_t : \{e_1, \dots, e_n\} \rightarrow \{0, \dots, b\}$ that is defined as follows for all $i \in \{1, \dots, n\}$:

$$\psi_t(e_i) = \begin{cases} \psi_s(e_i) - \psi_{s'}(e_1) \bmod b + 1, & \text{if } e_i \neq e \\ \psi_s(e_i) - \psi_{s'}(e_i) + 1 \bmod b + 1, & \text{else.} \end{cases}$$

For convenience, we identify $\psi_t = (\psi_t(e_1), \dots, \psi_t(e_n))$.

By the following lemma, we can use the fundamental cycles to generate abstract signatures of A :

Lemma 5.11. If A' is a spanning tree of a TS A with chords t_1, \dots, t_k then $abs \in \mathbb{Z}_{b+1}^n$ is an abstract signature of A if and only if $\psi_{t_i} \cdot abs = 0$ for all $i \in \{1, \dots, k\}$. Two different spanning trees A' and A'' provide equivalent systems of equations.

Proof:

We start with proving the first statement. *If:* Let $abs \in \mathbb{Z}_{b+1}^n$ such that $\psi_{t_i} \cdot abs = 0$ for all $i \in \{1, \dots, k\}$ and $sup(\iota) \in \{0, \dots, b\}$. Let $sup(\iota) \in \{0, \dots, b\}$ be arbitrary but fixed and, for all $s \in S$, let $sup(s) = sup(\iota) + \psi_s \cdot abs$. We show that (sup, sig) is an abstract region of A , that is, for all edges $t = s \xrightarrow{a} s'$ of A holds $sup(s') = sup(s) + abs(a) \bmod b + 1$: By definition, we have $sup(s) = sup(\iota) + \psi_s \cdot abs$ and $sup(s') = sup(\iota) + \psi_{s'} \cdot abs$. If t is not a chord, then $\psi_{s'}(a) = \psi_s(a) + 1 \bmod b + 1$ and $\psi_{s'}(e) = \psi_s(e)$ for all $e \in \{e_1, \dots, e_n\} \setminus \{a\}$. This implies $sup(s') = sup(\iota) + \psi_s \cdot abs + abs(a) \bmod b + 1$ and thus $sup(s') = sup(s) + abs(a) \bmod b + 1$.

Otherwise, if t is a chord of A' , then it holds $\psi_t(a) = \psi_s(a) - \psi_{s'}(a) + 1$ and the following implications (considered modulo $b + 1$) are true:

$$\begin{aligned}
0 &= \psi_t \cdot abs && \iff \\
0 &= \sum_{i=1}^n ((\psi_s(e_i) - \psi_{s'}(e_i)) \cdot abs(e_i) + abs(a)) && \iff \\
0 &= \sum_{i=1}^n \psi_s(e_i) \cdot abs(e_i) - \sum_{i=1}^n \psi_{s'}(e_i) \cdot abs(e_i) + abs(a) && \iff \\
\psi_{s'} \cdot abs &= \psi_s \cdot abs + abs(a) && \iff \\
sup(\iota) + \psi_{s'} \cdot abs &= sup(\iota) + \psi_s \cdot abs + abs(a) && \iff \\
sup(s') &= sup(s) + abs(a)
\end{aligned}$$

Hence, abs is an abstract signature of A and the proof shows how to get a corresponding abstract region (sup, abs) of A .

Only-if: If abs is an abstract region of A then we have $sup(s') = sup(s) + abs(e)$ for every transition in A . Hence, if $t = s \xrightarrow{e} s'$ is a chord of a spanning tree A' of A then working backwards the equivalent equalities above proves $\psi_t \cdot abs = 0$.

The second statement is implied by the first: If A', A'' are two spanning trees of A with fundamental cycles $\psi_{t_1}^{A'}, \dots, \psi_{t_k}^{A'}$ and $\psi_{t_1}^{A''}, \dots, \psi_{t_k}^{A''}$, respectively, then we have for $abs \in \mathbb{Z}_{b+1}^n$ that $\psi_{t_i}^{A'} \cdot abs = 0, i \in \{1, \dots, k\}$ if and only if abs is an abstract signature of A if and only if $\psi_{t_i}^{A''} \cdot abs = 0, i \in \{1, \dots, k\}$. \square

In the following, justified by Lemma 5.11, we assume A' to be a fixed spanning tree of A with chords t_1, \dots, t_k . By $M_{A'}$ we denote the system of equations that consists of $\psi_{t_i} \cdot abs = 0$ for all $i \in \{1, \dots, k\}$. A spanning tree of A is computable in polynomial time: As δ_A is a function, A has at most $|E||S|^2$ edges and A' contains $|S| - 1$ edges. Thus, by $2 \leq |S|$, A' has at most $|E||S|^2 - 1$ chords. Consequently, a spanning tree A' of A is computable in time $\mathcal{O}(|E||S|^3)$ [25].

To get polynomial time solvable systems of equations, we have restricted ourselves to equations like Equation 16 or its reformulated version Equation 17. This restriction results in the challenge to compute abstract signatures of A . By Lemma 5.11, abstract signatures of A are solutions of $M_{A'}$. We get an (abstract) τ -region (sup, abs) of A from $sup(\iota)$ and abs by defining $sup(\iota)$ and $sup(s) = sup(\iota) + \psi_s \cdot abs$ for all $s \in S$. However, if (s, s') is an SSA of A then $sup(s) \neq sup(s')$ is not implied. Moreover, by definition, to solve an ESSA (e, s) , we need (concrete) τ -regions (sup, sig) such that $sig : E \rightarrow E_\tau$. The next section shows how to extend $M_{A'}$ to get such solving τ -regions.

5.2. The Proof of Theorem 5.1

This section shows how to extend $M_{A'}$ for a given (E)SSA α to get a system M_α , whose solution yields a region solving α if there is one. But first we need the following lemma that tells us how to obtain abstract regions from (concrete) regions:

Lemma 5.12. If (sup, sig) is a τ -region of a TS $A = (S, E, \delta, \iota)$ then we obtain a corresponding abstract τ -region (sup, abs) by defining abs for $e \in E$ as follows: If $sig(e) = (m, n)$ then $abs(e) = -m + n \bmod (b + 1)$ and, otherwise, if $sig(e) \in \{0, \dots, b\}$ then $abs(e) = sig(e)$.

Proof:

We have to show that $s \xrightarrow{e} s'$ in A entails $sup(s) \xrightarrow{abs(e)} sup(s')$ in τ . If $abs(e) = sig(e) \in \{0, \dots, b\}$ this is true as (sup, sig) is a τ -region.

If $sig(e) = (m, n)$ then, by definition, we have $sup(s') = sup(s) - m + n \bmod (b + 1)$ implying $sup(s') - sup(s) = -m + n \bmod (b + 1)$. By $abs(e) = -m + n \bmod (b + 1)$ and symmetry, we get $-m + n = abs(e) \bmod (b + 1)$ and, by transitivity, we obtain $sup(s') - sup(s) = abs(e) \bmod (b + 1)$ which implies $sup(s') = sup(s) + abs(e) \bmod (b + 1)$. Thus $sup(s) \xrightarrow{abs(e)} sup(s')$. \square

If α is an SSA (s, s') then we only need to assure that the (abstract) region (sup, abs) built on a solution of $M_{A'}$ satisfies $sup(s) \neq sup(s')$. By $sup(s) = sup(\iota) + \psi_s \cdot abs$ and $sup(s') = sup(\iota) + \psi_{s'} \cdot abs$, it is sufficient to extend $M_{A'}$ in a way that ensures $\psi_s \cdot abs \neq \psi_{s'} \cdot abs$. The next lemma proves this claim.

Lemma 5.13. If $\tau \in \{\tau_{\mathbb{Z}PT}^b, \tau_{\mathbb{Z}PPT}^b, \tau_{\mathbb{R}ZPT}^b\}$ then an SSA (s, s') of $A = (S, E, \delta, \iota)$ is τ -solvable if and only if there is an abstract signature abs of A with $\psi_s \cdot abs \neq \psi_{s'} \cdot abs$.

Proof:

If: If abs is an abstract signature with $\psi_s \cdot abs \neq \psi_{s'} \cdot abs$ then the τ -region (sup, abs) with $sup(\iota) = 0$ and $sup(s) = \psi_s \cdot abs$ satisfies $sup(s) \neq sup(s')$. *Only-if:* If (sup, sig) is a τ -region then we obtain a corresponding abstract τ -region (sup, abs) as defined in Lemma 5.12. Clearly, abs is an abstract signature and satisfies the path equations. Consequently, by $sup(s_0) + \psi_s \cdot abs = sup(s) \neq sup(s') = sup(s_0) + \psi_{s'} \cdot abs$, we have that $\psi_s \cdot abs \neq \psi_{s'} \cdot abs$. \square

The next lemma applies Lemma 5.13 to get a polynomial time algorithm which decides the τ -SSP if $\tau \in \{\tau_{\mathbb{Z}PT}^b, \tau_{\mathbb{Z}PPT}^b, \tau_{\mathbb{R}ZPT}^b\}$.

Lemma 5.14. If $\tau \in \{\tau_{\mathbb{Z}PT}^b, \tau_{\mathbb{Z}PPT}^b, \tau_{\mathbb{R}ZPT}^b\}$ then to decide whether a TS $A = (S, E, \delta, \iota)$ has the τ -SSP is doable in time $\mathcal{O}(|E|^3 \cdot |S|^6)$.

Proof:

If $\alpha = (s, s')$ is an SSA of A then the (basic) part $M_{A'}$ of M_α consists of at most $|E| \cdot |S|^2 - 1$ equations for the fundamental cycles. To satisfy $\psi_s \cdot abs \neq \psi_{s'} \cdot abs$, we add the equation $(\psi_s - \psi_{s'}) \cdot abs = q$, where initially $q = 1$, and get (the first possible) M_α . A solution of M_α provides an abstract region satisfying $\psi_s \neq \psi_{s'}$. By Lemma 5.13, this proves the solvability of α . If M_α is not solvable then we modify M_α to M'_α simply by incrementing q and try to solve M'_α . Either we get a solution or we modify M'_α to M''_α by incrementing q again. By Lemma 5.13, if (s, s') is solvable then there is a $q \in \{1, \dots, b\}$ such that the corresponding (modified) system has a solution. Hence, after at most b iterations we can decide whether (s, s') is solvable or not. Consequently, we have to solve at most b linear systems with at most $|E| \cdot |S|^2$ equations for (s, s') . The value b is not part of the input. Thus,

by Lemma 5.3, this is doable in $\mathcal{O}(|E|^3 \cdot |S|^4)$ time. We have at most $|S|^2$ different SSA to solve. Hence, we can decide the τ -SSP in time $\mathcal{O}(|E|^3 \cdot |S|^6)$. \square

As a next step, we let $\tau = \tau_{RZPT}^b$ and prove the polynomial time decidability of τ -ESSP. Let α be an ESSA (e, s) and let s_1, \dots, s_k be the sources of e in A . By definition, a τ -region (sup, sig) solves α if and only if $sig(e) = (m, n)$ and $\neg sup(s) \xrightarrow{sig(e)}$ for a $(m, n) \in E_\tau$. By definition of τ , every element $(m, n) \in E_\tau$ occurs at exactly one state in τ and this state is m . Hence, $sup(s_1) = \dots = sup(s_k) = m$ and $sup(s) \neq m$. We base the following lemma on this simple observation. It provides necessary and sufficient conditions that an *abstract* region must fulfill to imply a *solving* (concrete) region.

Lemma 5.15. Let $\tau = \tau_{RZPT}^b$ and $A = (S, E, \delta, \iota)$ be a TS and let $s_1 \xrightarrow{e} s'_1, \dots, s_k \xrightarrow{e} s'_k$ be the e -labeled transitions in A , that is, if $s' \in S \setminus \{s_1, \dots, s_k\}$ then $\neg s' \xrightarrow{e}$. The atom (e, s) is τ -solvable if and only if there is an event $(m, n) \in E_\tau$ and an abstract region (sup, abs) of A such that the following conditions are satisfied:

1. $abs(e) = -m + n \bmod (b + 1)$,
2. $\psi_{s_1} \cdot abs = m - sup(\iota) \bmod (b + 1)$,
3. $(\psi_{s_1} - \psi_{s_i}) \cdot abs = 0 \bmod (b + 1)$ for $i \in \{2, \dots, k\}$
4. $(\psi_{s_1} - \psi_s) \cdot abs \neq 0 \bmod (b + 1)$.

Proof:

If: Let (sup, abs) be an abstract region that satisfies the conditions 1-4. We obtain a τ -solving region (sup, sig) with (the same support and) the signature sig defined by $sig(e') = abs(e')$ if $e' \neq e$ and $sig(e') = (m, n)$ if $e' = e$. To argue that (sup, sig) is a τ -region we have to argue that $q \xrightarrow{e'} q'$ in A implies $sup(q) \xrightarrow{sig(e')} sup(q')$. As (sup, abs) is an abstract region this is already clear for transitions $q \xrightarrow{e'} q'$ where $e' \neq e$. Moreover, (sup, abs) satisfies $\psi_{s_1} \cdot abs = m - sup(\iota) \bmod (b + 1)$ and the path equation holds, that is, $sup(s_1) = sup(\iota) + \psi_{s_1} \cdot abs \bmod (b + 1)$ which implies $sup(s_1) = m$. Consequently, by definition of τ , we have $sup(s_1) \xrightarrow{(m, n)} n$ in τ . Furthermore, by $abs(e) = -m + n \bmod (b + 1)$ we have $m + abs(e) = n \bmod (b + 1)$. Hence, by $sup(s_1) \xrightarrow{abs(e)} sup(s'_1)$, we conclude $sup(s'_1) = n$ and, thus, $sup(s_1) \xrightarrow{(m, n)} sup(s'_1)$. By $(\psi_{s_1} - \psi_{s_i}) \cdot abs = 0 \bmod (b + 1)$ for $i \in \{2, \dots, k\}$, we obtain that $sup(s_1) = \dots = sup(s_k) = m$. Therefore, similar to the discussion for $s_1 \xrightarrow{e} s'_1$, we obtain by $sup(s_i) \xrightarrow{abs(e)} sup(s'_i)$ that the transitions $sup(s_i) \xrightarrow{(m, n)} sup(s'_i)$ are present in τ for $i \in \{2, \dots, k\}$. Consequently, (sup, sig) is a τ -region.

Finally, by $(\psi_{s_1} - \psi_s) \cdot abs \neq 0 \bmod (b + 1)$, have that $sup(s_1) \neq sup(s)$ and thus $\neg sup(s) \xrightarrow{sig(e)}$. This proves (e, s) to be τ -solvable by (sup, sig) .

Only-if: Let (sup, sig) be a τ -region that solves (e, s) implying, by definition, $\neg sup(s) \xrightarrow{sig(e)}$. We use (sup, sig) to define a corresponding abstract τ -region (sup, abs) in accordance to Lemma 5.12. If $sig(e) \in \{0, \dots, b\}$ then $sup(s) \xrightarrow{sig(e)}$, a contradiction. Hence, it is $sig(e) = (m, n) \in E_\tau$ such that $sup(s_i) \xrightarrow{(m, n)}$ for $i \in \{1, \dots, k\}$ and $\neg sup(s) \xrightarrow{(m, n)}$. This immediately implies $sup(s) \neq sup(s_1)$ and, hence, $(\psi_{s_1} - \psi_s) \cdot abs \neq 0 \pmod{b+1}$. By $sup(s_i) \xrightarrow{(m, n)} sup(s'_i)$ and definition of τ , we have that $sup(s_i) = m$ and $sup(s'_i) = n$ for $i \in \{1, \dots, k\}$ implying $(\psi_{s_1} - \psi_{s_i}) \cdot abs = 0 \pmod{b+1}$ for $i \in \{2, \dots, k\}$. Moreover, by $sup(s_1) \xrightarrow{abs(e)} sup(s'_1)$ we have $abs(e) = sup(s'_1) - sup(s_1) \pmod{b+1}$. Hence, it is $abs(e) = -m + n \pmod{b+1}$. Finally, by the path equation, we have $sup(s_1) = sup(\iota) + \psi_{s_1} \cdot abs \pmod{b+1}$ which with $sup(s_1) = m$ implies $\psi_{s_1} \cdot abs = m - sup(\iota) \pmod{b+1}$. This proves the lemma. \square

The proof of the following lemma exhibits a polynomial time decision algorithm for the τ_{RZPT}^b -ESSP: Given a TS $A = (S, E, \delta, \iota)$ and a corresponding ESSA α , the system $M_{A'}$ is extended to a system M_α . If M_α has a solution abs , then it implies a region (sup, abs) satisfying the conditions of Lemma 5.15 and thus implies the τ -solvability of α . Conversely, if α is solvable, then there is an abstract region (sup, abs) that satisfies the four conditions of by Lemma 5.15. The abstract signature abs is the solution of a corresponding equation system M_α . Hence, we get a solvable M_α if and only if α is solvable. We argue that the number of possible systems is bounded polynomially in the size of A . The solvability of every system is also decidable in polynomial time. Consequently, by the at most $|E| \cdot |S|$ ESSA to solve, this yields the announced decision procedure.

Lemma 5.16. If a TS $A = (S, E, \delta, \iota)$ has the τ_{RZPT}^b -ESSP is decidable in time $\mathcal{O}(|E|^4 \cdot |S|^5)$.

Proof:

To estimate the computational complexity of deciding the τ_{RZPT}^b -ESSP for A observe that A has at most $|S| \cdot |E|$ ESSA to solve. Hence, the maximum costs of deciding the τ_{RZPT}^b -ESSP for A equals $|S| \cdot |E|$ times the maximum effort for a single atom.

In order to decide the τ -solvability of a single ESSA (e, s) , we compose systems in accordance to Lemma 5.15. The maximum costs can be estimated as follows: The (basic) part $M_{A'}$ of M_α has at most $|E| \cdot |S|^2$ equations. Moreover, e occurs at most at $|S| - 1$ states. This makes at most $|S|$ equations to ensure that e 's sources will have the same support, the third condition of Lemma 5.15. According to the first and the second condition, we choose an event $(m, n) \in E_\tau$, a value $sup(\iota) \in \{0, \dots, b\}$, define $abs(e) = -m + n \pmod{b+1}$ and add the corresponding equation $\psi_{s_1} \cdot abs = m - sup(\iota)$. For the fourth condition we choose a fixed value $q \in \{1, \dots, b\}$ and add the equation $(\psi_{s_1} - \psi_s) \cdot abs = q$. Hence, the system has at most $2 \cdot |E| \cdot |S|^2$ equations.

By Lemma 5.3, one checks in time $\mathcal{O}(|E|^3 \cdot |S|^4)$ if such a system has a solution. Notice, we use that $2 \cdot |E| \cdot |S|^2 = \max\{|E|, 2 \cdot |E| \cdot |S|^2\}$. There are at most $(b+1)^2$ possibilities to choose a corresponding $(m, n) \in E_\tau$ and only $b+1$ possible values for x and for q , respectively. Hence, for a fixed atom (e, s) , we have to solve at most $(b+1)^4$ such systems and b is not part of the input. Consequently, we can decide in time $\mathcal{O}(|E|^3 \cdot |S|^4)$ if (e, s) is solvable. A provides at most $|S| \cdot |E|$ ESSA. Hence, the τ_{RZPT}^b -ESSP for A is decidable in time $\mathcal{O}(|E|^4 \cdot |S|^5)$. \square

The following lemma completes the proof of Theorem 5.1 and, moreover, shows that τ_{RZPT}^b -SYNTHESIS is solvable in polynomial time.

Corollary 5.17. There is an algorithm that constructs, for a TS $A = (S, E, \delta, \iota)$, a τ_{RZPT}^b -net N with a state graph A_N isomorphic to A if it exists in time $\mathcal{O}(|E|^3 \cdot |S|^5 \cdot \max\{|E|, |S|\})$.

Proof:

By [6], if \mathcal{R} is a set of regions of A containing for each ESSP and SSA of A a solving region, respectively, then the τ -net $N_A^{\mathcal{R}} = (\mathcal{R}, E(A), f, M_0)$, where $f((sup, sig), e) = sig(e)$ and $M_0((sup, sig)) = sup(\iota)$ for $(sup, sig) \in \mathcal{R}, e \in E(A)$, has a state graph isomorphic to A . Hence, the corollary follows from Lemma 5.14 and Lemma 5.16. \square

Example 5.18. We pick up our running example TS A and its spanning tree of Figure 7. We present two steps of the method given by Lemma 5.16 for the type τ_4^2 and check τ_4^2 -solvability of the ESSA $(c, 1)$.

For a start, we choose $(m, n) = (0, 1)$ and $sup(0) = 0$ and determine $abs(c) = -0 + 1 = 1$ which yields $abs = (x_a, x_b, 1, x_d)$. We have to add $\psi_0 \cdot abs = m - sup(0) = 0$ which, by $\psi_0 = (0, 0, 0, 0)$, is always true and do not contribute to the system. Moreover, for $i \in \{0, 2, 3, 4, 5, 6\}$, we add the equation $(\psi_0 - \psi_i) \cdot abs = 0$. We have $\psi_0 - \psi_6 = (0, 0, -2, 0)$ and $(0, 0, -2, 0) \cdot abs = 0 \cdot x_a - 0 \cdot x_b - 2 - 0 \cdot x_d = 0$ yields a contradiction. Hence, $(c, 1)$ is not solvable by a region (sup, sig) where $sup(0) = 0$ and $sig(c) = (0, 1)$. Similarly, we obtain that the system corresponding to $sup(0) \in \{1, 2\}$ and $sig(c) = (0, 1)$ is also not solvable.

For another try, we choose $(m, n) = (2, 2)$ and $sup(0) = 2$. In accordance to the first and the second condition of Lemma 5.15 this determines $abs = (x_a, x_b, 0, x_d)$ and yields the equation $\psi_0 \cdot abs = m - sup(0) = 2 - 2 = 0$ which is always true. For the fourth condition, we pick $q = 2$ and add the equation $(\psi_0 - \psi_1) \cdot abs = 2 \cdot x_a = 2$. Finally, for the third condition, we add for $i \in \{0, 2, 3, 4, 5, 6\}$ the equation $(\psi_0 - \psi_i) \cdot abs = 0$ and obtain the following system of equations modulo $(b + 1)$:

$$\begin{array}{rclcl}
\psi_t \cdot abs & = & 2 \cdot x_b & + x_d & = 0 \\
(\psi_0 - \psi_1) \cdot abs & = & 2 \cdot x_a & & = 2 \\
(\psi_0 - \psi_2) \cdot abs & = & 2 \cdot x_a & + 2 \cdot x_b & = 0 \\
(\psi_0 - \psi_3) \cdot abs & = & 2 \cdot x_a & + 2 \cdot x_b & + 2 \cdot 0 & = 0 \\
(\psi_0 - \psi_4) \cdot abs & = & 2 \cdot x_a & + 2 \cdot x_b & + 1 \cdot 0 & = 0 \\
(\psi_0 - \psi_5) \cdot abs & = & & & 2 \cdot 0 & = 0 \\
(\psi_0 - \psi_6) \cdot abs & = & & & 1 \cdot 0 & = 0
\end{array}$$

This system is solvable by $abs = (1, 2, 0, 2)$. We construct a region in accordance to the proof of Lemma 5.15: By $sup(0) = 2$ we obtain $sup(1) = 2 + \psi_1 \cdot abs = 2 + (1, 0, 0, 0) \cdot (1, 2, 0, 2) = 0$. Similarly, by $sup(i) = 2 + \psi_i \cdot abs$ for $i \in \{2, \dots, 7\}$ we obtain $sup(2) = sup(3) = sup(4) = sup(5) = sup(6) = 2$ and $sup(7) = 0$. Hence, by defining $sig(c) = (2, 2)$, $sig(a) = 1$, $sig(b) = 2$ and $sig(d) = 2$ we obtain a fitting τ_{RZPT}^b -region (sup, sig) that solves $(c, 1)$.

6. Conclusion

In this paper, for all $b \in \mathbb{N}$, we completely characterize the computational complexity of τ -SSP and τ -ESSP and τ -SOLVABILITY for the types of pure b -bounded P/T-nets, b -bounded P/T-nets and their corresponding \mathbb{Z}_{b+1} -extensions. This answers an open problem posed by Schlachter et al. in [17].

Some open problems in the field of Petri net synthesis concern the computational complexity of τ -synthesis up to language equivalence (τ -LANGUAGE SYNTHESIS) and τ -synthesis from modal TS (τ -MODAL SYNTHESIS):

τ -LANGUAGE SYNTHESIS is the task to find for a given TS $A = (S, E, \delta, \iota)$ a τ -net N whose state graph A_N has the same language as A , that is, $L(A_N) = L(A)$. If there is a sought τ -net N for A , then A is called τ -solvable up to language equivalence. To attack this problem, in [6, p. 164], the language $L(A)$ of A is viewed as the TS $L_A = (L(A), E, \delta_L, \varepsilon)$ where $\delta_L(w, e) = we$ if and only if $we \in L(A)$. By the result of [6, p. 164], there is a τ -net N that solves A up to language equivalence if and only if the TS L_A has the τ -ESSP. Since there might be exponentially (or even infinite) many paths in A , computing L_A and then checking the ESSP yields an algorithm that, in general, is at least exponential in the size of A . Anyway, the exact computational complexity of τ -language synthesis has not yet been proven, and, so far, there has been also no lower bound. For $\tau \in \{\tau_{PT}^b, \tau_{PPT}^b\}$, our results imply a lower bound, to be seen as follows: If $A = s_0 \xrightarrow{e_1} s_1 \xrightarrow{e_2} \dots \xrightarrow{e_n} s_n$ is a linear TS, then $L_A = \varepsilon \xrightarrow{e_1} e_1 \xrightarrow{e_2} \dots \xrightarrow{e_n} e_1 \dots e_n$ (the states of L_A are e_1 and $e_1 e_2$ and \dots and $e_1 \dots e_n$). In particular, it is easy to see that A and L_A are isomorphic. Consequently, by [6, p. 164], a linear TS A is τ -solvable up to language equivalence if and only if it has the τ -ESSP. Thus, by Theorem 4.1, τ -language synthesis is NP-hard, since there is a trivial reduction from τ -ESSP to τ -LANGUAGE SYNTHESIS.

τ -modal synthesis [17] is the task to find for a given modal TS M a τ -net N such that the state graph A_N implements A : A modal TS $M = (S, E, \delta_{must}, \delta_{may}, s_0)$ has a set of states S , events E , an initial state s_0 , a (partial) function $\delta_{must} : S \times E \rightarrow S$ that defines the *must*-edges and a (partial) function $\delta_{may} : S \times E \rightarrow S$ that defines the *may* edges of A ; moreover, δ_{must} and δ_{may} satisfy that if $\delta_{must}(s, e) = s'$, then $\delta_{may}(s, e) = s'$, that is, every *must*-arc is a *may*-arc, but not every *may*-arc is necessarily a *must*-arc. A TS A that has the same event set as M implements M if a relation $R \subseteq M(S) \times A(S)$ exists such that $(s_{0,M}, \iota) \in R$ and for all $(s, q) \in R$ and $e \in E(M) = E(A)$ the following holds:

1. If $\delta_{must}(s, e) = s'$, then there is a $q' \in S(A)$ such that $\delta_A(q, e) = q'$ and $(s', q') \in R$.
2. If $\delta_A(q, e) = q'$, then there is a $s' \in S(M)$ such that $\delta_{may}(s, e) = s'$ and $(s', q') \in R$.

If there is a searched net N for M , then M is called τ -implementable. The computational complexity of τ -modal synthesis has been stated as an open problem in [17]. While at least an (exponential) upper bound is given in [17], a lower bound has not yet been stated. Our results imply a lower bound for $\tau \in \{\tau_{PT}^b, \tau_{PPT}^b\}$. This can be seen as follows: Every TS A can be interpreted as a modal TS where the *must*-edges and the *may*-edges coincide. For such a TS, the just introduced implementation relation then reduces to the well known relation of bisimulation [26, p. 22]. Moreover, it is also known

that deterministic TS A_0 and A_1 are bisimilar if and only if they are language equivalent (also-called *trace equivalent*) [26, p. 26].

Altogether, we have justified that a linear TS A has the τ -ESSP if and only if it is τ -solvable up to language equivalence if and only if, interpreted as modal TS, it is implementable by the state graph A_N of a τ -net N . Thus, for $\tau \in \{\tau_{PT}^b, \tau_{PPT}^b\}$, the following theorem is a corollary of Theorem 4.1 and, at least, gives lower bounds for the computational complexity of both τ -language synthesis and τ -modal synthesis:

Theorem 6.1. Let $\tau \in \{\tau_{PT}^b, \tau_{PPT}^b\}$. Deciding for a TS A if it is τ -solvable up to language equivalence or deciding for a modal TS M if it is τ -implementable is NP-hard.

It remains for future work to settle the exact complexity of τ -language synthesis and τ -modal synthesis. Moreover, one might investigate if τ -SOLVABILITY and τ -ESSP remain NP-complete for 1-grade TS if $\tau \in \{\tau_{\mathbb{Z}PT}^b, \tau_{\mathbb{Z}PPT}^b\}$.

References

- [1] Badouel E, Caillaud B, Darondeau P. Distributing Finite Automata Through Petri Net Synthesis. *Formal Asp. Comput.*, 2002. **13**(6):447–470. doi:10.1007/s001650200022.
- [2] van der Aalst WMP. *Process Mining - Discovery, Conformance and Enhancement of Business Processes*. Springer, 2011. doi:10.1007/978-3-642-19345-3.
- [3] Holloway LE, Krogh BH, Giua A. A Survey of Petri Net Methods for Controlled Discrete Event Systems. *Discrete Event Dynamic Systems*, 1997. **7**(2):151–190. doi:10.1023/A:1008271916548.
- [4] Cortadella J, Kishinevsky M, Kondratyev A, Lavagno L, Yakovlev A. A region-based theory for state assignment in speed-independent circuits. *IEEE Trans. on CAD of Integrated Circuits and Systems*, 1997. **16**(8):793–812. doi:10.1109/43.644602.
- [5] Badouel E, Bernardinello L, Darondeau P. Polynomial Algorithms for the Synthesis of Bounded Nets. In: TAPSOFT, volume 915 of *Lecture Notes in Computer Science*. Springer, 1995 pp. 364–378. doi:10.1007/3-540-59293-8_207.
- [6] Badouel E, Bernardinello L, Darondeau P. *Petri Net Synthesis*. Texts in Theoretical Computer Science. An EATCS Series. Springer, 2015. doi:10.1007/978-3-662-47967-4.
- [7] Badouel E, Bernardinello L, Darondeau P. The Synthesis Problem for Elementary Net Systems is NP-Complete. *Theor. Comput. Sci.*, 1997. **186**(1-2):107–134. doi:10.1016/S0304-3975(96)00219-8.
- [8] Schmitt V. Flip-Flop Nets. In: STACS, volume 1046 of *Lecture Notes in Computer Science*. Springer, 1996 pp. 517–528. doi:10.1007/3-540-60922-9_42.
- [9] Tredup R, Rosenke C. The Complexity of Synthesis for 43 Boolean Petri Net Types. In: Gopal TV, Watada J (eds.), *Theory and Applications of Models of Computation - 15th Annual Conference, TAMC 2019, Kitakyushu, Japan, April 13-16, 2019, Proceedings*, volume 11436 of *Lecture Notes in Computer Science*. Springer, 2019 pp. 615–634. doi:10.1007/978-3-030-14812-6_38.
- [10] Tredup R, Rosenke C. On the Hardness of Synthesizing Boolean Nets. In: ATAED@Petri Nets/ACSD, volume 2371 of *CEUR Workshop Proceedings*. CEUR-WS.org, 2019 pp. 71–86.

- [11] Tredup R, Rosenke C, Wolf K. Elementary Net Synthesis Remains NP-Complete Even for Extremely Simple Inputs. In: Petri Nets, volume 10877 of *Lecture Notes in Computer Science*. Springer, 2018 pp. 40–59. doi:10.1007/978-3-319-91268-4_3.
- [12] Tredup R. The Complexity of Synthesizing nop-Equipped Boolean Nets from g -Bounded Inputs (Technical Report), accepted for ToPNoC 2020. *CoRR*, 2019. **abs/1911.05834**. 1911.05834, URL <http://arxiv.org/abs/1911.05834>.
- [13] Pietkiewicz-Koutny M. Transition Systems of Elementary Net Systems with Inhibitor Arcs. In: Azéma P, Balbo G (eds.), Application and Theory of Petri Nets 1997, 18th International Conference, ICATPN '97, Toulouse, France, June 23-27, 1997, Proceedings, volume 1248 of *Lecture Notes in Computer Science*. Springer, 1997 pp. 310–327. doi:10.1007/3-540-63139-9_43.
- [14] Montanari U, Rossi F. Contextual Nets. *Acta Inf.*, 1995. **32**(6):545–596. doi:10.1007/BF01178907.
- [15] Badouel E, Darondeau P. Trace Nets and Process Automata. *Acta Inf.*, 1995. **32**(7):647–679. doi:10.1007/BF01186645. URL <https://doi.org/10.1007/BF01186645>.
- [16] Kleijn J, Koutny M, Pietkiewicz-Koutny M, Rozenberg G. Step semantics of boolean nets. *Acta Inf.*, 2013. **50**(1):15–39. doi:10.1007/s00236-012-0170-2.
- [17] Schlachter U, Wimmel H. k -Bounded Petri Net Synthesis from Modal Transition Systems. In: CONCUR, volume 85 of *LIPICs*. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2017 pp. 6:1–6:15. doi:10.4230/LIPICs.CONCUR.2017.6.
- [18] Hiraishi K. Some Complexity Results on Transition Systems and Elementary Net Systems. *Theor. Comput. Sci.*, 1994. **135**(2):361–376. doi:10.1016/0304-3975(94)90112-0.
- [19] Tredup R, Rosenke C. Narrowing down the Hardness Barrier of Synthesizing Elementary Net Systems. In: CONCUR, volume 118 of *LIPICs*. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2018 pp. 16:1–16:15. doi:10.4230/LIPICs.CONCUR.2018.16.
- [20] Moore C, Robson JM. Hard Tiling Problems with Simple Tiles. *Discrete & Computational Geometry*, 2001. **26**(4):573–590. doi:10.1007/s00454-001-0047-6.
- [21] Tredup R. Hardness Results for the Synthesis of b -bounded Petri Nets. In: Donatelli S, Haar S (eds.), Application and Theory of Petri Nets and Concurrency - 40th International Conference, PETRI NETS 2019, Aachen, Germany, June 23-28, 2019, Proceedings, volume 11522 of *Lecture Notes in Computer Science*. Springer, 2019 pp. 127–147. doi:10.1007/978-3-030-21571-2_9.
- [22] Tredup R. Fixed Parameter Tractability and Polynomial Time Results for the Synthesis of b -bounded Petri Nets. In: Donatelli S, Haar S (eds.), Application and Theory of Petri Nets and Concurrency - 40th International Conference, PETRI NETS 2019, Aachen, Germany, June 23-28, 2019, Proceedings, volume 11522 of *Lecture Notes in Computer Science*. Springer, 2019 pp. 148–168. doi:10.1007/978-3-030-21571-2_10.
- [23] Garey MR, Johnson DS. Computers and Intractability: A Guide to the Theory of NP-Completeness. W. H. Freeman, 1979. ISBN 0-7167-1044-7.
- [24] Goldmann M, Russell A. The Complexity of Solving Equations over Finite Groups. *Inf. Comput.*, 2002. **178**(1):253–262. doi:10.1006/inco.2002.3173.
- [25] Tarjan RE. Finding optimum branchings. *Networks*, 1977. **7**(1):25–35. doi:10.1002/net.3230070103.
- [26] Gorrieri R. Process Algebras for Petri Nets - The Alphabetization of Distributed Systems. Monographs in Theoretical Computer Science. An EATCS Series. Springer, 2017. ISBN 978-3-319-55558-4. doi:10.1007/978-3-319-55559-1.