arXiv:2112.03605v4 [cs.FL] 11 Oct 2022

Some Basic Techniques Allowing Petri Net Synthesis: Complexity and Algorithmic Issues

Raymond Devillers

Département d'Informatique - Université Libre de Bruxelles

Boulevard du Triomphe, B1050 Brussels, Belgium

rdevil@ulb.ac.be

Ronny Tredup[*]

Institut Für Informatik - Universität Rostock

Albert-Einstein-Straße 22, D18059 Rostock, Germany

ronny.tredup@uni-rostock.de

Abstract. In Petri net synthesis we ask whether a given transition system A can be implemented by a Petri net N. Depending on the level of accuracy, there are three ways how N can implement A: an *embedding*, the least accurate implementation, preserves only the diversity of states of A; a *language simulation* already preserves exactly the language of A; a *realization*, the most accurate implementation, realizes the behavior of A exactly. However, whatever the sought implementation, a corresponding net does not always exist. In this case, it was suggested to modify the input behavior – of course as little as possible. Since transition systems consist of states, events and edges, these components appear as a natural choice for modifications. In this paper we show that the task of converting an unimplementable transition system into an implementable one by removing as few states or events or edges as possible is NP-complete –regardless of what type of implementation we are aiming for; we also show that the corresponding parameterized problems are $W[2]$-hard, where the number of removed components is considered as the parameter; finally, we show there is no c-approximation algorithm (with a polynomial running time) for neither of these problems, for every constant $c \geq 1$.

---

[*]Address of correspondence: Institut Für Informatik - Universität Rostock, Albert-Einstein-Straße 22, D18059 Rostock, Germany

1.   Introduction

Petri nets are a widely accepted framework for modeling and validating concurrent and distributed systems. In general, there are two ways to deal with the behavior of Petri nets: *Analysis* starts from a given Petri net and investigates if its behavior satisfies some properties such as, for example, *liveness*, *reachability of deadlocks*, *reachability of stable markings* or the *fireability of a transition* [1]. In *synthesis*, we deal with the opposite direction: starting from a regular behavior, given as a *labeled transition system* (LTS, for short), we try to find a Petri net that implements this behavior.

Synthesis of Petri nets has practical applications in numerous areas such as, for example, data and process mining [2, 3], digital hardware design [4, 5] and discovering of concurrency and distributability [6, 7]. On the other hand, Petri net synthesis has also been the subject of theoretical studies that, for example, aim at characterizing the complexity of synthesis [8] or look for structural properties that classify an LTS as implementable by subclasses of Petri nets such as, for example, *marked graphs*, and thus allow improved synthesis procedures [9]. It also led to various synthesis tools [10, 11, 12, 13, 14].

LTSs have *states*, *events* and labeled *edges*, i.e., "source-event-target" triplets: the occurrence of the event at the source triggers a change of state to the target. They have an *initial state*, from which, triggered by an event-sequence, any other state is reachable. Petri nets have *places* containing *tokens*, and an overall token distribution is considered as a *marking* (i.e., a global state) of the net; nets have *transitions*, connected with places, which possibly can *fire*: the token distribution of their connected places may allow the firing. A firing of a transition (locally) changes the token distribution (of its connected places) and thus (globally) the marking of the net. They have an *initial marking*, from which, triggered by the firing of a sequence of transitions, any other reachable marking is obtained. The global behavior of a net is captured by its *reachability graph*, which is a transition system, where reachable markings become states, transitions become events and edges correspond to "marking-transition-marking"-triplets. A Petri net N implements an LTS A, if the events of A and the transitions of N coincide and, moreover, if (the states of) A and (the states of) the reachability graph of N can be related by a mapping, which satisfies certain requirements.

According to the properties of the mappings, implementations with various degrees of "accuracy" are possible: Such a mapping is first required to be a *simulation*, which means that every allowed sequence of events (starting at the initial state) can be simulated by a (fireable) sequence of transitions (starting at the initial marking). However, finding a net that allows a simulation is not a challenge: If N is the net without places that has a transition e for every event e of A, then N simulates A, since it can simply fire every sequence of events of A. Moreover, this net can obviously simulate *every* LTS that has the same events as A. From this point of view, N simulates A with the greatest inaccuracy and every information about the (forbidden) original behavior is lost. On the other end of the spectrum, N simulates A most accurately if the simulation is an *isomorphism*: then N is an (exact) *realization* of the behavior defined by A. Unfortunately, not every LTS can be realized by a Petri net. However, this is actually not always necessary, depending on the application. Therefore, *embedding* and *language simulation* have been discussed as other possible implementations in the literature, which –in a certain sense– are less accurate, but still acceptable: an embedding preserves at least the diversity of states, that is, the simulation map is injective; a language simulation preserves exactly the allowed event sequences of A, that is, A and N are language-equivalent. Unfortunately,

although these implementations are less restrictive, they also do not ensure the existence of a corresponding net. In order to achieve implementability, various techniques have been proposed in the literature that modify the components of the input behavior, i.e., its states, events and edges [15, 16]. One of the most discussed approaches among them is what is known as *label-splitting*: events are split into several (new) events and edges are relabeled so that edges that are initially labeled with the same event then are labeled with different events, which originate from the same event by splitting. This method is relatively well-understood from the practical point of view –in the sense of available algorithms [17, 4, 18]– and from the theoretical point of view as well. In particular, it was recently shown that achieving implementability by splitting as few events as possible is a problem hard to solve, namely NP-complete, regardless of the implementation kind [19, 20].

In a natural way, the question arises whether there are other (small) suitable modifications to obtain implementable behaviors. The answer is given by the nature of implementations itself: States of the input are related with reachable markings of the implementing net, and if the behavior is not implementable, then for some of its states no reachable marking may exist. Hence, the *state removal* of such states (as well as the part of the LTS which becomes unreachable) may lead to an implementable LTS. Occurrences of events at (source-) states of the input correspond to the firing of transitions in markings that are associated with the sources. If the behavior is not implementable, then the firing of the transitions in the corresponding markings may not be possible. Hence, removing such occurrences (i.e, the corresponding edges) may then yield an implementable behavior. If the latter is an option, then there are two distinct ways to put the focus on the removal: On the one hand, the modeler may allow the *edge removal* of several edges that affect several events and, simultaneously, demand that some occurrences of every event remain –if this is possible. On the other hand, the modeler may come to the conclusion that some events are less interesting than others and thus prefer the complete *event removal* of the former (i.e., all of their occurrences) and the complete preservation of the latter. Just like label-splitting, the removal of states, events and edges is a powerful transformation, since each of it is able to produce an implementable behavior: For instance, when A is degenerated to a single state, the resulting behavior is implementable. Surely, however, this solution is not desirable. Instead, we are interested in the corresponding optimization problems, that is, given an LTS A, we are looking for a modification A' of A, such that the number κ of removed edges, events or states is as small as possible –depending on the technique applied. If we turn the number κ into a part of the input, then we obtain the corresponding decision version of the optimization problem. Obviously, if we can solve the optimization problem, then we can solve the decision version as well (with only polynomial overhead). Hence, the characterization of the computational complexity of the latter problem provides a lower bound of the complexity of the former.

In this paper, which is an extended version of [21], we completely characterize the computational complexity of *state removal*, *event removal* and *edge removal* for all thinkable implementations they can aim for, i.e., embedding, language-simulation and realization. In particular, we show that all of these decision problems are NP-complete and thus their optimization variant is also hard to solve.

Although the problem of finding an optimal edge, event, or state distance is difficult to solve, we still want to find such implementable modifications of LTS. There are at least two common ways to deal with such difficult problems: First, one can investigate whether the decision variants can be solved exactly by an algorithm in which the exponential explosion of its running time depends functionally

on only a (reasonable) parameter of the input, while the input length contributes only by a polynomial factor, i.e., one can ask whether the problem is *fixed-parameter-tractable* (fpt for short). In this paper, we show that, unfortunately, for all problems and all implementations this question is to be answered negatively if we consider the number κ of removed components as the parameter. In particular, we show that all our problems parameterized by this natural parameter are at least $W[2]$-hard.

On the other hand, instead of an exact solution, one could also look for an approximate solution. In this case we wish that such an approximate solution is provably of at least some constant quality. That is, we look for a so-called c-approximation algorithm, where $c \geq 1$ is a constant, which for any given LTS A outputs in polynomial time an implementable LTS B which emerges from A by deleting at most c times as many components as it would be optimally necessary. In this paper, for all such problems, and all implementations, we show that such an algorithm does not exist, unless P=NP.

The paper is structured as follows: Section 2 provides the basic notions and supports them with some examples. After that, Section 3, Section 4, and Section 5 provide the NP-completeness of *edge removal*, *event removal*, and *state removal*, respectively. After that, in Section 6, we argue that these problems are $W[2]$-hard when parameterized by the number of removed components. In Section 7, we argue that none of the optimization versions of these problems allows for a c-approximation algorithm. Finally, Section 8 briefly concludes and suggests subsequent possible developments.

2. Preliminaries

This section provides the basic notions that we use throughout the paper and supports them with examples.

2.1. Transition systems, and their modifications

The overall starting point for the synthesis of Petri nets is a behavior that is given by a transition system:

Definition 2.1. (Transition System)
A (deterministic, labeled) *transition system* (LTS, for short) $A = (S, E, \delta, \iota)$ consists of two disjoint sets of *states* S and *events* E, a partial *transition function* $\delta : S \times E \longrightarrow S$ and an *initial state* $\iota \in S$. An event e *occurs* at state s, denoted by $s \xrightarrow{e}$, if $\delta(s, e)$ is defined. By $s \xrightarrow{\neg e}$ we denote that $\delta(s, e)$ is not defined. We abridge $\delta(s, e) = s'$ by $s \xrightarrow{e} s'$ and call the latter an *edge* with *source* s and *target* s'. By $s \xrightarrow{e} s' \in A$, we denote that the edge $s \xrightarrow{e} s'$ is present in A. A sequence $s_0 \xrightarrow{e_1} s_1, s_1 \xrightarrow{e_2} s_2, \ldots, s_{n-1} \xrightarrow{e_n} s_n$ of edges is called a (directed labeled) *path* (from s_0 to s_n in A). A is called *reachable*, if there is a path from ι to s for every state $s \in S$. The *language* of A is the set of words $L(A) = \{e_1 \ldots e_n \in E^* \mid \exists s \in S : \iota \xrightarrow{e_1} \ldots \xrightarrow{e_n} s\} \cup \{\varepsilon\}$, where ε denotes the empty word.

In the remainder of this paper, we always assume that LTSs are *reachable*. In this paper, we relate LTSs with the same set of events by so-called *simulations*:

Definition 2.2. A *simulation* between an LTS $A = (S, E, \delta, \iota)$ and an LTS $B = (S', E, \delta', \iota')$ is a mapping $\varphi : S \to S'$ such that $\varphi(\iota) = \iota'$ and if $s \overset{e}{\longrightarrow} s' \in A$, then $\varphi(s) \overset{e}{\longrightarrow} \varphi(s') \in B$; φ is called an *embedding*, denoted by $A \hookrightarrow B$, if it is injective, that is, if $s \neq s'$, then $\varphi(s) \neq \varphi(s')$; φ is a *language-simulation*, denoted by $A \triangleright B$, if $s \overset{\neg e}{\longrightarrow}$ implies $\varphi(s) \overset{\neg e}{\longrightarrow}$; φ is an *isomorphism*, denoted by $A \cong B$, if it is bijective and $s \overset{e}{\longrightarrow} s' \in A$ if and only if $\varphi(s) \overset{e}{\longrightarrow} \varphi(s') \in B$.

It is known from the literature that if $A \triangleright B$, then $L(A) = L(B)$ [15]; if $A \cong B$, then A and B are basically the same –but possibly for the names of their states. An LTS describes a behavior that is implementable or not. In the latter case, we may apply the following *modifications* in order to obtain an implementable LTS:

Definition 2.3. (Edge Removal)
Let $A = (S, E, \delta, \iota)$ be an LTS. An LTS $B = (S', E', \delta', \iota)$ with state set $S' \subseteq S$ and event set $E' \subseteq E$ is an *edge removal* of A if, for all $e \in E'$ and all $s, s' \in S'$, holds: if $s \overset{e}{\longrightarrow} s' \in B$, then $s \overset{e}{\longrightarrow} s' \in A$. By $\mathfrak{K} = \{s \overset{e}{\longrightarrow} s' \in A \mid s \overset{e}{\longrightarrow} s' \notin B\}$ we refer to the (set of) removed edges.

Definition 2.4. (Event Removal)
Let $A = (S, E, \delta, \iota)$ be an LTS. An LTS $B = (S', E', \delta', \iota)$ with state set $S' \subseteq S$ and event set $E' \subseteq E$ is an *event removal* of A if for all $e \in E'$ the following is true: $s \overset{e}{\longrightarrow} s' \in B$ if and only if $s \overset{e}{\longrightarrow} s' \in A$ for all $s, s' \in S$. By $\mathfrak{E} = E \setminus E'$ we refer to the (set of) removed events.

Definition 2.5. (State Removal)
Let $A = (S, E, \delta, \iota)$ be an LTS. An LTS $B = (S', E', \delta', \iota)$ with states $S' \subseteq S$ and events $E' \subseteq E$ is a *state removal* of A if the following two conditions are satisfied: (1) $s \overset{e}{\longrightarrow} s' \in B$ if and only if $s \overset{e}{\longrightarrow} s' \in A$ for all $e \in E'$ and all $s, s' \in S'$; (2) if $s \overset{e}{\longrightarrow} s' \in A$ and $s \overset{e}{\longrightarrow} s' \notin B$, then $s \notin S'$ or $s' \notin S'$. By $\mathfrak{S} = S \setminus S'$ we refer to the (set of) removed states.

Notice that neither of these modifications is "functional", since, generally, there are several LTS that can be considered as a suitable modification of A. Moreover, edge removal is the most general modification introduced, since every event- or state removal is also an edge removal. However, not every edge removal is an event removal or a state removal, not every event removal is a state removal, and not every state removal is an event removal. In particular, there are substantial differences between these modifications that focus on different aspects of the LTS: If B is an edge removal, then there could possibly be an event $e \in E'$ for which there is an edge $s \overset{e}{\longrightarrow} s'$ in A that is not in B. In contrast, if B is an event removal and $e \in E'$, then every e-labeled edge of A has to be present in B. Furthermore, if B is a state removal, then an edge $s \overset{e}{\longrightarrow} s'$ of A can only be missing in B if its source s or its target s' is removed. By contrast, the latter is not necessarily the case if B is an event- or an edge removal.

Example 2.6. Consider the LTS A of Figure 1. The LTS B of Figure 2 is a state removal of A resulting by removing the state s_3, i.e., $\mathfrak{S} = \{s_3\}$. B is also an edge removal, where $\mathfrak{K} = \{s_2 \overset{x}{\longrightarrow} s_3\}$. However, this LTS is not an event removal, since x belongs to B, but not all x-labeled edges of A
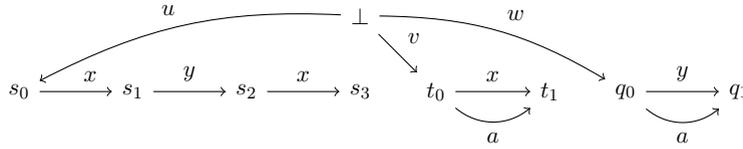
Figure 1.   The LTS A.
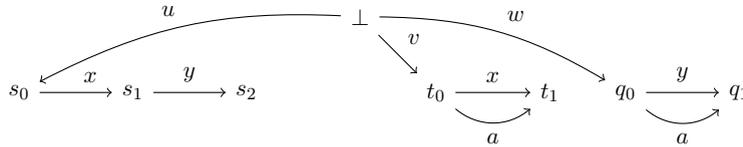
Figure 2.   The state removal B of A that results by removing the state s_3.
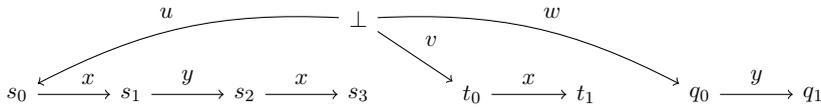
Figure 3.   The event removal C of A that results by removing the event a.

are present. The LTS C of Figure 3 is an event removal of A such that $\mathfrak{E} = \{a\}$. C is also an edge removal and $\mathfrak{K} = \{t_0 \xrightarrow{a} t_1, q_0 \xrightarrow{a} q_1\}$, but is is not a state removal.

Note that, if A is reachable, this is not always the case for B. For instance, in Figure 1, if we remove event x, or state s_1 or edge $s_0 \xrightarrow{x} s_1$, the result is not reachable. In the following, we shall only consider reachable removals, however. On the contrary, if A is finite, so is B; in the following we shall only consider finite LTSs.

2.2. Petri nets, and implementations

Petri nets are the target model with which we want to implement LTS:

Definition 2.7. (Petri Nets)
A *(weighted) Petri net* $N = (P, T, f, M_0)$ consists of finite and disjoint sets of *places* P and *transitions* T, a (total) *flow* $f : ((P \times T) \cup (T \times P)) \to \mathbb{N}$ and an *initial marking* $M_0 : P \to \mathbb{N}$. A transition $t \in T$ can *fire* or *occur* in a marking $M : P \to \mathbb{N}$, denoted by $M \xrightarrow{t}$, if $M(p) \geq f(p, t)$ for all places $p \in P$. The firing of t in marking M leads to the marking $M'(p) = M(p) - f(p, t) + f(t, p)$ for all $p \in P$, denoted by $M \xrightarrow{t} M'$. This notation extends to sequences $w \in T^*$ and the *reachability set* $RS(N) = \{M \mid \exists w \in T^* : M_0 \xrightarrow{w} M\}$ contains all the reachable markings of N. The *reachability graph* of N is the LTS $A_N = (RS(N), T, \delta, M_0)$, where, for every reachable marking M of N and transition $t \in T$, the transition function δ of A_N is defined by $\delta(M, t) = M'$ if and only if $M \xrightarrow{t} M'$.

Simulations between A and A_N define how a net N implements an LTS A:

Definition 2.8. (Implementations)
If A is an LTS and N is a Petri net, then N is an *embedding* of A if $A \hookrightarrow A_N$; N is a *language-simulation* of A, if $A \triangleright A_N$, and N is a *realization* of A, if $A \cong A_N$. We say N *implements* A, if it is an embedding or a language-simulation or a realization of A.

2.3. Regions, separation properties, and synthesized nets

If a Petri net N implements an LTS A, then the events of A are the transitions of N. We obtain the remaining components of N, that is, places, flow and initial marking, by regions of A:

Definition 2.9. (Region)
A *region* $R = (sup, con, pro)$ of an LTS $A = (S, E, \delta, \iota)$ consists of the mappings *support* $sup :$ $S \to \mathbb{N}$ and *consume* and *produce* $con, pro : E \to \mathbb{N}$ such that if $s \xrightarrow{e} s'$ is an edge of A, then $con(e) \leq sup(s)$ and $sup(s') = sup(s) - con(e) + pro(e)$.

Remark 2.10. A region $R = (sup, con, pro)$ is *implicitly* completely defined by $sup(\iota)$, con and pro: Since A is reachable, there is a path $\iota \xrightarrow{e_1} \ldots \xrightarrow{e_n} s_n$ such that $s = s_n$ for every state $s \in S$. Consequently, we inductively obtain $sup(s_{i+1})$ by $sup(s_{i+1}) = sup(s_i) - con(e_{i+1}) + pro(e_{i+1})$ for all $i \in \{0, \ldots, n-1\}$ and $s_0 = \iota$. Hence, for the sake of simplicity, we often present regions only implicitly, since sup and thus R can be obtained from $sup(\iota)$, con and pro (one has to check however that, if two edges lead to the same state, the yielded support is the same). For an even more compact presentation, for $c, p \in \mathbb{N}$, we group events with the same "behavior" together by $\mathcal{T}_{c,p}^R = \{e \in E \mid con(e) = c \text{ and } pro(e) = p\}$.

Definition 2.11. (Effect)
Let $A = (S, E, \delta, \iota)$ be a LTS, and (sup, con, pro) be a region of A. If $e \in E$, then we say that $\text{eff}(e) = -con(e) + pro(e)$ (in $\mathbb{Z}$) is the *effect* of e.

Lemma 2.12. Let $A = (S, E, \delta, \iota)$ be a LTS, and (sup, con, pro) be a region of A. If $s_0 \xrightarrow{e_1} \ldots \xrightarrow{e_n} s_n$ is a path in A, then it holds $sup(s_n) = sup(s_0) + \sum_{i=0}^{n} \text{eff}(e_i)$.

Proof:
Since R is a region, we have $sup(s_{i+1}) = sup(s_i) - con(e_{i+1}) + pro(e_{i+1})$, and thus $sup(s_{i+1}) = sup(s_i) + \text{eff}(e_{i+1})$ for all $i \in \{0, \ldots, n-1\}$. Hence, the claim follows easily by induction. □

If there is an implementing net N for A, then each place correspond to a region $R = (sup, con, pro)$ of A: $con(e)$ and $pro(e)$ model $f(R, e)$ and $f(e, R)$ for all transitions e, respectively, and $sup(\iota)$ models the initial marking $M_0(R)$. In particular, every set of regions defines a *synthesized net*:

Definition 2.13. (Synthesized Net)
A set $\mathcal{R}$ of regions of LTS $A = (S, E, \delta, \iota)$ defines the *synthesized net* $N_A^{\mathcal{R}} = (\mathcal{R}, E, f, M_0)$, where $f(R, e) = con(e)$ and $f(e, R) = pro(e)$ and $M_0(R) = sup(\iota)$ for all $R = (sup, con, pro) \in \mathcal{R}$ and $e \in E$.

If the synthesized net is an embedding or a realization of A, then distinct states of A correspond to distinct markings of the net. The net $N_A^{\mathcal{R}}$ satisfies this requirement if the set $\mathcal{R}$ of regions present the *state separation property*:

Definition 2.14. (State Separation Property)
A pair (s, s') of distinct states of LTS $A = (S, E, \delta, \iota)$ defines a *states separation atom* (SSA). A region $R = (sup, con, pro)$ *solves* (s, s') if $sup(s) \neq sup(s')$. We say a state s is *solvable* if, for every $s' \in S \setminus \{s\}$, there is a region that solves the SSA (s, s'). If every SSA or, equivalently, every state of A is solvable, then A has the *state separation property* (SSP).

If the net is a language-simulation or a realization, then the firing of a transition e must be inhibited at a marking M whenever the event e does not occur at the state s that correspond to M via φ. This is ensured if $\mathcal{R}$ witnesses the *event/state separation property*:

Definition 2.15. (Event/State Separation Property)
A pair (e, s) of event e and state s of LTS $A = (S, E, \delta, \iota)$ such that $s \xrightarrow{\neg e}$ defines an *event/state separation atom* (ESSA). A region $R = (sup, con, pro)$ *solves* (e, s) if $sup(s) < con(e)$. We say an event e is *solvable* if, for every $s \in S$ with $s \xrightarrow{\neg e}$, there is a region that solves the ESSA (e, s). If every ESSA or, equivalently, every event of A is solvable, then A has the *event/state separation property* (ESSP).

A set $\mathcal{R}$ of regions of A is called a *witness* for the SSP or the ESSP (of A) if, for every SSA or ESSA, there is a region in $\mathcal{R}$ that solves it. The next lemma is based on [15, p. 162] and [15, p. 214 ff.] and states in which case the existence of a witness and the existence of an implementation are equivalent; this will allow us to formulate our decision problems rather on the notion of witnesses than on the notion of implementations (notice that Petri nets correspond to the type of nets τ_{PT} in [15, p. 130]):

Lemma 2.16. ([15])
Let A be an LTS and N a Petri net.

1. $A \hookrightarrow A_N$ if and only if there is a witness $\mathcal{R}$ for the SSP of A and $N = N_A^{\mathcal{R}}$;

2. $A \triangleright A_N$ if and only if there is a witness $\mathcal{R}$ for the ESSP of A and $N = N_A^{\mathcal{R}}$;

3. $A \cong A_n$ if and only if there is a witness $\mathcal{R}$ for both the SSP and the ESSP of A and $N = N_A^{\mathcal{R}}$.

4. Whether A has the SSP or the ESSP can be decided and, in case of a positive decision, a witness can be computed in polynomial time.

Example 2.17. Let $A = (Z, E, \delta, \bot)$ be the LTS of Figure 1. The following implicitly defined region $R = (sup, con, pro)$ solves all SSA of A: $sup(\bot) = 8$ and $\mathcal{T}_{5,0}^R = \{v\}$, $\mathcal{T}_{7,0}^R = \{w\}$ and $\mathcal{T}_{1,0}^R = E \setminus \{v, w\}$. According to Remark 2.10, one obtains R explicitly: $sup(s_i) = 7-i$ for all $i \in \{0, \ldots, 3\}$ and $sup(t_0) = 3$, $sup(t_1) = 2$, $sup(q_0) = 1$ and $sup(q_1) = 0$. The set $\mathcal{R} = \{R\}$ witnesses the SSP of A and the net $N = N_A^{\mathcal{R}}$ is an embedding of A. Figure 4 shows N (top) and its reachability graph A_N (bottom). The injective simulation map φ is defined by $\varphi(\bot) = 8$, $\varphi(s_i) = 7-i$ for all $i \in \{0, \ldots, 3\}$ and $\varphi(t_0) = 3$, $\varphi(t_1) = 2$, $\varphi(q_0) = 1$ and $\varphi(q_1) = 0$.
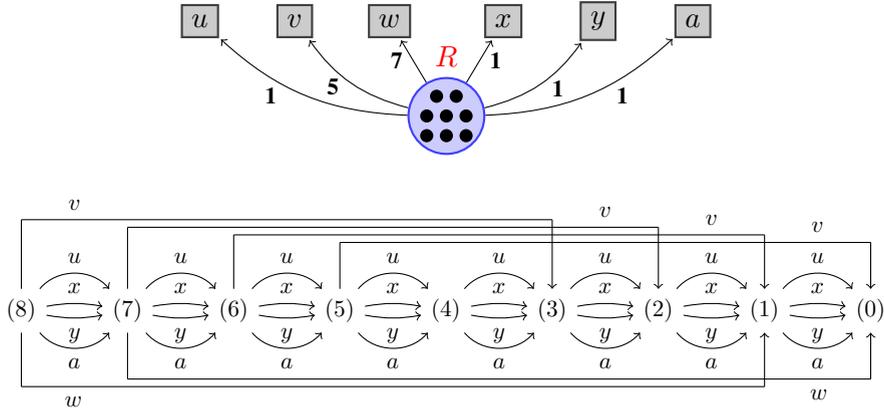
Figure 4.   The net $N = N_A^{\mathcal{R}}$ and its reachability graph A_N according to Example 2.17.

Example 2.18. The LTS $A = (Z, E, \delta, \perp)$ of Figure 1 does not have the ESSP, since the ESSA $\alpha = (x, s_1)$ is not solvable. This can be seen as follows: Assume $R = (sup, con, pro)$ is a region that solves α, that is, $con(x) > sup(s_1)$. Since x occurs at s_0, we have $con(x) \leq sup(s_0)$. By $sup(s_1) = sup(s_0) - con(x) + pro(x)$ and $con(x) > sup(s_1)$, this implies $con(x) > pro(x)$. By $t_0 \xrightarrow{x} t_1$, this also implies $sup(t_0) > sup(t_1)$ and thus $con(a) > pro(a)$ by $t_0 \xrightarrow{a} t_1$. On the other hand, x occurs at s_2, which implies $con(x) \leq sup(s_2)$. By $con(x) > sup(s_1)$ and $s_1 \xrightarrow{y} s_2$, this is only possible if $con(y) < pro(y)$, which implies $sup(q_0) < sup(q_1)$ and thus $con(a) < pro(a)$. This is a contradiction. Hence, α is not solvable.

Example 2.19. The LTS of Figure 1 does not have the ESSP, since the ESSA $\alpha = (x, s_1)$ is not solvable, as we just saw. However, for the LTS B of Figure 2, which is a state removal for A, there is a region $R = (sup, con, pro)$ that solves the ESSA (x, s_1), which is implicitly defined as follows: $sup(\perp) = 2$ and $\mathcal{T}_{2,1}^R = \{x\}$ and $\mathcal{T}_{1,0}^R = \{a, y\}$ and $\mathcal{T}_{0,0}^R = E \setminus \{a, y, x\}$. One finds out that the remaining ESSA of B are also solvable. Hence, B has the ESSP and the SSP. If the modeler comes to the conclusion that the events x and y, their corresponding edges and all of their sources and targets are essential for the modeled behavior, then a realizable behavior can also be obtained as the event removal C of A as defined in Figure 3. A Region $R' = (sup', con', pro')$ solving (x, s_1) in C, is then implicitly defined as follows: $sup(\perp) = 1$ and $\mathcal{T}_{1,0}^{R'} = \{x\}$ and $\mathcal{T}_{0,1}^{R'} = \{y\}$ and $\mathcal{T}_{0,0}^{R'} = E \setminus \{x, y\}$.

3.   The complexity of Edge Removal

According to Lemma 2.16, the question whether a particular implementation for a given LTS exists is equivalent to the question whether the LTS has the separation properties that correspond to the implementation. In this section, we are interested in modifying an LTS into an implementable one by the removal of a bounded number of edges. In particular, we are interested in the computational complexity of the following decision problems:

EDGE REMOVAL FOR EMBEDDING

Input: An LTS $A = (S, E, \delta, \iota)$, a natural number κ.

Question: Does there exist an edge removal B for A by $\mathfrak{K}$ that has the SSP and satisfies $|\mathfrak{K}| \leq \kappa$?

EDGE REMOVAL FOR LANGUAGE-SIMULATION

Input: An LTS $A = (S, E, \delta, \iota)$, a natural number κ.

Question: Does there exist an edge removal B for A by $\mathfrak{K}$ that has the ESSP and satisfies $|\mathfrak{K}| \leq \kappa$?

EDGE REMOVAL FOR REALIZATION

Input: An LTS $A = (S, E, \delta, \iota)$, a natural number κ.

Question: Does there exist an edge removal B for A by $\mathfrak{K}$ that has the ESSP and the SSP and satisfies $|\mathfrak{K}| \leq \kappa$?

3.1.   Edge Removal aiming at language-simulation or realization

The following theorem characterizes the complexity of both Edge Removal for Language-simulation, and Edge Removal for realization:

Theorem 3.1. Edge Removal for Language-simulation as well as Edge Removal for realization are NP-complete.

It is easy to see that the addressed problems are in NP: If there is an adequate edge removal for A, then a Turing machine can guess $\mathfrak{K}$ by a non-deterministic computation in a time polynomial in the size of the input. After that, the machine can deterministically (and polynomially) compute B and, since the size of B is bounded by the size of A, it can compute a witness for the relevant property of B in a time polynomial in the size of the input by Lemma 2.16. Hence, in order to complete the NP-completeness part of Theorem 3.1, it remains to prove the NP-hardness of the problems. This proof is based on a polynomial reduction of the problem HITTING SET:

HITTING SET (HS)

Input: A triple $(\mathfrak{U}, M, \lambda)$ that consist of a finite set $\mathfrak{U}$, a set $M = \{M_0, \ldots, M_{m-1}\}$ of subsets of $\mathfrak{U}$ and a natural number λ.

Question: Does there exist a hitting set Z for $(\mathfrak{U}, M)$, that is, $Z \subseteq \mathfrak{U}$ and $Z \cap M_i \neq \emptyset$ for all $i \in \{0, \ldots, m-1\}$, that satisfies $|Z| \leq \lambda$?

Example 3.2. The instance $(\mathfrak{U}, M, 4)$ such that $\mathfrak{U} = \{X_0, \ldots, X_5\}$ and $M = \{M_0, \ldots, M_8\}$, where $M_0 = \{X_0, X_1\}$, $M_1 = \{X_0, X_3\}$, $M_2 = \{X_0, X_5\}$, $M_3 = \{X_1, X_2\}$, $M_4 = \{X_1, X_5\}$, $M_5 = \{X_2, X_3\}$, $M_6 = \{X_2, X_4\}$, $M_7 = \{X_3, X_4\}$, $M_8 = \{X_4, X_5\}$, has, for example, the hitting set $Z = \{X_0, X_2, X_3, X_5\}$ and thus allows a positive decision.

Without loss of generality, we may assume that $\lambda \leq |\mathfrak{U}|$, and, for all $i \in \{0, \ldots, m - 1\}$, $|\mathfrak{U}| > |M_i| > 1$. Indeed, if $M_i = \emptyset$ we know there is no solution, if $M_i = \{x\}$ we know x must be in Z and the problem may be reduced to another, smaller, one.

Theorem 3.3. ([22])
Hitting Set is NP-complete.

In the following, unless explicitly stated otherwise, let $(\mathfrak{U}, M, \lambda)$ be an arbitrary but fixed input of HS such that $\mathfrak{U} = \{X_0, \ldots, X_{n-1}\}$ and $M = \{M_0, \ldots, M_{m-1}\}$, where $M_i = \{X_{i_0}, \ldots, X_{i_{m_i-1}}\}$ (and thus $|M_i| = m_i$) for all $i \in \{0, \ldots, m - 1\}$. For technical reasons, we assume without loss of generality that $i_0 < \cdots < i_{m_i-1}$ for the elements $X_{i_0}, \ldots, X_{i_{m_i-1}}$ of the set M_i for all $i \in \{0, \ldots, m - 1\}$.

The General Idea of the Reduction.   We reduce $(\mathfrak{U}, M, \lambda)$ to a pair (A, κ) of LTS A and a natural number κ in such a way that the following implications are satisfied: If there is an edge removal B of A that removes at most κ edges and has the ESSP, then there is a hitting set with at most λ elements for $(\mathfrak{U}, M)$. Conversely, if there is a hitting set with λ elements for $(\mathfrak{U}, M)$, then there is an edge removal B of A that removes at most κ edges and has both the ESSP and the SSP. Notice that such a reduction proves the hardness of both EDGE REMOVAL FOR LANGUAGE-SIMULATION and EDGE REMOVAL FOR REALIZATION. The announced LTS A consists of several components. Just as it is common in the world of reductions, we refer to these components as *gadgets*.

The Reduction.   For a start, we define $\kappa = \lambda$. Moreover, for every $i \in \{0, \ldots, m - 1\}$, and for every $j \in \{0, \ldots, \kappa\}$, the LTS A has the following gadget $T_{i,j}$, that uses the elements of $M_i = \{X_{i_0}, \ldots, X_{m_i-1}\}$ as events:

$$T_{i,j} = t_{i,j,0} \xrightarrow{X_{i_0}} t_{i,j,1} \xrightarrow{X_{i_1}} \cdots \xrightarrow{X_{i_{m_i-1}}} t_{i,j,m_i} \xrightarrow{X_{i_0}} t_{i,j,m_i+1}$$

Notice that X_{i_0} is both the first, and the last event along the path $T_{i,j}$, while all the other events of $M_i \setminus \{X_{i_0}\}$ occur exactly once. Moreover, since there are essentially $\kappa + 1$ copies of the same sequence of events $X_{i_0} \ldots X_{m_i-1} X_{i_0}$, after the removal of at most κ edges, there is a least one of these sequences left.

For every $i \in \{0, \ldots, n - 1\}$, the LTS has the following gadget F_i that has X_i as event, as well as an a_ℓ-labeled edge, for all $\ell \in \{0, \ldots, \kappa\}$, with the same direction:
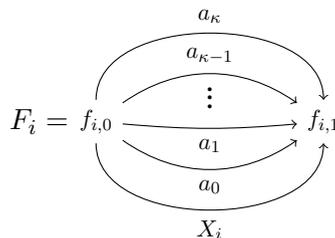
Finally, the LTS A has the initial state ι, and uses the following edges to connect ι to the just introduced gadgets:

- For all $i \in \{0, \ldots, m-1\}$, and all $j \in \{0, \ldots, \kappa\}$, the LTS A has the edge $\iota \xrightarrow{u_i^j} t_{i,j,0}$.

- For all $i \in \{0, \ldots, n-1\}$, the LTS A has the edge $\iota \xrightarrow{v_i} f_{i,0}$.

The resulting LTS is $A = (S, E, \delta, \iota)$. It is easy to check that A is reachable. Figure 5 provides (essentially) the gadgets of the LTS A that are based on the input of Example 3.2. In what follows, we prove that A satisfies the conditions stated above in the description of the general idea of reduction. (Recall that $\mathfrak{K}$ refers to the set of edges removed from A, where the edge removal is clear from the context).
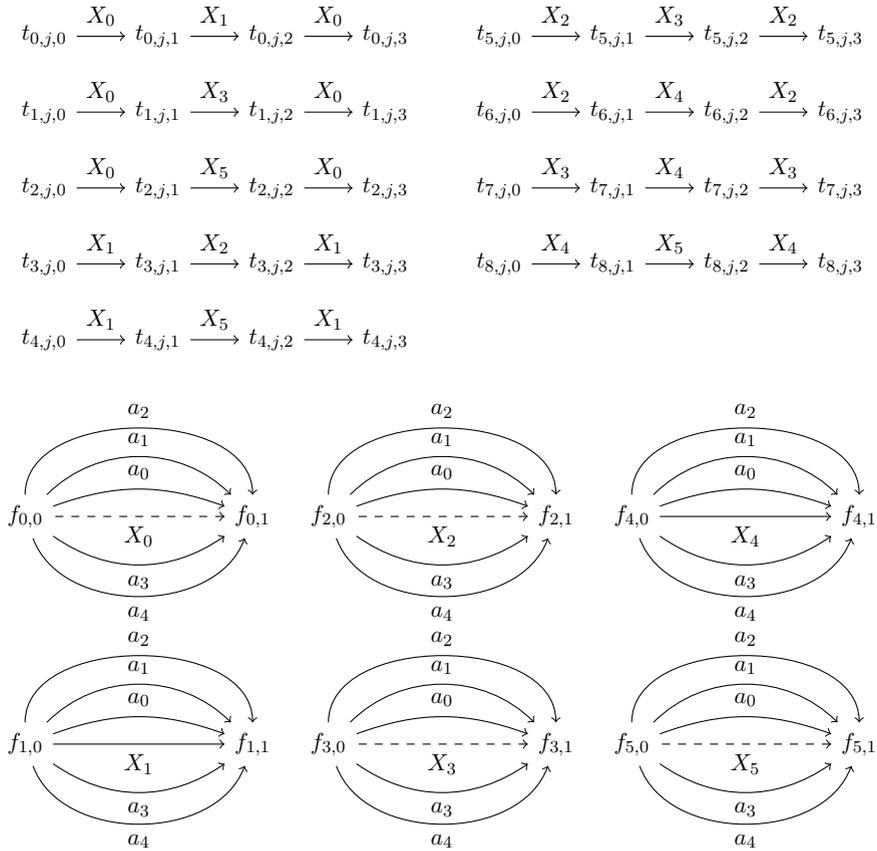
$$t_{0,j,0} \xrightarrow{X_0} t_{0,j,1} \xrightarrow{X_1} t_{0,j,2} \xrightarrow{X_0} t_{0,j,3} \qquad t_{5,j,0} \xrightarrow{X_2} t_{5,j,1} \xrightarrow{X_3} t_{5,j,2} \xrightarrow{X_2} t_{5,j,3}$$

$$t_{1,j,0} \xrightarrow{X_0} t_{1,j,1} \xrightarrow{X_3} t_{1,j,2} \xrightarrow{X_0} t_{1,j,3} \qquad t_{6,j,0} \xrightarrow{X_2} t_{6,j,1} \xrightarrow{X_4} t_{6,j,2} \xrightarrow{X_2} t_{6,j,3}$$

$$t_{2,j,0} \xrightarrow{X_0} t_{2,j,1} \xrightarrow{X_5} t_{2,j,2} \xrightarrow{X_0} t_{2,j,3} \qquad t_{7,j,0} \xrightarrow{X_3} t_{7,j,1} \xrightarrow{X_4} t_{7,j,2} \xrightarrow{X_3} t_{7,j,3}$$

$$t_{3,j,0} \xrightarrow{X_1} t_{3,j,1} \xrightarrow{X_2} t_{3,j,2} \xrightarrow{X_1} t_{3,j,3} \qquad t_{8,j,0} \xrightarrow{X_4} t_{8,j,1} \xrightarrow{X_5} t_{8,j,2} \xrightarrow{X_4} t_{8,j,3}$$

$$t_{4,j,0} \xrightarrow{X_1} t_{4,j,1} \xrightarrow{X_5} t_{4,j,2} \xrightarrow{X_1} t_{4,j,3}$$

Figure 5. For a fixed $j \in \{0, \ldots, 3\}$, the gadgets $T_{0,j}, \ldots, T_{8,j}$ and $F_0, \ldots, F_5$ of the LTS A (Section 3) based on Example 3.2. Dashed lines correspond to edges that are removed in accordance to the edge removal B defined for the proof of Lemma 3.5 and correspond to the hitting set $\{X_0, X_2, X_3, X_5\}$.

Lemma 3.4. If there is an edge removal B of A by $\mathfrak{K}$ that satisfies $|\mathfrak{K}| \leq \kappa$ and has the ESSP, then there is a hitting set of size at most $\lambda = \kappa$ for $(\mathfrak{U}, M)$.

Proof:

Let $B = (S', E', \delta', \iota)$ be an edge removal of A by $\mathfrak{K}$ that satisfies $|\mathfrak{K}| \leq \kappa$ and has the ESSP. First, we may observe that, without loss of generality, we may assume that all the members of $\mathfrak{K}$ are of the form $f_{i,0} \xrightarrow{X_i} f_{i,1}$, so that in particular B remains reachable.

Indeed, for each $i \in \{0, \ldots, m-1\}$, since there are more than κ copies of $T_{i,0}$, at least one of them is completely present in B (including the edge from ι). Let us assume it is $T_{i,j}$ for $j \in \{0, \ldots, \kappa\}$. Then we may remove from $\mathfrak{K}$ all the members of the form $t_{i,k,h} \xrightarrow{X_\ell} t_{i,k,h+1}$ or $\iota \xrightarrow{u_i^k} t_{i,k,0}$ and reinsert them in $T_{i,k}$. This reduces the size of $\mathfrak{K}$ as well as the set of ESSAs and, by choosing $con(u_i^k) = con(u_i^j)$, and $pro(u_i^k) = pro(u_i^j)$, we get regions that solve the same ESSAs in $T_{i,k}$ as in $T_{i,j}$.

Similarly, since the set $\{a_0, \ldots, a_\kappa\}$ contains $\kappa + 1$ elements and $|\mathfrak{K}| \leq \kappa$, there is an index $h \in \{0, \ldots, \kappa\}$ such that all a_h-labeled edges and thus particularly the edge $f_{i,0} \xrightarrow{a_h} f_{i,1}$ is present in B for all $i \in \{0, \ldots, n-1\}$. Hence, if we reinsert in B all the missing a_ℓ-labeled edges, $\ell \in \{0, \ldots, \kappa\}$, as well as the missing v_i-labeled ones, each region of the old B extends immediately in a region of the new one, solving all the needed remaining ESSAs.

In the following, we show that this implies that the set $Z = \{X_i \in \mathfrak{U} \mid f_{i,0} \xrightarrow{X_i} f_{i,1} \in \mathfrak{K}\}$ defines a hitting set, with at most $\lambda = \kappa$ elements, for $(\mathfrak{U}, M)$.

Since B has the ESSP, for each $i \in \{0, \ldots, m-1\}$ and $j \in \{0, \ldots, \kappa\}$, there is a region that solves $\alpha = (X_{i_0}, t_{i,j,1})$. Let $R = (sup, con, pro)$ be such a region. We argue that there is some $\ell \in \{0, \ldots, m_i - 1\}$ such that the X_{i_ℓ}-labeled edge $f_{i_\ell,0} \xrightarrow{X_{i_\ell}} f_{i_\ell,1} \in \mathfrak{K}$, i.e., is not present in B.

Assume, for a contradiction, that $f_{i_\ell,0} \xrightarrow{X_{i_\ell}} f_{i_\ell,1} \in B$ for all $\ell \in \{0, \ldots, m_i - 1\}$. By $t_{i,j,0} \xrightarrow{X_{i_0}}$, we have $con(X_{i_0}) \leq sup(t_{i,j,0})$; since R solves α, we have that $con(X_{i_0}) > sup(t_{i,j,1})$. This implies $con(X_{i_0}) > pro(X_{i_0})$, and thus $sup(f_{i_0,0}) > sup(f_{i_0,1})$. Moreover, by $t_{i,j,m_i} \xrightarrow{X_{i_0}}$, we have that $con(X_{i_0}) \leq sup(t_{i,j,m_i})$, which, by $con(X_{i_0}) > sup(t_{i,j,1})$, implies

$$sup(t_{i,j,1}) < sup(t_{i,j,m_i}) = sup(t_{i,j,1}) + \sum_{\ell=1}^{m_i-1} \text{eff}(X_{i_\ell})$$

and thus $\sum_{\ell=1}^{m_i-1} \text{eff}(X_{i_\ell}) > 0$. In particular, there is an $\ell \in \{1, \ldots, m_i - 1\}$ (so that $i_\ell \neq i_0$) such that $con(X_{i_\ell}) < pro(X_{i_\ell})$, which implies $sup(f_{i_\ell,0}) < sup(f_{i_\ell,1})$.

Since the edges $f_{i_0,0} \xrightarrow{a_0} f_{i_0,1}$ and $f_{i_\ell,0} \xrightarrow{a_0} f_{i_\ell,1}$ are both present in B, by $sup(f_{i_0,0}) > sup(f_{i_0,1})$, we obtain $con(a_0) > pro(a_0)$ and, by $sup(f_{i_\ell,0}) < sup(f_{i_\ell,1})$, we get $con(a_0) < pro(a_0)$, which is a contradiction. Consequently, $\{f_{i_0,0} \xrightarrow{X_{i_0}} f_{i_0,1}, f_{i_\ell,0} \xrightarrow{X_{i_\ell}} f_{i_\ell,1}\} \cap \mathfrak{K} \neq \emptyset$ and thus $M_i \cap Z \neq \emptyset$. Since i was arbitrary, we have $M_i \cap Z \neq \emptyset$ for all $i \in \{0, \ldots, m-1\}$ and Z defines a hitting set for $(\mathfrak{U}, M)$. $\square$

Conversely, we show:

Lemma 3.5. If there is a hitting set with at most λ elements for $(\mathfrak{U}, M)$, then there is an edge removal B of A with $|\mathfrak{K}| \leq \lambda$ that has the ESSP and the SSP.

Proof:

Let $Z = \{X_{j_0}, \ldots, X_{j_{\lambda-1}}\}$, where $j_0, \ldots, j_{\lambda-1} \in \{0, \ldots, n-1\}$, be a hitting set with λ elements for $(\mathfrak{U}, M)$. (Notice that every hitting set with at most λ elements for $(\mathfrak{U}, M)$ can be easily extended to a hitting set with exactly λ elements.)Moreover, let $B = (S, E, \delta', \iota)$ be the LTS that originates from A by removing, for all $\ell \in \{0, \ldots, \lambda-1\}$, the edges $f_{j_\ell,0} \xrightarrow{X_{j_\ell}} f_{j_\ell,1}$, and nothing else. Obviously, B is an edge removal of A such that $|\mathfrak{K}| \leq \lambda = \kappa$. In the following, we argue that B has the SSP and the ESSP. Let $I = \bigcup_{i=0}^{m-1}(\bigcup_{j=0}^{\kappa}\{t_{i,j,0}\}) \cup \bigcup_{i=0}^{n-1}\{f_{i,0}\}$ be the set of the initial states of the gadgets of B, and let $E' = \mathfrak{U} \cup \{a_0, \ldots, a_\kappa\}$ be the set of the events of them.

For a start, it is easy to see that if s and s' are states of different gadgets (or if one of it equals ι), then the SSA (s, s') is solvable. Likewise, one easily finds out that if e is an event and s is a state that do not belong to the same gadget of B, then the ESSA (e, s) is solvable. Hence, in the following, we restrict our attention to the solvability of separation atoms whose components belong to the same gadgets. Notice that we define the following regions implicitly according to Remark 2.10. For an even more compact representation, until explicitly stated otherwise, we assume ι to be mapped to 0, and we define additionally only the support of the initial states of the gadgets, and con, and pro for the events of the gadgets. According to Remark 2.10, one computes the support values of the remaining states of the gadgets. Moreover, since the u_i^j's, and the v_i's are unique, it is easy to see that there con and pro values can always be chosen suitably.

The following region $R_0 = (sup_0, con_0, pro_0)$ solves all the remaining SSA and, for all $i \in \{0, \ldots, n-1\}$ and all $e \in \mathfrak{U} \cup \{a_0, \ldots, a_\kappa\}$, the ESSA $(f_{i,1}, e)$: $sup_0(\iota) = 0$; for all $i \in \{0, \ldots, m-1\}$, $\mathcal{T}_{0,m_i}^{R_0} = \{u_i^0, \ldots, u_i^\kappa\}$; $\mathcal{T}_{0,1}^{R_0} = \{v_0, \ldots, v_{n-1}\}$ and $\mathcal{T}_{1,0} = \mathfrak{U} \cup \{a_0, \ldots, a_\kappa\}$.

Hence, it remains to consider the ESSA of the $T_{i,j}$'s. Let $i \in \{0, \ldots, m-1\}$ and $j \in \{0, \ldots, \kappa\}$ be arbitrary but fixed.

In the following, we investigate the cases $X_{i_0} \in Z$ and $X_{i_0} \notin Z$ separately.

Case $X_{i_0} \notin Z$: Unless explicitly stated otherwise, let $\ell \in \{1, \ldots, m_i - 1\}$ be the smallest index such that $X_{i_\ell} \in Z$ (which exists since Z is a hitting set.) Notice that there are $\ell - 1$ variable events between the first occurrence of X_{i_0}, at the beginning of $T_{i,j}$, and X_{i_ℓ}, and that there are $m_i - (\ell + 1)$ variable events between X_{i_ℓ} and the second occurrence of X_{i_0} at the end of $T_{i,j}$.

The following region $R_1 = (sup_1, con_1, pro_1)$ solves (X_{i_0}, s) for all $s \in \{t_{i,j,1}, \ldots, t_{i,j,\ell},\} \cup \{t_{i,j,m_i+1}\}$: for all $s \in I$, if $s = t_{i,j,0}$, then $sup(t_{i,j,0}) = \ell$, otherwise $sup(s) = \ell \cdot \sum_{i=0}^{m-1} m_i$ (that is, $sup(s)$ is just chosen "big enough") and $\mathcal{T}_{\ell,\ell-1}^{R_1} = \{X_{i_0}\}$ (this implies $sup(t_{i,j,1}) = \ell - 1$); $\mathcal{T}_{1,0}^{R_1} = \mathfrak{U} \setminus \{X_{i_\ell}\} \cup \{a_0, \ldots, a_\kappa\}$ and $\mathcal{T}_{0,m_i-1}^{R_1} = \{X_{i_\ell}\}$.

The following region $R_2 = (sup_2, con_2, pro_2)$ solves (X_{i_0}, s) for all $s \in \{t_{i,j,\ell+1}, \ldots, t_{i,j,m_i-1}\}$: for all $s \in I$, if $s = t_{i,j,0}$, then $sup(t_{i,j,0}) = m_i - \ell - 1$, otherwise $sup_2(s) = \ell \cdot \sum_{i=0}^{m-1} m_i$ (again, $sup(s)$ is "big enough") and $\mathcal{T}_{m_i-\ell-1,m_i-\ell}^{R_2} = \{X_{i_0}\}$; $\mathcal{T}_{0,1}^{R_2} = \mathfrak{U} \setminus \{X_{i_\ell}\} \cup \{a_0, \ldots, a_\kappa\}$ and $\mathcal{T}_{m_i,0}^{R_2} = \{X_{i_\ell}\}$. Notice that R_2 implies $sup_2(t_{i,j,\ell+1}) = 0$ by $m_i - \ell - 1 + (\ell - 1) = m_i$ and thus $sup_2(t_{i,j,m_i}) = m_i - \ell - 1$.

This proves the solvability of X_{i_0}. In the following, we argue for the solvability of the other variables events in $T_{i,j}$. Let $\ell \in \{1, \ldots, m-1\}$ be arbitrary but fixed (that is, X_{i_ℓ} is not necessarily in Z). Notice that X_{i_ℓ} is preceded by exactly ℓ variable events and followed by $m_i - \ell$ variable events (since X_{i_0} occurs twice).

The following region $R_3 = (sup_3, con_3, pro_3)$ solves (X_{i_ℓ}, s) for all $s \in \{t_{i,j,0}, \ldots, t_{i,j,\ell-1}\}$: for all $s \in I$, if $s = t_{i,j,0}$, then $sup(t_{i,j,0}) = 0$, otherwise $sup_3(s) = \ell \cdot \sum_{i=0}^{m-1} m_i$; $\mathcal{T}_{\ell,\ell+1}^{R_3} = \{X_{i_\ell}\}$ and $\mathcal{T}_{0,1}^{R_3} = E' \setminus \{X_{i_\ell}\}$.

The following region $R_4 = (sup_4, con_4, pro_4)$ solves (X_{i_ℓ}, s) for all $s \in \{t_{i,j,\ell+1}, \ldots, t_{i,j,m_i+1}\}$: for all $s \in I$, if $s = t_{i,j,0}$, then $sup(t_{i,j,0}) = m_i+1$, otherwise $sup_4(s) = \ell \cdot \sum_{i=0}^{m-1} m_i$; $\mathcal{T}_{m_i-\ell+1,m_i-\ell}^{R_4} = \{X_{i_\ell}\}$ and $\mathcal{T}_{1,0}^{R_4} = E' \setminus \{X_{i_\ell}\}$.

Since ℓ was arbitrary, this proves that all ESSA of $T_{i,j}$ are solvable.

Case $X_{i_0} \in Z$: Recall that there are $m_i - 1$ events between the two occurrences of X_{i_0} in $T_{i,j}$.

The following region $R_5 = (sup_5, con_5, pro_5)$ solves (X_{i_0}, s) for all $s \in \{t_{i,j,1}, \ldots, t_{i,j,m_i-1}\} \cup \{t_{i,j,m_i+1}\}$: for all $s \in I$, if $s = t_{i,j,0}$, then $sup_5(s) = m_i - 1$, otherwise $sup_4(s) = \ell \cdot \sum_{i=0}^{m-1} m_i$; $\mathcal{T}_{m_i-1,0}^{R_5} = \{X_{i_0}\}$ and $\mathcal{T}_{0,1}^{R_5} = E' \setminus \{X_{i_0}\}$.

Similar to the former case, one proves, for all (e, s) with $e \in \{X_{i_1}, \ldots, X_{i_{m_i-1}}\}$ and $s \in S(T_{i,j})$, if (e, s) is an ESSA of $T_{i,j}$, then (e, s) is solvable.

Altogether, by the arbitrariness of i and j, we conclude that B has the ESSP. This completes the proof. $\qquad\square$

3.2. Edge Removal aiming at embedding

In the conclusion of [19], Schlachter and Wimmel mention en passant that their reduction-technique imply that Edge Removal for embedding is NP-complete. The following theorem re-formulates this claim, and is followed by a full (and new) proof.

Theorem 3.6. EDGE REMOVAL aiming at embedding is NP-complete.

Like in the previous subsection, it is easy to see that Edge Removal aiming at embedding is in NP.

The General Idea of the Reduction. To prove hardness, in the following, we reduce $(\mathfrak{U}, M, \lambda)$ polynomially to a pair (A, κ) such that there is hitting set of size at most λ for $(\mathfrak{U}, M)$ if the removal of at most λ edges of A yields an LTS B that has the SSP. Conversely, if $(\mathfrak{U}, M, \lambda)$ has a fitting hitting set, then there are at most κ edges of A whose removal yields an LTS that has the SSP.

**The Reduction.** Again, we shall have $\kappa = \lambda$.

For every $i \in \{0, \ldots, m-1\}$, and for every $j \in \{0, \ldots, \kappa\}$, the LTS A has the following gadget $T_{i,j}$, that uses the event k_i, and the elements of $M_i = \{X_{i_0}, \ldots, X_{m_i-1}\}$ as events:

$$T_{i,j} = \; t_{i,j,0} \xrightarrow{\;k_i\;} t_{i,j,1} \xrightarrow{\;X_{i_0}\;} \cdots \xrightarrow{\;X_{i_{m_i-2}}\;} t_{i,j,m_i} \xrightarrow{\;X_{i_{m_i-1}}\;} t_{i,j,m_i+1}$$

Furthermore, for every $i \in \{0, \ldots, m-1\}$, and for every $j \in \{0, \ldots, \kappa\}$, the LTS A has the following directed cycle $D_{i,j}$, where the event a occurs m_i times in a row followed by k_i, which closes the cycle:

$$D_{i,j} = d_{i,j,0} \xrightarrow{a} d_{i,j,1} \xrightarrow{a} \cdots \xrightarrow{a} d_{i,j,m_i-1} \xrightarrow{a} d_{i,j,m_i}$$

$$k_i$$
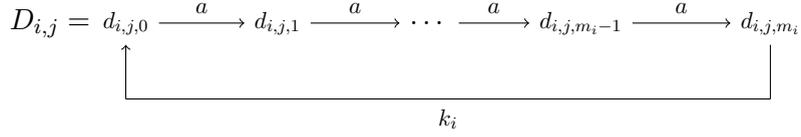
For all $i \in \{0, \ldots, n-1\}$, the LTS A also has following gadget F_i that synchronizes the event a with the variable event X_i:

$$F_i = f_{i,0} \underset{X_i}{\overset{a}{\rightleftarrows}} f_{i,1}$$
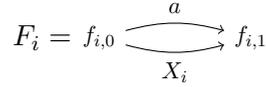
Finally, in order to connect the introduced gadgets, and to ensure reachability, the LTS A has the initial state ι, and the following edges:

- for all $i \in \{0, \ldots, m-1\}$, and all $j \in \{0, \ldots, \kappa\}$, the LTS A has the edges $\iota \xrightarrow{u_i^j} t_{i,j,0}$, and $\iota \xrightarrow{v_i^j} d_{i,j,0}$.

- For all $i \in \{0, \ldots, n-1\}$, the LTS A has the edge $\iota \xrightarrow{w_i} f_{i,0}$.

The result is the reachable LTS $A = (S, E, \delta, \iota)$.

Lemma 3.7. If there is an edge removal B of A that has the SSP and satisfies $|\mathfrak{K}| \leq \kappa$, then there is a hitting set with at most $\lambda = \kappa$ elements for $(\mathfrak{U}, M)$.

Proof:
Let B be an edge removal of A that has the SSP, and satisfies $|\mathfrak{K}| \leq \kappa$. Like in the proof of Lemma 3.7, without loss of generality, we may assume that no edge in $\mathfrak{K}$ arises from ι, nor from some $T_{i,j}$, nor from some $D_{i,j}$, and if one arises from some F_i the other edge from F_i is not in $\mathfrak{K}$, so that in particular B remains reachable. If some edge from ι has been removed, we may reintroduce it and choose its effect in order to keep the SSA we had before. Since there are more than κ copies $T_{i,j}$, at least one is completely in B; then we may complete the other copies in order to keep their SSAs. The same is true for $D_{i,j}$. Now, if both a and X_i are missing in some F_i, we may reintroduce a; indeed, when we separate $d_{i,j,0}$ from $d_{i,j,1}$, the effect of a is non-null, then $f_{i,0}$ will be separated from $f_{i,1}$ (it may be necessary to choose $pro(w_i)$ high enough in order to allow a in $f_{i,0}$).

In the following, we argue that

$$Z = \{X_i \mid i \in \{0, \ldots, n-1\} \text{ and } \{f_{i,0} \xrightarrow{X_i} f_{i,1}, f_{i,0} \xrightarrow{a} f_{i,1}\} \cap \mathfrak{K} \neq \emptyset\}$$

is a hitting set for $(\mathfrak{U}, M)$.

Let $i \in \{0, \ldots, m-1\}$ be arbitrary but fixed. From the hypothesis on B above, $T_{i,0}$, and $D_{i,0}$ are completely present in B. By the SSP of B, there is a region that solves the SSA $\alpha = (t_{i,0,0}, t_{i,0,m_i+1})$. Let $R = (sup, sig)$ be such a region. We argue that $Z \cap M_i \neq \emptyset$: Assume, for a contradiction,

that the opposite is true, that is, $Z \cap M_i = \emptyset$. Then, for all $\ell \in \{0, \ldots, m_i - 1\}$, the gadget F_ℓ is completely present in B. By Lemma 2.12, we have that $sup(d_{i,0,m_i}) = sup(d_{i,0,0}) + m_i \cdot \text{eff}(a)$, and $sup(d_{i,0,0}) = sup(d_{i,0,m_i}) + \text{eff}(k_i)$, which implies $\text{eff}(k_i) = -m_i \cdot \text{eff}(a)$. Moreover, since the gadget F_ℓ is present in B, we have that $\text{eff}(X_{i_\ell}) = \text{eff}(a)$ for all $\ell \in \{0, \ldots, m_i - 1\}$. Finally, again by Lemma 2.12, we have

$$\begin{aligned}
sup(t_{i,0,m_i+1}) &= sup(t_{i,0,0}) + \text{eff}(k_i) + \sum_{\ell=0}^{m_i-1} \text{eff}(X_{i_\ell}) \\
&= sup(t_{i,0,0}) - m_i \cdot \text{eff}(a) + m_i \cdot \text{eff}(a) \\
&= sup(t_{i,0,0}),
\end{aligned}$$

which contradicts that R solves α. Hence, there is an $\ell \in \{0, \ldots, m_i - 1\}$, such that $f_{i_\ell,0} \xrightarrow{X_{i_\ell}} f_{i_\ell,1} \in \mathfrak{K}$ or $f_{i_\ell,0} \xrightarrow{a} f_{i_\ell,1} \in \mathfrak{K}$, which implies $Z \cap M_i \neq \emptyset$. Since i was arbitrary, this is simultaneously true for all sets of M and Z is a hitting set for $(\mathfrak{U}, M)$, which completes the proof. $\qquad \square$

Conversely, we show:

Lemma 3.8. If there is a hitting set with at most λ elements for $(\mathfrak{U}, M)$, then there is an edge removal B of A that has the SSP, and satisfies $|\mathfrak{K}| \leq \lambda$.

Proof:
Let $Z = \{X_{\ell_0}, \ldots, X_{\ell_{\lambda-1}}\}$, where $\ell_0, \ldots, \ell_{\lambda-1} \in \{0, \ldots, n-1\}$, be a hitting set with λ elements for $(\mathfrak{U}, M)$ (which exists whenever there is a hitting set with at most λ elements). Moreover, let $B = (S, E, \delta', \iota)$ be the LTS that originates from A by removing, for all $j \in \{0, \ldots, \lambda-1\}$, the edge $f_{\ell_j,0} \xrightarrow{X_{\ell_j}} f_{\ell_j,1}$. One easily verifies that B is a well-defined edge removal of A that satisfies $|\mathfrak{K}| \leq \lambda = \kappa$, and that B is reachable. The definition of B implies that B and A differ (only) with respect to the gadgets $F_{\ell_0}, \ldots, F_{\ell_{\lambda-1}}$. However, by a little abuse of notation, and for the sake of readability, in the following, we will refer to $F_0, \ldots, F_{n-1}$ also as the gadgets of B, where we always keep in mind that the edge $f_{\ell_j,0} \xrightarrow{X_{\ell_j}} f_{\ell_j,1}$ is not present in F_{ℓ_j} for all $j \in \{0, \ldots, \lambda-1\}$.

First of all, it is easy to see, that if G and G' are two distinct gadgets of A, then, for all states $s \in S(G) \cup \{\iota\}$ and $s' \in S(G')$, the SSA (s, s') is solvable. Hence, it remains to argue that every SSA (s, s') of B is solvable if s and s' belong to the same gadget of B.

The following region $R = (sup, con, pro)$ solves all the corresponding SSAs in one blow: For all $i \in \{0, \ldots, m-1\}$, and all $j \in \{0, \ldots, n-1\}$, let $sup(d_{i,j,0}) = 0$, and $sup(t_{i,j,0}) = m_i$; for all $i \in \{0, \ldots, n-1\}$, let $sup(f_{i,0}) = 0$; for all $e \in \{a\} \cup (\mathfrak{U} \setminus Z)$, let $(con(a), pro(a)) = (0, 1)$; for all $i \in \{0, \ldots, m-1\}$, let $(con(k_i), pro(k_i)) = (m_i, 0)$; for all $j \in \{0, \ldots, \lambda-1\}$, let $(con(X_{\ell_j}), pro(X_{\ell_j})) = (0, \sum_{i=0}^{m-1} m_i)$. Obviously, R solves all SSA of the $D_{i,j}$'s, and the F_i's.

Let $i \in \{0, \ldots, m-1\}$, and $j \in \{0, \ldots, \kappa\}$ be arbitrary but fixed and let $h \in \{0, \ldots, m_i - 1\}$ be the smallest index such that X_{i_h} belongs to the hitting set Z, that is, $X_{i_h} \in Z$. Then the definition of R ensures that $sup(d_{i,j,1}) < \cdots < sup(d_{i,j,i_h+1}) < sup(d_{i,j,0}) < sup(d_{i,j,i_h+2}) < \cdots < sup(d_{i,j,m_i+1})$. By the arbitrariness of i, and j, this completes the proof. $\qquad \square$

4.   The complexity of Event Removal

In this section, we are looking for implementable event removals that remove only a bounded number of events. By the connection between implementations and separation properties stated by Lemma 2.16, the following problems arise:

EVENT REMOVAL FOR EMBEDDING

Input: An LTS $A = (S, E, \delta, \iota)$, a natural number κ.

Question: Does there exist an event removal B for A by $\mathfrak{E}$ that has the SSP and satisfies $|\mathfrak{E}| \leq \kappa$?

EVENT REMOVAL FOR LANGUAGE-SIMULATION

Input: An LTS $A = (S, E, \delta, \iota)$, a natural number κ.

Question: Does there exist an event removal B for A by $\mathfrak{E}$ that has the ESSP and satisfies $|\mathfrak{E}| \leq \kappa$?

EVENT REMOVAL FOR REALIZATION

Input: An LTS $A = (S, E, \delta, \iota)$, a natural number κ.

Question: Does there exist an event removal B for A by $\mathfrak{E}$ that has the ESSP and the SSP and satisfies $|\mathfrak{E}| \leq \kappa$?

The following theorem states the main result of this section:

Theorem 4.1. Event Removal for embedding, Event Removal for Language Simulation, and Event Removal for realization are NP-complete.
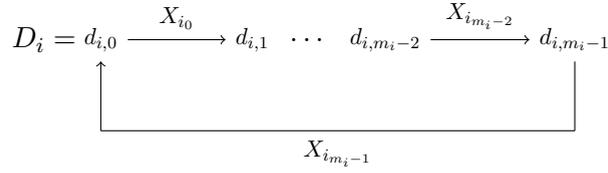
Obviously, the problems addressed by Theorem 4.1 belong to NP: If there is a fitting event removal B for A, then a Turing machine T can guess $\mathfrak{E}$ non-deterministically in time polynomial in the size of A, since $|\mathfrak{E}| \leq |E|$. After that, T can deterministically compute B and a witness for the property in question in polynomial time by Lemma 2.16, since the size of B is bounded by the size of A.

Hence, to complete the proof of the NP-completeness part of Theorem 4.1, it remains to prove the hardness-part. In order to do that, we provide a reduction of HITTING SET that reduces the input $(\mathfrak{U}, M, \lambda)$ to an instance (A, κ) with LTS A and natural number κ such that is a hitting set of size at most λ for $(\mathfrak{U}, M)$ if and only if A can be made implementable by removing at most κ events.

For a start, we first define $\kappa = \lambda$. Moreover, for every $i \in \{0, \ldots, m - 1\}$, the LTS A has the following directed path on which the elements of $M_i = \{X_{i_0}, \ldots, X_{m_i-1}\}$ occur as event, which are encompassed by the event k_i:

$$T_i = t_{i,0} \xrightarrow{\ k_i\ } t_{i,1} \xrightarrow{\ X_{i_0}\ } \cdots \xrightarrow{\ X_{i_{m_i-1}}\ } t_{i,m_i+1} \xrightarrow{\ k_i\ } t_{i,m_i+2}$$

Moreover, for every $i \in \{0, \ldots, m - 1\}$, the LTS A has the following directed cycle at which the elements of $M_i = \{X_{i_0}, \ldots, X_{m_i-1}\}$ occur consecutively as events:

$$D_i = d_{i,0} \xrightarrow{X_{i_0}} d_{i,1} \quad \cdots \quad d_{i,m_i-2} \xrightarrow{X_{i_{m_i-2}}} d_{i,m_i-1}$$

$$X_{i_{m_i-1}}$$

In order to connect the gadgets, and to ensure that every (reasonable) event removal of A is a reachable LTS, that is, every state is reachable from the initial state ι by a directed path, we add the following edges:

- for every $i \in \{0, \ldots, m-1\}$, and every $j \in \{0, \ldots, m_i + 2\}$, the LTS A has the edge $\iota \xrightarrow{u_i^j} t_{i,j}$;

- for every $i \in \{0, \ldots, m-1\}$, and every $j \in \{0, \ldots, m_i - 1\}$, the LTS A has the edge $\iota \xrightarrow{v_i^j} d_{i,j}$.

The result is the LTS $A = (S, E, \delta, \iota)$. In the following, we prove the functionality of A. (Recall that $\mathfrak{E}$ refers to the events removed from A.)

Lemma 4.2. If there is an event removal $B = (S, E', \delta', \iota)$ of A by $\mathfrak{E}$ that has the ESSP or the SSP, and satisfies $|\mathfrak{E}| \leq \kappa$, then there is a hitting set of size at most λ for $(\mathfrak{U}, M)$.

Proof:
In the following, we argue that the set $Z = \mathfrak{U} \cap \mathfrak{E}$ defines a hitting set with at most λ elements for $(\mathfrak{U}, M)$. In order to do that, we show that, for every $i \in \{0, \ldots, m-1\}$, there has to be at least one event $X \in M_i$ that does not occur in B, and thus is removed from A. Note first that we may assume that no u_i^j or v_i^j belongs to $\mathfrak{E}$, since otherwise it is always possible to reintroduce them in B, with adequate con/pro, while keeping the regions solving the ESSAs and SSAs. We may thus assume that B is reachable.

Assume that B has the ESSP (the SSP). Let $i \in \{0, \ldots, m-1\}$ be arbitrary but fixed. Since B has the ESSP (the SSP), there is a region that solves the ESSA $\alpha = (k_i, t_{i,1})$ (the SSA $\beta = (t_{i,1}, t_{i,m_i+1})$). Let $R = (sup, con, pro)$ be such a region. Assume, for a contradiction, that $M_i \cap \mathfrak{E} = \emptyset$. By Lemma 2.12, we have that $sup(d_{i,0}) = sup(d_{i,0}) + \sum_{j=0}^{m_i-1} \text{eff}(X_{i_j})$, which implies $\sum_{j=0}^{m_i-1} \text{eff}(X_{i_j}) = 0$. Since R solves α (solves β), and k_i occurs at t_{i,m_i+1}, we have that $con(k_i) > sup(t_{i,1})$, and $con(k_i) \leq sup(t_{i,m_i+1})$ (we have $sup(t_{i,1}) \neq sup(t_{i,m_i+1})$). Hence, we have that $sup(t_{i,m_i+1}) = sup(t_{i,1}) + \sum_{j=0}^{m_i-1} \text{eff}(X_{i_j}) \neq sup(t_{i,1})$, and thus $\sum_{j=0}^{m_i-1} \text{eff}(X_{i_j}) \neq 0$, which is a contradiction. Hence, we have that $M_i \cap \mathfrak{E} \neq \emptyset$. Since i was arbitrary, and $|Z| = |\mathfrak{U} \cap \mathfrak{E}| \leq |\mathfrak{E}| \leq \kappa = \lambda$, this implies that Z is a hitting set of size at most λ for $(\mathfrak{U}, M)$. The claim follows. $\square$

Lemma 4.3. If there is a hitting set of size at most λ for $(\mathfrak{U}, M)$, then there is an event removal $B = (S, E', \delta', \iota)$ of A by $\mathfrak{E}$ that has the ESSP and the SSP, and satisfies $|\mathfrak{E}| \leq \kappa$.

Proof:
Let $U = \{u_i^j \mid i \in \{0, \ldots, m-1\}, j \in \{0, \cdots m_i + 2\}\}$, and $V = \{v_i^j \mid i \in \{0, \ldots, m-1\}, j \in \{0, \cdots m_i - 1\}\}$. Let Z be a hitting set of size λ for $(\mathfrak{U}, M)$. Let B be the LTS that results from A by removing exactly the events of Z (and the corresponding edges), and nothing else. Obviously, B is a suitable event removal of A that satisfies $|\mathfrak{E}| \leq \lambda = \kappa$, and is reachable.

For a start, the following region $R_1 = (sup_1, con_1, pro_1)$ solves (s, e), and (ι, s) for all $s \in S \setminus \{\iota\}$, and all $e \in U \cup V$: $sup_1(\iota) = 1$, $con_1(e) = 1$ for each $e \in U \cup V$, all the other images are 0.

The following observation is crucial for our further argumentation: Since Z is a hitting set, and thus intersects with the event set of T_i, and D_i for all $i \in \{0, \ldots, m-1\}$, we observe that $B \setminus \{s \xrightarrow{e} s' \mid e \in U \cup V\}$ consists of simple directed paths, say $P = s_0^0 \xrightarrow{e_1^0} \ldots \xrightarrow{e_{n_0}^0} s_{n_0}^0, \ldots, P_y = s_0^y \xrightarrow{e_1^y} \ldots \xrightarrow{e_{n_y}^y} s_{n_y}^y$ for some $y \in \mathbb{N}$, on which every event of B occurs at most once.

The following region $R_2 = (sup_2, con_2, pro_2)$ solves all the remaining SSA of B: let $sup(\iota) = 0$, and, for all $i \in \{0, \ldots, y\}$, and all $j \in \{0, \ldots, n_i\}$, let $sup(s_j^i) = \sum_{\ell=0}^{i-1} n_\ell + i + j$, and, for all $e \in (\mathfrak{U} \setminus Z) \cup \{k_0, \ldots, k_{m-1}\}$, let $con(e) = 0$, and $pro(e) = 1$, and, for all $e \in U \cup V$, if $\iota \xrightarrow{e} s$ is the unique e-labeled edge in B, then $con(e) = 0$, and $pro(e) = sup(s)$. (That is, we successively increment the support values of the states of B by one, and the supports of the s_0^i's are chosen in order to get different values for all the simple paths.)

Let $a \in (\mathfrak{U} \setminus Z) \cup \{k_0, \ldots, k_{m-1}\}$ be arbitrary but fixed, and let $P_{i_0}, \ldots, P_{i_\ell}$ be the paths in which a occurs. The following region $R_4 = (sup_4, con_4, pro_4)$ solves (a, s) for all states $s \in S \setminus (\bigcup_{j=0}^{\ell} S(P_{i_j}))$: for all $s \in S$, if $s \in \bigcup_{j=0}^{\ell} S(P_{i_j})$, then $sup(s) = 1$, otherwise $sup(s) = 0$; for all $e \in (\mathfrak{U} \setminus Z) \cup \{k_0, \ldots, k_{m-1}\}$, if $e = a$, then $(con(e), pro(e)) = (1, 1)$, otherwise $(con(e), pro(e)) = (0, 0)$; for all $e \in U \cup V$, if $\iota \xrightarrow{e} s$ is the unique e-labeled edge in B, then $(con(e), pro(e)) = (0, sup(s))$.

It remains to argue that the events of $(\mathfrak{U} \setminus Z) \cup \{k_0, \ldots, k_{m-1}\}$ can be separated within the paths at which they occur. This can be seen as follows: Let $i \in \{0, \ldots, y\}$, and $j \in \{1, \ldots, n_i\}$ be arbitrary but fixed. We observe that there are mappings $R_5 = (sup_5, con_5, pro_5)$ such that $sup_5(s_{j-1}^i) = 1$, and $sup_5(s) = 0$ for all $s \in S(P_i) \setminus \{s_{j-1}^i\}$, and $(con_5(e_j^i), pro_5(e_j^i)) = (1, 0)$, and $(con_5(e_{j-1}^i), pro_5(e_{j-1}^i)) = (0, 1)$ if $j \geq 2$, and $(con_5(e), pro_5(e)) = (0, 0)$ for all $e \in E(P_i) \setminus \{e_{j-1}^i, e_j^i\}$. Moreover, since every path of $P_0, \ldots, P_{i-1}, P_{i+1}, \ldots, P_y$ contains every event of B at most once, and since the events of $U \cup V$ occur exactly once in B, it is easy to see that R_5 can be extended to a region of B, which, for all $s \in S(P_i) \setminus \{s_{j-1}^i\}$, solves (e_j^i, s) by $con(e_j^i) > s$. By the arbitrariness of i, this completes the proof. $\square$

5. The complexity of State Removal

In this section, we are interested in finding implementable state removals of A that remove only a restricted number of states. Again justified by Lemma 2.16, this task corresponds to the following decision problems:

STATE REMOVAL FOR EMBEDDING

Input: An LTS $A = (S, E, \delta, \iota)$, a natural number κ.

Question: Does there exist a state removal B for A by $\mathfrak{S}$ that has the SSP and satisfies $|\mathfrak{S}| \leq \kappa$?

STATE REMOVAL FOR LANGUAGE-SIMULATION

Input: An LTS $A = (S, E, \delta, \iota)$, a natural number κ.

Question: Does there exist a state removal B for A by $\mathfrak{S}$ that has the ESSP and satisfies $|\mathfrak{S}| \leq \kappa$?

STATE REMOVAL FOR REALIZATION

Input: An LTS $A = (S, E, \delta, \iota)$, a natural number κ.

Question: Does there exist a state removal B for A by $\mathfrak{S}$ that has the ESSP and the SSP and satisfies $|\mathfrak{S}| \leq \kappa$?

In the following Section 5.1, we show that state removal aiming at embedding or realization is NP-complete. After that, we show that this is also true if we aim at language simulation in Section 5.2.

5.1. State Removal aiming at language-simulation or realization

Theorem 5.1. State Removal for Language-Simulation and State Removal for Realization are NP-complete.

First of all, the problems are in NP: If there is a suitable state removal, then the set $\mathfrak{S}$ can be guessed non-deterministically in polynomial time; after that, B and witnesses for the corresponding separation properties can be computed deterministically in polynomial time by Lemma 2.16. The proof of the hardness-part is based again on a reduction of HITTING SET that transforms the input $(\mathfrak{U}, M, \lambda)$ into an instance (A, κ). In particular, we will reuse the reduction that has been used for the proof of Theorem 3.1. So let $\kappa = \lambda$, and let A be the LTS that has been defined in Section 3.1.

Lemma 5.2. If there is a state removal B of A by $\mathfrak{S}$ that has the ESSP and satisfies $|\mathfrak{S}| \leq \kappa$, then there is a hitting set of size at most λ for $(\mathfrak{U}, M)$.

Proof:
Let B be a state removal of A by $\mathfrak{S}$ that has the ESSP and satisfies $|\mathfrak{S}| \leq \kappa$. Without loss of generality, we may assume that $\mathfrak{S} \subseteq \{f_{j,1} | j \in \{0, \ldots, n-1\}\}$, so that in particular B is reachable. Indeed, for each $i \in \{0, \ldots, m-1\}$, there are $\kappa + 1$ copies of T_i, hence at least one of them is entirely in B, but then we may keep all of them without destroying the ESSP. For each $j \in \{0, \ldots, n-1\}$, if $f_{j,0}$ and $f_{j,1}$ are both in $\mathfrak{S}$, we may suppress $f_{j,0}$ without destroying ESSP; if $f_{j,0}$ is in $\mathfrak{S}$ but not $f_{j,1}$, we may replace the first one by the latter in $\mathfrak{S}$, still without destroying ESSP.

In the following, we will show that

$$Z = \{X_j \mid j \in \{0, \ldots, n-1\} \text{ and } f_{j,1} \in \mathfrak{S}\}$$

defines a suitable hitting set for $(\mathfrak{U}, M)$, that is, $Z \cap M_i \neq \emptyset$ for each $i \in \{0, \ldots, m-1\}$.

Assume, for a contradiction, that for some $i \in \{0, \ldots, m-1\}$ and for all $\ell \in \{0, \ldots, m_i - 1\}$, F_{i_ℓ} is entirely in B. Since B has the ESSP, there is a region that solves $\alpha = (X_{i_0}, t_{i,0,1})$. Let

$R = (sup, con, pro)$ be such a region. We argue that there is some $\ell \in \{0, \ldots, m_i - 1\}$ such that the X_{i_ℓ}-labeled edge $f_{i_\ell,0} \xrightarrow{X_{i_\ell}} f_{i_\ell,1}$ is not present in B.

Assume that this is not true and $f_{i_\ell,0} \xrightarrow{X_{i_\ell}} f_{i_\ell,1} \in B$ for all $\ell \in \{0, \ldots, m_i - 1\}$. By $t_{i,0,0} \xrightarrow{X_{i_0}}$, we have $con(X_{i_0}) \leq sup(t_{i,0,0})$; since R solves α, we have that $con(X_{i_0}) > sup(t_{i,0,1})$. This implies $con(X_{i_0}) > pro(X_{i_1})$, and thus $sup(f_{i_0,0}) > sup(f_{i_0,1})$. Moreover, by $t_{i,0,m_i} \xrightarrow{X_{i_0}}$, we have that $con(X_{i_0}) \leq sup(t_{i,0,m_i})$, which, by $con(X_{i_0}) > sup(t_{i,0,1})$, implies $sup(t_{i,0,1}) < sup(t_{i,0,m_i}) = sup(t_{i,0,1}) + \sum_{\ell=1}^{m_i-1} \mathrm{eff}(X_{i_\ell})$, and thus $\sum_{\ell=1}^{m_i-1} \mathrm{eff}(X_{i_\ell}) > 0$. In particular, there is an $\ell \in \{1, \ldots, m_i - 1\}$ (so that $i_\ell \neq i_0$) such that $con(X_{i_\ell}) < pro(X_{i_\ell})$, which implies $sup(f_{i_\ell,0}) < sup(f_{i_\ell,1})$.

Since the edges $f_{i_0,0} \xrightarrow{a_0} f_{i_0,1}$ and $f_{i_\ell,0} \xrightarrow{a_0} f_{i_\ell,1}$ are both present in B, by $sup(f_{i_0,0}) > sup(f_{i_0,1})$, we obtain $con(a_0) > pro(a_0)$ and, by $sup(f_{i_\ell,0}) < sup(f_{i_\ell,1})$, we get $con(a_0) < pro(a_0)$, which is a contradiction. Consequently, either $f_{i_0,1} \in \mathfrak{S}$ or $f_{i_\ell,1} \in \mathfrak{S}$ and thus $M_i \cap Z \neq \emptyset$. Since i was arbitrary, we have $M_i \cap Z \neq \emptyset$ for all $i \in \{0, \ldots, m - 1\}$ and Z defines a hitting set for $(\mathfrak{U}, M)$. The claim follows.   □

Conversely, the following lemma states that, if $(\mathfrak{U}, M, \lambda)$ allows a positive decision, then so does (A, κ):

Lemma 5.3. If there is a hitting set of size at most λ for $(\mathfrak{U}, M)$, then there is a state removal B of A by $\mathfrak{S}$ that has the ESSP and satisfies $|\mathfrak{S}| \leq \kappa$.

Proof:
Let $Z = \{X_{\ell_0}, \ldots, X_{\ell_{\lambda-1}}\}$, where $\ell_0, \ldots, \ell_{\lambda-1} \in \{0, \ldots, n-1\}$, be a hitting set for $(\mathfrak{U}, M)$, and let $B = (S', E, \delta', \iota)$ be the LTS that originates from A by removing, for all $j \in \{0, \ldots, \lambda - 1\}$, exactly the state $f_{\ell_j,1}$ (and the affected edges), and nothing else. One easily verifies that B is a state removal of A by $\mathfrak{S}$ that satisfies $|\mathfrak{S}| \leq \kappa = \lambda$. In particular, for every gadget of A, its initial state is present in B. Hence, every region defined in for the proof of Lemma 3.5 is also valid for B. Consequently, the claim follows similarly to the proof of this lemma.   □

5.2. State Removal aiming at embedding

Removing as few states of an LTS as possible in order to make it implementable is also hard if we are aiming at embedding:

Theorem 5.4. State Removal for Embedding is NP-complete.

Similar to Theorem 5.1, one argues that STATE REMOVAL FOR LANGUAGE SIMULATION is in NP. We prove the hardness part by reusing the reduction of Section 3.2, that is, we define $\kappa = \lambda$, and let A be defined as in Section 3.2.

Lemma 5.5. If there is a state removal B for A by $\mathfrak{S}$ that has the SSP and satisfies $|\mathfrak{S}| \leq \kappa$, then there is a hitting set with at most $\lambda = \kappa$ elements for $(\mathfrak{U}, M)$.

Proof:

Let B be a state removal of A that has the SSP and satisfies $|\mathfrak{S}| \leq \kappa$. Again, without loss of generality, we may assume that $\mathfrak{S} \subseteq \{f_{j,1}|j \in \{0, \ldots, n-1\}\}$, so that in particular B is reachable. The proof is similar to the one of Lemma 5.2.

In the following, we argue that

$$Z = \{X_i \mid i \in \{0, \ldots, n-1\} \text{ and } f_{i,1} \in \mathfrak{S}\}$$

defines a suitable hitting set for $(\mathfrak{U}, M)$.

Let $i \in \{0, \ldots, m-1\}$ be arbitrary but fixed. $T_{i,0}$, and $D_{i,0}$ are completely present in B. Since B has the SSP, there is a region $R = (sup, con, pro)$ that solves the SSA $(t_{i,0,0}, t_{i,0,m_i+1})$. Similar to the proof of Lemma 3.7, one argues that this implies that there is an $\ell \in \{1, \ldots, m_i - 1\}$, such that the edge $f_{i_\ell,0} \xrightarrow{X_{i_\ell}} f_{i_\ell,1}$ is not in B, i.e., $X_{i_\ell} \in Z$. By the arbitrariness of i, we get $Z \cap M_i \neq \emptyset$ for all $i \in \{0, \ldots, m-1\}$, and $|Z| \leq \kappa = \lambda$. $\square$

Conversely, we show:

Lemma 5.6. If there is a hitting set with at most λ elements for $(\mathfrak{U}, M)$, then there is a state removal B for A by $\mathfrak{S}$ that has the SSP and satisfies $|\mathfrak{S}| \leq \kappa$,

Proof:

Let $Z = \{X_{\ell_0}, \ldots, X_{\ell_{\lambda-1}}\}$, where $\ell_0, \ldots, \ell_{\lambda-1} \in \{0, \ldots, n-1\}$ be a hitting set for $(\mathfrak{U}, M)$, and let B be the LTS that originates from A by the removal of the state $f_{\ell_j,1}$ (and the corresponding edges) for all $j \in \{0, \ldots, \lambda - 1\}$, and nothing else. Obviously, B is A with at most $\kappa = \lambda$ state removals. Moreover, for every gadget of A, its initial states is present in B. Hence, the regions defined for the proof of Lemma 3.8 are also valid here. In particular, these regions witness the SSP of B. The claim follows. $\square$

6. Parameterized complexity

In the previous sections, we have shown that all the decision problems studied in this paper are not efficiently solvable (unless P=NP), i.e., there is no solution algorithm for these problems whose running time depends polynomially on the length of the input. However, the input length alone is a fairly crude measure of the complexity of a problem and may make it appear more difficult than it actually is.

In parameterized complexity, we examine the complexity of a problem not only in terms of the input length n, but also consider an additional (meaningful) parameter κ. The question now is whether the problem is solvable by an algorithm whose "non-polynomial behavior" is functionally bounded by the parameter κ. Such investigations are of great importance from both practical and theoretical points of view: In the case of a positive decision, we usually obtain algorithms that can be considered efficient if κ is small compared to the input length. In the case of a negative decision, we know that we have to look for other solution strategies. This prevents us from spending hours searching for algorithms

that provably cannot exist. In any case, from a theoretical standpoint, we get a better understanding of complexity at a more fine-grained level.

In the following, we introduce some necessary (basic) notions:

Definition 6.1. Let Σ be an alphabet. A *parameterized problem* is a language $L \subseteq \Sigma^* \times \mathbb{N}$. For an instance $(\omega, \kappa) \in \Sigma^* \times \mathbb{N}$, we call κ the parameter.

For all our decision problems, we consider κ as the parameter, that is, (A, κ) is an input of our original problems if and only if $((A, \kappa), \kappa)$ is an input of the parameterized versions.

Definition 6.2. A parameterized problem $L \subseteq \Sigma^* \times \mathbb{N}$ is called *fixed-parameter-tractable* (fpt for short), if there is a constant c, a polynomially computable function $f : \mathbb{N} \to \mathbb{N}$, and an algorithm that correctly decides, for every $(\omega, \kappa) \in \Sigma^* \times \mathbb{N}$, in time $f(\kappa)|\omega|^c$, whether $(\omega, \kappa) \in L$. The class of all fpt problems is called *FPT*.

Usually, before one tries to investigates whether a problem belongs to FPT, one proves that if κ is fixed, then the problem is solvable in polynomial time, that is, it belongs the complexity class *XP* (for *slice-wise polynomial*). It is easy to see, that if κ is fixed, then all the synthesis up to removal problems we have considered here are solvable in polynomial time; indeed, we only have to check for at most $\mathcal{O}((|S| + |E|)^\kappa)$ different modifications of A whether they are implementable, and each check can be done in polynomial time in the size of A [8, 15].

In classical complexity, we prove, usually by means of a polynomial time reduction, that a problem is NP-hard to show that it (most likely) cannot be solved efficiently. Similarly, we proceed in parameterized complexity theory to argue that a parameterized problem is not fpt. For every $i \in \mathbb{N}^+$, $\mathrm{W}[i]$ defines a class of parameterized problems and the following inclusions hold:

$$\mathrm{FPT} \subseteq \mathrm{W}[1] \subseteq \mathrm{W}[2] \subseteq \mathrm{W}[3] \cdots \subseteq \mathrm{XP}$$

The precise definition of the class $\mathrm{W}[i]$, $i \in \mathbb{N}^+$, is beyond the scope of this paper. We therefore refer the reader to [23]. It can be proved that $\mathrm{FPT} \neq \mathrm{XP}$. Moreover, according to the generally accepted working hypothesis, even $\mathrm{FPT} \neq \mathrm{W}[1]$ and $\mathrm{W}[i] \neq \mathrm{W}[i+1]$ holds for all $i \in \mathbb{N}^+$. Consequently, for all $i \geq 1$, we assume that any $\mathrm{W}[i]$-hard problem is not fpt. Thereby, we call a parameterized problem L_2 $\mathrm{W}[i]$-hard if every parameterized problem L_1 from $\mathrm{W}[i]$ can be reduced to L_2 by means of a *parameterized reduction*:

Definition 6.3. Let $L_1, L_2 \subseteq \Sigma^* \times \mathbb{N}$ be two parameterized problems. We say L_1 can be *reduced* to L_2 by a *parameterized reduction*, if there are a constant c, polynomially computable functions $g, h : \mathbb{N} \to \mathbb{N}$ and a function $f : \Sigma^* \times \mathbb{N} \to \Sigma^* \times \mathbb{N}$, such that, for all $(\omega, \kappa) \in \Sigma^* \times \mathbb{N}$, the following conditions are true:

1. $(\omega, \kappa) \in L_1$ if and only if $f(\omega, \kappa) = (\omega', \kappa') \in L_2$;

2. f is computable in time $g(\kappa) \cdot |\omega|^c$;

3. $\kappa' \leq h(\kappa)$.

Moreover, if the problem also belongs to $W[i]$, then we say it is W[i]-complete. To show that a parameterized problem L_2 is not fpt, a classical technique is to show for a known $W[i]$-hard problem L_1, where $i \in \mathbb{N}^+$, that it can be reduced to L_2 by a parameterized reduction. Because L_1 is $W[i]$-hard and hence (most likely) not fpt, it holds according to the following lemma that L_2 is also not fpt:

Lemma 6.4. ([23])
If L_1 and L_2 are two parameterized problems such that L_1 is reducible to L_2 by a parameterized reduction and $L_2 \in$ FPT, then $L_1 \in$ FPT.

By the following theorem, the problem HITTING SET is most likely not fpt:

Theorem 6.5. ([23])
Hitting Set parameterized by λ is $W[2]$-complete.

Obviously, the reductions from the previous sections are parameterized reductions, since $\kappa = \lambda$. Hence, by Lemma 6.4, and Theorem 6.5, we obtain the following result:

Theorem 6.6. For every of the edge-, event- and state removal modifications, and for every of the implementations embedding, language-simulation and realization, the parameterized problem of deciding, for a given LTS A, and a natural number κ, whether A can be made implementable by removing at most κ components belongs to the complexity class $W[P]$, and is $W[2]$-hard, where κ is considered as the parameter.

7.   Inapproximability

For each modification technique *edge removal*, *event removal* and *state removal*, and for each implementation *embedding*, *language-simulation* and *realization*, our goal is to modify A as little as possible according to the chosen modification kind, so that the result is an implementable LTS. Since a lower bound of this optimization problem is defined by the complexity of its decision version, there is little hope that an implementable modification can be found efficiently according to the results of the previous sections. Moreover, by Theorem 6.6, they are most likely not fpt for the natural parameter κ.

On the other hand, instead of an exact solution, an approximate solution could be satisfactory, at least if a fixed approximation ratio c can be guaranteed. This leads to the search for a so-called c-approximation algorithm, i.e., an algorithm that, for a given LTS A, yields in polynomial time an implementable modification B of A with at most $c \cdot \kappa$ removed components when the optimal solution is κ, where $c \geq 1$ is a fixed constant. In this section we will argue that no such algorithm exists for any of our problems.

For a start, we provide some necessary basic definitions.

Definition 7.1. (Optimization Problem [24])
An *NP optimization problem* Q is a 4-tuple $(I, sol, m, type)$ such that the following hold:

1. I is the set of *instances* of Q, and it is recognizable in polynomial time.

2. Given an instance $x \in I$, $sol(x)$ denotes the set of *feasible solutions* of x. The solutions of x are short, that is, there exists a polynomial p such that, for every $y \in sol(x)$, it holds $|y| \leq p(|x|)$ and, for every x and every y with $|y| \leq p(|x|)$, it is decidable in polynomial time whether $y \in sol(x)$.

3. For $x \in I$, and $y \in sol(x)$, $m(x, y)$ denotes the positive integer *measure* of y. The function m is computable in polynomial time.

4. $type \in \{min, max\}$.

The *class* NPO is the set of all NP optimization problems, and MIN NPO is the set of all minimization NPO problems. In this paper, we deal only with problems of MIN NPO. The goal of a MIN NPO problem $Q = (I, sol, m, min)$ with respect to an instance x is to find an *minimum solution*, that is, an $y^* \in sol(x)$, such that $m(x, y^*) = min\{m(x, y) \mid y \in sol(x)\}$. In the following, opt will denote the function that maps an instance x of Q to the measure of a minimum solution.

Definition 7.2. (Ratio)
Let $Q = (I, sol, m, min)$ be an Min-NPO problem. For every $x \in I$ and $y \in sol(x)$, by

$$R(x, y) = \frac{m(x, y)}{\mathtt{opt}(x)}$$

we define the *performance ratio R* of y with respect to x.

Definition 7.3. (c-Approximation Algorithm)
Given an MIN NPO problem Q, and an arbitrary constant $c \geq 1$, we say that an algorithm T is a c-approximation algorithm for Q, if it runs in polynomial time and, for every instance x of Q with $sol(x) \neq \emptyset$, outputs a solution $T(x) \in sol(x)$ such that $R(x, T(x)) \leq c$.

The following notion of *L-reducibility* can be used to show that, for any constant $c \geq 1$, a MIN NPO problem does not have a c-approximation algorithm:

Definition 7.4. (L-reduction [24])
Let $Q_1 = (I_{Q_1}, sol_{Q_1}, m_{Q_1}, min)$ and $Q_2 = (I_{Q_2}, sol_{Q_2}, m_{Q_2}, min)$ be two MIN NPO problems. Q_1 is said to be *L-reducible* to Q_2, in symbols $Q_1 \leq Q_2$, if there exist two functions g, f and two positive constants α, β such that:

1. For every $x \in I_{Q_1}$, $f(x) \in I_{Q_2}$ is computable in polynomial time.

2. For every $x \in I_{Q_1}$ and every $y \in sol_{Q_2}(f(x))$, $g(x, y) \in sol_{Q_1}(x)$ is computable in polynomial time.

3. For every $x \in I_{Q_1}$, $\mathtt{opt}_{Q_2}(f(x)) \leq \alpha \cdot \mathtt{opt}_{Q_1}(x)$

4. For every $x \in I_{Q_1}$ and every $y \in sol_{Q_2}(f(x))$,

$$|\mathtt{opt}_{Q_1}(x) - m_{Q_1}(x, g(x, y))| \leq \beta \cdot |\mathtt{opt}_{Q_2}(f(x)) - m_{Q_2}(f(x), y)|$$

The 4-tuble (f, g, α, β) is called an *L-reduction* from Q_1 to Q_2.

It is known from [24, p. 1765] that if $Q_1 \leq_L Q_2$ and there exists a c-approximation algorithm for Q_2, where $c \geq 1$ is a constant, then there exists a c'-algorithm for Q_1, where $c' = 1 + \alpha \cdot \beta \cdot (c - 1)$.

By Corollary 1.5 of [25], for every $c \geq 1$, there does not exist a c-approximation algorithm of the problem MINIMUM SET COVER, unless P = NP:

MINIMUM SET COVER

Input: A finite set $\mathfrak{U}$, and a system M of subsets of $\mathfrak{U}$.

Solution: $Y \subseteq M$ such that $\bigcup_{y \in Y} y = \mathfrak{U}$.

Measure: $|Y|$

Hence, by the well-known equivalence of MINIMUM SET COVER and MINIMUM HITTING SET, the same is true for the latter problem:

MINIMUM HITTING SET

Input: A finite set $\mathfrak{U}$, and a system M of subsets of $\mathfrak{U}$.

Solution: A hitting set $Z \subset \mathfrak{U}$ for $(\mathfrak{U}, M)$.

Measure: $|Z|$

In the following, we argue that, for every component X of EDGE, EVENT or STATE and for every implementation Y of EMBEDDING, LANGUAGE SIMULATION or REALIZATION, there is an L-reduction from MINIMUM HITTING SET to the following MIN NPO problem X-REMOVAL FOR Y, and thus prove that the latter does not allow a c-approximation algorithm for any constant $c \geq 1$:

MINIMUM X-REMOVAL FOR Y

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

Solution: A Y-implementable X-REMOVAL B of A.

Measure: Number of components removed (from A to obtain B according to X).

Notice that the measure is actually a function that depends on A and B since, for example, $m(A, B) = |E \setminus E'|$ if B is an event removal of A.

The proof of the following lemma is based on the observation that the polynomial time reductions of the previous sections can be extended to suitable L-reductions:

Lemma 7.5. For every component X of {EDGE, EVENT, STATE} and for every implementation Y of {EMBEDDING, LANGUAGE SIMULATION, REALIZATION}, there is an L-reduction from MINIMUM HITTING SET to X-REMOVAL FOR Y.

Proof:
For a start, we argue for the problem MINIMUM EDGE REMOVAL FOR REALIZATION; the argument will be similar for the other cases.

We obtain (f, g, α, β) as follows: f is the function that maps an input $(\mathfrak{U}, M)$ to the LTS A in accordance to the reduction of Section 3.1; for every solution B of A, $g(A, B) = \{X \in \mathfrak{U} \mid \exists s \xrightarrow{X} s' \in \mathfrak{K}\}$; finally, $\alpha = \beta = 1$. Obviously, f and g can be computed in polynomial time and, by

the proof arguments of Lemma 3.4, $g(A, B)$ defines actually a hitting set of $(\mathfrak{U}, M)$. We argue that the optimum values of both problems are equal:

Let Z be an optimum hitting set for $(\mathfrak{U}, M)$, that is, $|Z| = \mathtt{opt}((\mathfrak{U}, M))$, and let $B = (S, E', \delta', \iota)$ be the edge removal of A obtained in accordance to the proof arguments of Lemma 3.5 with the set of removed edges $\mathfrak{K}$. By the arguments of the same lemma, B is a feasible solution for A, which satisfies $|\mathfrak{K}| = |Z|$. Assume that there is a realizable solution $B' = (S, E'', \delta'', \iota)$ of A with set of removed edges $\mathfrak{K}'$ such that $|\mathfrak{K}'| < |\mathfrak{K}|$, i.e., B' removes strictly less edges than B. By the proof of Lemma 3.4, the set $Z' = \{X \in \mathfrak{U} \mid s \xrightarrow{X} s' \in \mathfrak{K}'\}$ defines a hitting set for $(\mathfrak{U}, M)$. This implies $|Z'| \leq |\mathfrak{K}'| < |\mathfrak{K}| = |Z|$ and thus contradicts the choice of Z. Hence, $\mathtt{opt}((\mathfrak{U}, M)) = \mathtt{opt}(A)$, which obviously implies $\mathtt{opt}(A) \leq \alpha \cdot \mathtt{opt}((\mathfrak{U}, M))$. In particular, (f, g, α, β) satisfies Condition 3 of Definition 7.4.

Moreover, if $B' = (S, E'', \delta'', \iota)$ is an arbitrary but fixed implementable solution of A with a set of removed edges $\mathfrak{K}'$, then $g(A, B') = Z' = \{X \in \mathfrak{U} \mid s \xrightarrow{X} s' \in \mathfrak{K}'\}$ and we get:

$$|\mathtt{opt}((\mathfrak{U}, M)) - |Z'|| = |\mathtt{opt}(A) - |Z'|| \leq |\mathtt{opt}(A) - |\mathfrak{K}'|| = \beta \cdot |\mathtt{opt}(A) - |\mathfrak{K}'||$$

since $0 \leq \mathtt{opt}((\mathfrak{U}, M)) = \mathtt{opt}(A) \leq |Z'| \leq |\mathfrak{K}'|$. Hence, (f, g, α, β) satisfies Condition 4 of Definition 7.4, and thus is a valid L-reduction.

The arguments for the other cases are similar: For every remaining combination of component X and implementation Y, one finds out that there is an L-reduction $(f, g, 1, 1)$, such that

- f is defined in accordance to the corresponding reduction that maps $(\mathfrak{U}, M)$ to LTS A, and

- for any Y-implementable X-removal B of A, g maps A, and B to the set $Z = g(A, B) \subseteq \mathfrak{U}$ of elements that is implied by the set $\mathfrak{R}$ of removed components, and

- $\mathtt{opt}((\mathfrak{U}, M)) = \mathtt{opt}(A)$, and $|\mathtt{opt}((\mathfrak{U}, M)) - |Z|| \leq |\mathtt{opt}(A) - |\mathfrak{R}||$.   □

Altogether, we get the main result of this section:

Theorem 7.6. For every component X of EDGE, EVENT, STATE, for every implementation Y of EMBEDDING, LANGUAGE SIMULATION, REALIZATION and for every constant $c \geq 1$, the problem X-REMOVAL FOR Y does not allow a c-approximation algorithm, unless P = NP.

8.   Conclusion

In this paper, we showed that converting an unimplementable LTS into an implementable one by removing as few of its states, events or edges as possible, is intractable from the point of view of classical complexity theory. This solves a problem that was left open in [19]. Notice that the reductions for edge- and event removal also work if these modifications are defined in a way that require all original states to be preserved. However, in general, they could then not always produce an implementable LTS, since there are unimplementable trees. Moreover, we also show that these problems are also

intractable from the point of view of parameterized complexity, and approximability as well. However, a complete characterization of the parameterized complexity of the problems is still open, and future work should address finding the exact upper bounds for them. Future work could also address other techniques of modifications that were suggested in the literature such as, for example, state- or event-refinement, state fusion or edge addition (like in [26]), or a mixture of these modifications.

Acknowledgements.

References

[1] Kordon F, Garavel H, Hillah LM, Hulin-Hubard F, Berthomieu B, Ciardo G, Colange M, Dal Zilio S, Amparore E, Beccuti M, Liebke T, Meijer J, Miner A, Rohr C, Srba J, Thierry-Mieg Y, van de Pol J, Wolf K. Complete Results for the 2017 Edition of the Model Checking Contest. http://mcc.lip6.fr/2017/results.php, 2017.

[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] de San Pedro J, Cortadella J. Mining structured Petri nets for the visualization of process behavior. In: Ossowski S (ed.), Proceedings of the 31st Annual ACM Symposium on Applied Computing, Pisa, Italy, April 4-8, 2016. ACM, 2016 pp. 839–846. doi:10.1145/2851613.2851645.

[4] Cortadella J, Kishinevsky M, Kondratyev A, Lavagno L, Yakovlev A. A region-based theory for state assignment in speed-independent circuits. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 1997. **16**(8):793–812. doi:10.1109/43.644602.

[5] Cortadella J, Kishinevsky M, Kondratyev A, Lavagno L, Yakovlev A. Logic Synthesis for Asynchronous Controllers and Interfaces. Springer Publishing Company, Incorporated, 2013. ISBN:3642627765.

[6] 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.

[7] Best E, Darondeau P. Petri Net Distributability. In: Ershov Memorial Conference, volume 7162 of *Lecture Notes in Computer Science*. Springer, 2011 pp. 1–18. doi:10.1007/978-3-642-29709-0_1.

[8] 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.

[9] Best E, Devillers RR. Synthesis and reengineering of persistent systems. *Acta Inf.*, 2015. **52**(1):35–60. doi:10.1007/s00236-014-0209-7.

[10] Best E, Schlachter U. Analysis of Petri Nets and Transition Systems. In: Proceedings 8th Interaction and Concurrency Experience, ICE 2015, Grenoble, France, 4-5th June 2015. 2015 pp. 53–67. doi:10.4204/EPTCS.189.6.

[11] Borde D, Dierkes S, Ferrari R, Gieseking M, Göbel V, Grunwald R, von der Linde B, Lückehe D, Schlachter U, Schierholz C, Schwammberger M, Spreckels V. Projektgruppe APT: Analyse von Petri-Netzen und Transitionssystemen. Final report of a project assignment, 2013. URL `https://github.com/CvO-Theory/apt/blob/master/doc/APT.pdf`.

[12] Caillaud B. Synet : A Synthesizer of Distributable Bounded Petri-Nets from Finite Automata, 1999. URL `https://www.irisa.fr/s4/tools/synet/`.

[13] Carmona J, Cortadella J, Kishinevsky M. Genet: A Tool for the Synthesis and Mining of Petri Nets. In: Ninth International Conference on Application of Concurrency to System Design, ACSD 2009, Augsburg, Germany, 1-3 July 2009. 2009 pp. 181–185. doi:10.1109/ACSD.2009.6.

[14] Cortadella J, Kishinevsky M, Kondratyev A, Lavagno L, Yakovlev A. Petrify: A Tool for Manipulating Concurrent Specifications and Synthesis of Asynchronous Controllers, 1997.

[15] 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.

[16] Verbeek HMW, Pretorius AJ, van der Aalst WMP, van Wijk JJ. Visualizing State Spaces with Petri Nets. Eindhoven University of Technology, Eindhoven, The Netherlands, 2007. URL `https://publications.rwth-aachen.de/record/715007`.

[17] Carmona J. The Label Splitting Problem. *Trans. Petri Nets Other Model. Concurr.*, 2012. **6**:1–23. doi: 10.1007/978-3-642-35179-2_1.

[18] Schlachter U, Wimmel H. Relabelling LTS for Petri Net Synthesis via Solving Separation Problems. *Trans. Petri Nets Other Model. Concurr.*, 2019. **14**:222–254. doi:10.1007/978-3-662-60651-3_9.

[19] Schlachter U, Wimmel H. Optimal Label Splitting for Embedding an LTS into an arbitrary Petri Net Reachability Graph is NP-complete. *CoRR*, 2020. **abs/2002.04841**. 2002.04841, URL `https://arxiv.org/abs/2002.04841`.

[20] Tredup R. Finding an Optimal Label-Splitting to Make a Transition System Petri Net Implementable: a Complete Complexity Characterization. In: ICTCS, volume 2756 of *CEUR Workshop Proceedings*. CEUR-WS.org, 2020 pp. 131–144.

[21] Tredup R. Edge, Event and State Removal: The Complexity of Some Basic Techniques that Make Transition Systems Petri Net Implementable. In: Application and Theory of Petri Nets and Concurrency - 42nd International Conference, PETRI NETS 2021, Virtual Event, June 23-25, 2021, Proceedings. 2021 pp. 253–273. doi:10.1007/978-3-030-76983-3_13.

[22] Karp RM. Reducibility Among Combinatorial Problems. In: Miller RE, Thatcher JW (eds.), Proceedings of a symposium on the Complexity of Computer Computations, held March 20-22, 1972, at the IBM Thomas J. Watson Research Center, Yorktown Heights, New York, USA, The IBM Research Symposia Series. Plenum Press, New York, 1972 pp. 85–103. doi:10.1007/978-1-4684-2001-2_9.

[23] Cygan M, Fomin FV, Kowalik L, Lokshtanov D, Marx D, Pilipczuk M, Pilipczuk M, Saurabh S. Parameterized Algorithms. Springer, 2015. doi:10.1007/978-3-319-21275-3.

[24] Crescenzi P, Kann V, Silvestri R, Trevisan L. Structure in Approximation Classes. *SIAM J. Comput.*, 1999. **28**(5):1759–1782. doi:10.1137/S0097539796304220.

[25] Dinur I, Steurer D. Analytical Approach to Parallel Repetition. *CoRR*, 2013. **abs/1305.1979**. 1305.1979, URL `http://arxiv.org/abs/1305.1979`.

[26] Schlachter U. Petri net synthesis and modal specifications. Ph.D. thesis, University of Oldenburg, Germany, 2018.