Complexity of Distributed Petri Net Synthesis

Raymond Devillers[1] and Ronny Tredup[2] (✉)

[1] Département d'Informatique - Université Libre de Bruxelles, Boulevard du Triomphe, B1050 Brussels, Belgium (rdevil@ulb.ac.be)
[2] Institut Für Informatik - Universität Rostock, Albert-Einstein-Straße 22, D18059 Rostock, Germany (ronny.tredup@uni-rostock.de)

Abstract. *Distributed Petri Net Synthesis* corresponds to the task to decide, for a transition system A (with event set E) and a natural number κ, whether there exists a surjective location map $\lambda : E \to \{1, \ldots, \kappa\}$ and a Petri net N (with transition set E) such that, if two transitions $e, e' \in E$ share a common pre-place, then they have the same location ($\lambda(e) = \lambda(e')$), whose reachability graph is isomorphic to A (in which case such a solution should be produced as well). In this paper, we show that this problem is NP-complete.

1 Introduction

Labeled transition systems, TS for short, are a widely used tool for describing the potential sequential behaviors of discrete-state event-driven systems such as, for example, Petri nets.

Petri net synthesis consists in deciding, for a given transition system A, whether there exists a Petri net N whose reachability graph A_N is isomorphic to A, i. e., whether the TS indeed describes the behavior of a Petri net. In case of a positive decision, a possible solution N should be constructed as well. In this case, many solutions may usually be exhibited, sometimes with very different structures, and we may try to find solutions in a structural subclass of Petri nets with a particular interest.

Petri net synthesis has numerous practical applications, for example, in the field of process discovery to reconstruct a model from its execution traces [1], in supervisory control for discrete event systems [8], and in the design and synthesis of speed-independent circuits [5].

One of the most important applications of Petri net synthesis is the extraction of concurrency and distributability data from the sequential behavior given for instance by a TS [3]: Although TS are used in particular to describe the behavior of concurrent systems like Petri nets [10], they reflect concurrency only implicitly by the non-deterministic interleaving of sequential sequences of events.

In a Petri net whose reachability graph is isomorphic to a TS, the events of the TS correspond to the transitions of the Petri net, and the pre-places of a transition (an event in the TS) model the resources necessary for the firing of the transition (the occurrence of the event in the TS). Accordingly, the pre-places of a transition control the executability of the latter and, following Starke [11], transitions may be considered to be potentially concurrent if the intersection of their presets is empty, i. e., if they do not require the same resources. The concurrency of events (of the TS) thus becomes explicitly visible through the empty intersection of the presets of their corresponding transitions (in a synthesized net).

The question whether a TS having the event set E allows a distributed implementation not only asks about the concurrency of events, but goes a step further and asks whether concurrent events can actually be implemented at different physical locations. More exactly, for a set $\mathcal{L}$ of locations, one wonders if there is a surjective mapping $\lambda : E \to \mathcal{L}$ that assigns a (physical) location to each event $e \in E$ of the TS such that no two events sent to different locations share an input place.

In particular, it should be emphasized that concurrency and distributability are not equivalent properties: As elaborated in [4], transitions can be concurrent, but still not distributable. This phenomenon occurs, for example, in the context of the problem known as *confusion*: Although two transitions, say a and b, do not share any pre-places (do not require the same resources), there is a third transition, say c, that requires both resources from a and resources from b, so that a, b, and c must always be assigned to the same physical location.

The distributability of a transition system can thus be reduced to the distributability of Petri nets [3]. Note however that a TS may have various kinds of synthesized nets, some of which may be more or less highly distributed, while other ones are not at all. If λ is a distribution over E (in the sense just described), we may then say that a TS is λ-distributable if it has a λ-distributable Petri net synthesis. It is known that the question whether, for a TS A with event set E and a location map $\lambda : E \to \mathcal{L}$, a corresponding λ-distributable Petri net exists can be decided in polynomial time if λ is fixed in advance [3]. However, it is not clear a priori how a set of locations, and a location map can be chosen such that they describe an optimal distributed implementation of A, i. e., such that they imply a solution of the following optimization problem:

Given a TS A with event set E, find the maximum number κ of locations, and a (surjective) location map $\lambda : E \to \{1, \dots, \kappa\}$ that allow a distributed implementation of A, i. e., such that there exists a λ-distributable Petri net N whose reachability graph is isomorphic to A.

Since location maps are surjective, $\kappa \leq |E|$, and sending all transitions to a single location is always valid so that $\kappa \geq 1$. Moreover, if we have a distribution over κ locations, by grouping some of them we can get location maps to any subset of them. Hence, we can reduce by dichotomy the previous problem to the following:

Given a TS A with event set E, and a natural number κ between 1 and $|E|$, decide whether there exists a (surjective) location map $\lambda : E \to \{1, \dots, \kappa\}$ allowing a λ-distributable Petri net N whose reachability graph is isomorphic to A.

In this paper, we shall show that the latter problem is NP-complete, hence also the optimal one (so that these problems most probably cannot be solved efficiently in all generality).

The remainder of this paper is organized as follows: The following Section 2 introduces the definitions, and some basic results used throughout the paper, and provides them with examples. After that, Section 3 analyzes the distribution problem and Section 4 provides the announced NP-completeness result. Finally, Section 5 briefly closes the paper. The appendix contains some figures to help the reader understand some of the proofs.

2 Preliminaries

In this paper, we consider only finite objects, i. e., sets of events, states, places, etc. are always assumed to be finite.

Definition 1 (Transition System). *A (deterministic, labeled)* transition system*, TS for short, $A = (S, E, \delta, \iota)$ consists of two disjoint sets of* states S *and* events E *and 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 $\xrightarrow{\neg e}$ 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), denoted by $s_0 \xrightarrow{e_1} s_1 \xrightarrow{e_2} \ldots \xrightarrow{e_n} s_n$.*

We assume that A is reachable*: there is a path from ι to s for every state $s \in S \setminus \{\iota\}$.*

Two transition systems $A_1 = (S_1, E, \delta_1, \iota_1)$ and $A_2 = (S_2, E, \delta_2, \iota_2)$ on the event set E are said isomorphic (denoted $A_1 \cong A_2$) if there is a bijection $\beta : S_1 \to S_2$ such that $\beta(\iota_1) = \iota_2$ and $\delta_1(s_1) = s_1'$ iff $\delta_2(\beta(s_1)) = \delta_2(s_1')$ for any $s_1, s_1' \in S_1$ (also meaning that $\delta_1(s_1)$ is undefined iff so is $\delta_2(\beta(s_1))$).

$$\begin{tikzcd} s_3 \arrow[r, "a"] & s_2 \\ \mathbf{s_0} \arrow[u, "b"] \arrow[r, "a"'] & s_1 \arrow[u, "b"'] \end{tikzcd} \qquad \begin{tikzcd} 101 \arrow[r, "a"] & 002 \\ \mathbf{110} \arrow[u, "b"] \arrow[r, "a"'] & 011 \arrow[u, "b"'] \end{tikzcd}$$

A $\qquad$ A'

Fig. 1: Two isomorphic TS A and A'; the initial states are indicated in bold.

Definition 2 (Petri net). *A (weighted)* Petri net $N = (P, E, f, \mathfrak{m}_0)$ *consists of finite and disjoint sets of* places P *and* transitions E, *a (total)* flow $f : ((P \times E) \cup (E \times P)) \to \mathbb{N}$ *and an* initial marking $\mathfrak{m}_0 : P \to \mathbb{N}$ *(more generally, a marking is any function $P \to \mathbb{N}$, interpreted as giving a number of* tokens *present in each place).*
The preset of a transition is the set ${}^\bullet e = \{p \in P | f(p, e) > 0\}$ of its pre-places. The same may be defined for places, as well as postsets.
A transition $e \in E$ can fire *or* occur *in a marking $\mathfrak{m} : P \to \mathbb{N}$, denoted by $\mathfrak{m}\xrightarrow{e}$, if $\mathfrak{m}(p) \geq f(p, e)$ for all places $p \in P$. The firing of e in marking $\mathfrak{m}$ leads to the marking $\mathfrak{m}'(p) = \mathfrak{m}(p) - f(p, e) + f(e, p)$ for all $p \in P$, denoted by $\mathfrak{m}\xrightarrow{e}\mathfrak{m}'$. This notation extends to sequences $w \in E^*$ and the* reachability set $RS(N) = \{\mathfrak{m} \mid \exists w \in E^* : \mathfrak{m}_0\xrightarrow{w}\mathfrak{m}\}$ *contains all of N's reachable markings. The* reachability graph *of N is the TS $A_N = (RS(N), E, \delta, \mathfrak{m}_0)$, where, for every reachable marking $\mathfrak{m}$ of N and transition $e \in E$ with $\mathfrak{m}\xrightarrow{e}\mathfrak{m}'$, the transition function δ of A_N is defined by $\delta(\mathfrak{m}, e) = \mathfrak{m}'$ ($\delta(\mathfrak{m}, e)$ is undefined if e cannot fire in $\mathfrak{m}$).*

Many subclasses of Petri nets may be defined, and we shall consider some examples in the next section.

Definition 3 (Petri net synthesis). *Petri net synthesis consists in deciding, for a given transition system A, whether there exists a Petri net N whose reachability graph A_N is isomorphic to A, i. e., whether the TS indeed describes the behavior of a Petri net.*
In the positive case, one usually wants to also build such a net, called a solution of the synthesis problem. In the negative case, it may be useful to exhibit one or more reasons of the failure.
It is also possible to restrict the target to some specific subclass of nets.

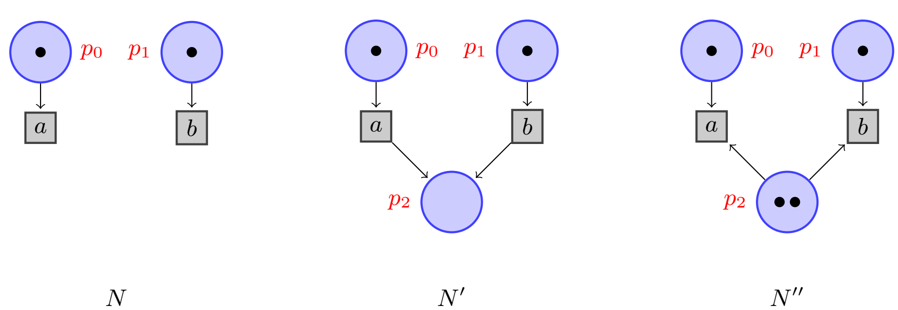

Fig. 2: Three different solutions of the TS A and A' in Figure 1. A' is the reachability graph of N'.

Classical synthesis procedures are linked to the notion of regions and to the solution of separation properties.

Definition 4 (Region). *Let $A = (S, E, \delta, \iota)$ be a TS. A* region $R = (sup, con, pro)$ *of A consists of three mappings* support $sup : S \to \mathbb{N}$*, as well as* consume *and* produce $con, pro : E \to \mathbb{N}$*, such that if $s \xrightarrow{e} s'$ is an edge of A, then $con(e) \le sup(s)$ and $sup(s') = sup(s) - con(e) + pro(e)$.*

A region may be seen as a place of a Petri net with transition set E, with *sup* giving the marking of the place at each reachable state as specified by A, $con(e)$ giving the number of tokens needed (and thus consumed when firing) by e in that place, and $pro(e)$ giving the number of tokens produced by e in that place when firing.

The *state separation property* ensures that different states may be differentiated by a region, i.e., be associated with different markings:

Definition 5 (State Separation Property). *Two distinct states $s, s' \in S$ define the* state separation atom*, SSA for short, (s, s') of A. A region $R = (sup, con, pro)$ solves (s, s') if $sup(s) \neq sup(s')$. A state $s \in S$ is called* solvable *if, for every $s' \in S \setminus \{s\}$, there is a region that solves the SSA (s, s'). If every state of A is solvable, then A has the* state separation property*, SSP for short.*

The *event state separation property* ensures that if an event e does not occur at a state s in A, that is $s \xrightarrow{\neg e}$, then the transition e cannot fire in the marking associated to s in some region:

Definition 6 (Event State Separation Property). *An event $e \in E$, and a state $s \in S$ of A such that $s \xrightarrow{\neg e}$ define the* event state separation atom*, ESSA for short, (e, s) of A. A region $R = (sup, con, pro)$ solves (e, s) if $con(e) > sup(s)$. An event $e \in E$ is called* solvable *if, for every state $s \in S$ such that $s \xrightarrow{\neg e}$, there is a region of A that solves the ESSA (e, s). If all events of A are solvable, then A has the* event state separation property*, ESSP for short.*

Definition 7 (Admissible Set). *Let $A = (S, E, \delta, \iota)$ be a TS. A set $\mathcal{R}$ of regions of A is called an* admissible set *if it* witnesses *the SSP and the ESSP of A, i. e., for every SSA, and for every ESSA of A, there is a region in $\mathcal{R}$ that solves it.*
If $\mathcal{R}$ is an (admissible) set of regions of A, $N_A^{\mathcal{R}}$ is the Petri net where E is the set of transitions, $\mathcal{R}$ is the set of places and, for each place $R = (sup, con, pro) \in \mathcal{R}$, the initial marking is $sup(\iota)$ and, for each transition $e \in E$, $f(R, e) = con(e)$ and $f(e, R) = pro(e)$.

A classical result about Petri net synthesis is then:

Theorem 1 ([6]). *A labeled transition system A has a weighted Petri net solution iff it has an admissible set $\mathcal{R}$ of regions. A possible solution is then $N_A^{\mathcal{R}}$.*

3 Distributability

The idea here is to bind the events of a transition system or a Petri net to certain (physical) locations.

Definition 8 (Location Map). *Let E be a set, and $\mathcal{L}$ a set of* locations*. A* location map *(over E and $\mathcal{L}$) is a surjective mapping $\lambda : E \to \mathcal{L}$.*

In the case of a Petri net, the intent is to separate the pre-sets of transitions sent to different locations:

Definition 9 (Distributable Petri net). *Let $N = (P, E, f, \mathfrak{m}_0)$ be a Petri net, $\mathcal{L}$ a set of locations, and $\lambda : E \to \mathcal{L}$ a location map. N will be called λ-*distributable *if the following condition is satisfied: for all transitions $e, e' \in E$ and every place $p \in P$, if $f(p, e) > 0$ and $f(p, e') > 0$, then $\lambda(e) = \lambda(e')$.*
*Let $\kappa \in \mathbb{N}$; N will be called κ-*distributable *(with $1 \leq \kappa \leq |E|$) if it is λ-distributable for some location map λ such that $|\mathcal{L}| = \kappa$.*

The last definition results from the observation that the exact identity of the locations is not important: what really matters is the partition of the transition set defined by λ, i.e., $\{\lambda^{-1}(e) | e \in \mathcal{L}\}$. Hence we may always choose $\mathcal{L} = [1...|\mathcal{L}|]$. Moreover, if π is a permutation of $\mathcal{L}$, we may use equivalently $\pi \circ \lambda$ instead of λ.

For instance, the nets N and N' in Figure 2 are λ-distributable with $\lambda(a) = 1$ and $\lambda(b) = 2$, hence also 2-distributable. On the contrary, N'' is only 1-distributable.

We may then consider the synthesis problems where the target is the class of λ-distributable Petri nets, for some location map λ, or the class of κ-distributable Petri nets, for some $\kappa \in [1...|E|]$.

Definition 10 (Localized Region). *Let $A = (S, E, \delta, \iota)$ be a TS, $\mathcal{L}$ a set of locations and $\lambda : E \to \mathcal{L}$ a location map. A λ-*localized region *is a region $R = (sup, con, pro)$ of A such that, if $con(e) > 0$, and $con(e') > 0$, then $\lambda(e) = \lambda(e')$.*

In other words, if $\lambda(e) \neq \lambda(e')$, then either $con(e) = 0$ or $con(e') = 0$ (or both).

Definition 11 (Localized admissible Set). *Let $A = (S, E, \delta, \iota)$ be a TS, $\mathcal{L}$ a set of locations, and $\lambda : E \to \mathcal{L}$ a location map. An admissible set $\mathcal{R}$ of regions of A will be said λ-*localized *if all its members are λ-localized. It will be said κ-localizable (for some $\kappa \in [1...|E|]$) if it is λ-localized for some location map λ with $|\lambda(E)| = \kappa$.*

The following result extends Theorem 1 to the localized context. It states that the question whether there is a λ-distributable (or a κ-distributable) Petri net whose reachability graph is isomorphic to A is equivalent to the question whether there is a λ-localized (or a κ-localizable) admissible set of regions of A:

Theorem 2 ([2]). *Let $A = (S, E, \delta, \iota)$ be a TS, $\mathcal{L}$ a set of locations, $\lambda : E \to \mathcal{L}$ a location map, and $\kappa \in [1...|E|]$ a degree of distribution. A has a λ-distributed (or a κ-distributed) Petri net solution iff it has an admissible λ-localized (or κ-localizable) set $\mathcal{R}$ of regions. A possible solution is then $N_A^{\mathcal{R}}$.*

If a TS A allows a λ-distributed (hence also a κ-distributed) Petri net solution N, it is possible to extend the location map to the places: if $p \in {}^\bullet e$, we may coherently state $\lambda(p) = \lambda(e)$. If a place p has an empty post-set, we may arbitrarily associate it to any location, for instance to $\lambda(e)$ if $e \in {}^\bullet p$ (if any), but here the location may rely on the particular choice of e. If we add the initial marking and the arcs between the connected places and transitions in each location, we shall then get $|\mathcal{L}|$ subnets $N_1, \ldots, N_{|\mathcal{L}|}$.

If these subnets are well separated, $N = \bigoplus_{i=1}^{|\mathcal{L}|} N_i$ is the disjoint sum of its various localized components, in the sense of [7], and then its reachability graph is isomorphic to the disjoint product of the reachability graphs of those components: if $A_i = RG(A_i)$ for each $i \in [1...|\mathcal{L}|]$, $A \cong RG(N) \cong \bigotimes_{i=1}^{|\mathcal{L}|} RG(N_i)$. This is the case for example for the net N in Figure 2, but not for N' while both nets are 2-distributable and solutions of the same TS A.

In general, however, each component N_i still has to send tokens to places belonging to other components, and the relationship on the reachability graphs is not so obvious. In [3], the authors show how to get around the difficulty. Albeit we shall not need it in the following, we sketch here their procedure. When components have to exchange tokens, it is not possible to read it in the corresponding transition systems, since the latter are considered up to isomorphisms, so that the markings disappear. Instead, the idea is to add special transitions materializing the sending or reception of a token to or from another component, but these extra transitions will be considered as invisible from outside. This leads to reachability graphs and transition systems with invisible events, but it is possible to define an equivalence, called branching bisimulation, which generalizes the isomorphism between transition systems without invisible events, and to combine disjoint transition systems with invisible events in such a way that the combination of the reachability graphs of the (extended) components N_i is branching bisimilar to the original TS A. For instance, for net N' in Figure 2, this leads to the components, reachability graphs and combination illustrated by Figures 3, 4, and 5.

4 Complexity Analysis

In [3], it is shown that the question whether, for a TS A with event set E and a location map $\lambda : E \to \mathcal{L}$, a corresponding λ-distributable Petri net exists can be decided in polynomial time. But this is only proved when λ is fixed in advance, and it is not clear a priori if this remains true if λ is left unknown, as in the decision problems mentioned in the Introduction, which may now be formalized as follows:

Optimal Distributability *Input:* A TS A with event set E and an integer $\kappa \in [1...|E|]$. *Question:* Is κ the maximal value such that A has a κ-distributable solution?

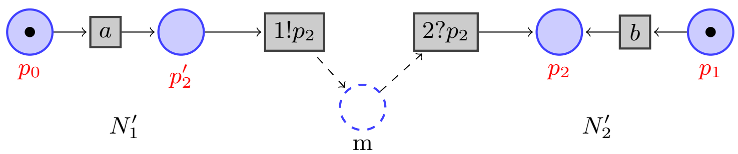

Fig. 3: N_1' and N_2' are the local components of N' associated to locations 1 and 2, respectively; $1!p_2$ is the invisible transition that sends asynchronously a token to p_2 in N_2' from location 1, and $2?p_2$ is the invisible transition that receives asynchronously in N_2' a token for p_2. We can think of this sending/receiving of tokens as follows: There is an additional (message) place m; the firing of $1!p_2$ produces a token on m (message "N_1' sends a token for p_2"); the firing of $2?p_2$ consumes a token from m, and produces a token on p_2 ("message received").

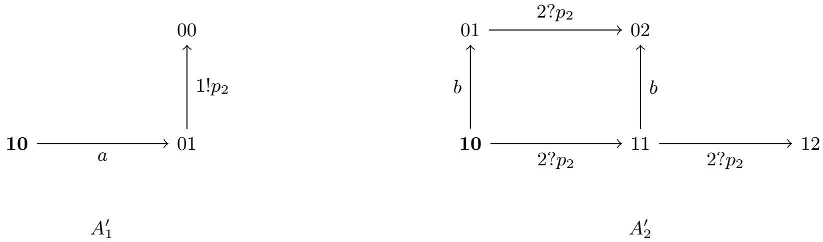

Fig. 4: A_1' and A_2' are the corresponding reachability graph (bounded by the maximal marking 2 of p_2 in N').

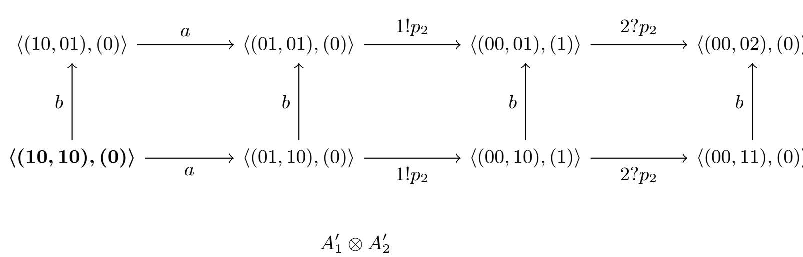

Fig. 5: Finally, $A_1' \otimes A_2'$ is the combination of A_1' and A_2' that is branching bisimilar to A. A state $(\langle(s, s')\rangle, (i))$ of $A_1' \otimes A_2'$ corresponds to the pair (s, s') of states of A_1' and A_2', respectively, and i is the number of messages sent but not yet received.

κ-Distributability
Input: A TS A with event set E and an integer $\kappa \in [1 \ldots |E|]$.
Question: Is there a κ-distributable solution?

We shall show in this section that the second problem (hence also the first one) is unfortunately NP-complete. First of all, we argue for the membership in NP: On the one hand, if, for a given TS A, and a natural number κ, there is a location map λ that allows a corresponding λ-admissible set, then a non-deterministic Turing machine can compute λ in polynomial time (by simply guessing $\lambda(e) \in [1 \ldots |E|]$ for all $e \in E$). On the other hand, as mentioned above, it is known that once λ is fixed, one can compute in polynomial time a corresponding λ-distributable admissible set $\mathcal{R}$ if it exists (and reject the input otherwise) [2]. Hence, κ-distributability is in NP.

Cubic Monotone 1 in 3 3Sat (CM1in33Sat)
Input: A pair $(\mathfrak{U}, M)$ that consists of a set $\mathfrak{U}$ of boolean variables and a set of 3-clauses $M = \{M_0, \ldots, M_{m-1}\}$ such that $M_i = \{X_{i_0}, X_{i_1}, X_{i_2}\} \subseteq \mathfrak{U}$ and $i_0 < i_1 < i_2$ for all $\in \{0, \ldots, m-1\}$. Every variable of $\mathfrak{U}$ occurs in exactly three clauses of M
Question: Does there exist a one-in-three model of $(\mathfrak{U}, M)$, i. e., a subset $\mathfrak{S} \subseteq \mathfrak{U}$ such that $|\mathfrak{S} \cap M_i| = 1$ for all $i \in \{0, \ldots, m-1\}$?

Theorem 3 ([9]). Cubic Monotone 1 in 3 3Sat *is* NP*-complete.*

Example 1. The instance $(\mathfrak{U}, M)$, where $\mathfrak{U} = \{X_0, X_1, X_2, X_3, X_4, X_5\}$, and $M = \{M_0, \ldots, M_5\}$ such that $M_0 = \{X_0, X_1, X_2\}$, $M_1 = \{X_0, X_1, X_3\}$, $M_2 = \{X_0, X_1, X_5\}$, $M_3 = \{X_2, X_3, X_4\}$, $M_4 = \{X_2, X_4, X_5\}$, and $M_5 = \{X_3, X_4, X_5\}$, allows a positive decision: $\mathfrak{S} = \{X_0, X_4\}$ defines a one-in-three model for $(\mathfrak{U}, M)$.

In the following, until explicitly stated otherwise, let $(\mathfrak{U}, M)$ be an arbitrary but fixed instance of CM1in33Sat such that $\mathfrak{U} = \{X_0, \ldots, X_{m-1}\}$, and $M = \{M_, \ldots, M_{m-1}\}$, where $M_i = \{X_{i_0}, X_{i_1}, X_{i_2}\} \subseteq \mathfrak{U}$, and $i_0 < i_1 < i_2$ for all $i \in \{0, \ldots, m-1\}$. Note that $|\mathfrak{U}| = |M|$ holds by the definition of a valid input.

Lemma 1. *If* $\mathfrak{S} \subseteq \mathfrak{U}$*, then* $\mathfrak{S}$ *is a one-in-three model of* $(\mathfrak{U}, M)$ *if and only if* $\mathfrak{S} \cap M_i \neq \emptyset$ *for all* $i \in \{0, \ldots, m-1\}$*, and* $m = 3 \cdot |\mathfrak{S}|$.

Proof. Every variable of $\mathfrak{U}$ occurs in exactly three distinct clauses. Hence, every set $\mathfrak{S} \subseteq \mathfrak{U}$ intersects with $3|\mathfrak{S}|$ (distinct) clauses $M_{i_0}, \ldots, M_{i_{3|\mathfrak{S}|-1}} \in M$ if and only if $|\mathfrak{S} \cap M_{i_j}| = 1$ for all $j \in \{0, \ldots, 3|\mathfrak{S}| - 1\}$. $\square$

We shall polynomially reduce $(\mathfrak{U}, M)$ to a TS $A = (S, E, \delta, \iota)$ and a number κ such that there is location map $\lambda : E \to \{1, \ldots, \kappa\}$, and a λ-localizable admissible set of A if and only if $(\mathfrak{U}, M)$ has a one-in-three model.

For a start, let $\kappa = \frac{2m}{3} + 3$, and $\mathcal{L} = \{1, \ldots, \frac{2m}{3} + 3\}$. (By Lemma 1, if $m \not\equiv 0 \bmod 3$, then $(\mathfrak{U}, M)$ has no one-in-three model.) We proceed with the construction of A, being the composition of several gadgets that are finally connected by some uniquely labeled edges. First of all, the TS A has the following gadget H that will allow to consider the ESSA $\alpha = (k, h_1)$:

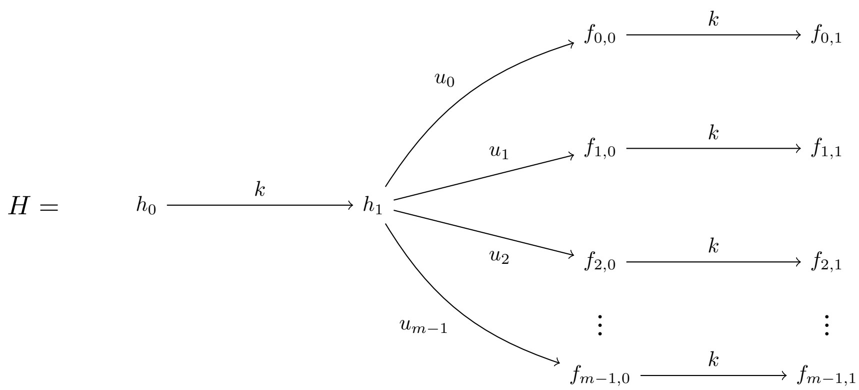

Moreover, for every $i \in \{0, \ldots, m-1\}$, the TS A has the following gadget T_i that represents the clause $M_i = \{X_{i_0}, X_{i_1}, X_{i_2}\}$ by using its variables as events, and uses the event u_i again.

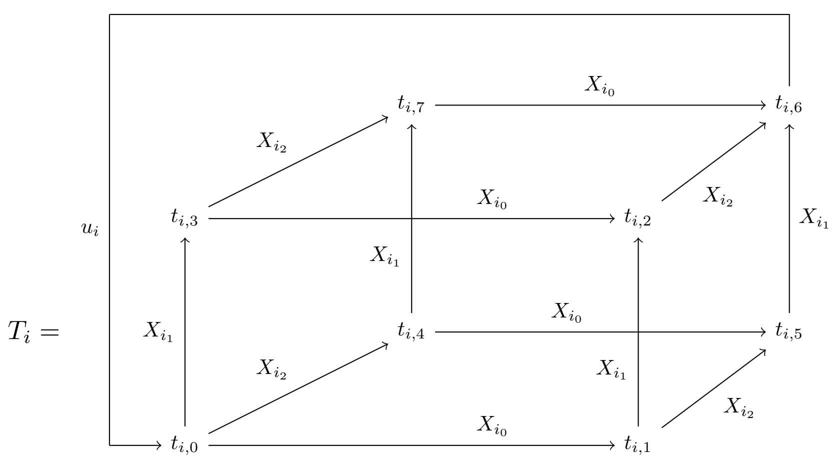

Finally, the TS $A = (S, E, \delta, \iota)$ has the initial state ι from which all introduced gadgets are reachable by unambiguous labeled edges: for every $i \in \{0, \ldots, m-1\}$, the TS A has the edge $\iota \xrightarrow{a_i} t_{i,0}$, and, moreover, it has the edge $\iota \xrightarrow{a_m} h_0$. Note that $E = \mathfrak{U} \cup \{k\} \cup \{a_0, \ldots, a_m\} \cup \{u_0, \ldots, u_{m-1}\}$, and $|E| = 3m + 2$. In the following, for any gadget G, we shall denote by $S(G)$ the set of all its states.

Lemma 2. *If there is a location map $\lambda : E \to \mathcal{L}$ and a λ-localizable admissible set $\mathcal{R}$ of A, i. e., for all $e \neq e' \in E$ and all $R = (sup, con, pro) \in \mathcal{R}$, if $con(e) > 0$ and $con(e') > 0$, then $\lambda(e) = \lambda(e')$, then there is a one-in-three model for $(\mathfrak{U}, M)$.*

Proof. We show that if $R = (sup, con, pro)$ is a λ-distributable region of $\mathcal{R}$ that solves (k, h_1), then the set $\mathfrak{S} = \{X \in \mathfrak{U} \mid con(X) > 0\}$ defines a one-in-three model of $(\mathfrak{U}, M)$.

We first argue that $\lambda(a_i) = \lambda(a_j)$ for all $i \neq j \in \{0, \dots, m\}$: Since $\mathcal{R}$ witnesses the ESSP of A, and $t_{j,0} \xrightarrow{\neg a_i}$, there is a region $R = (sup, con, pro) \in \mathcal{R}$ that solves the ESSA $(a_i, t_{j,0})$. By $\iota \xrightarrow{a_i}$, we have $con(a_i) \leq sup(\iota)$, and, since R solves $(a_i, t_{j,0})$, we have $con(a_i) > sup(t_{j,0}) \geq 0$. Together this implies $sup(\iota) > sup(t_{j,0})$, and thus $con(a_j) > pro(a_j) \geq 0$, since $sup(t_{j,0}) = sup(\iota) - con(a_j) + pro(a_j)$. Hence, by $con(a_i) > 0$, and $con(a_j) > 0$, we obtain $\lambda(a_i) = \lambda(a_j)$.

Similarly, one argues that if $i \neq j \in \{0, \dots, m-1\}$ are arbitrary but fixed, then $\lambda(u_i) = \lambda(u_j)$, which results from a region that solves $(u_i, f_{j,0})$. Hence, $\lambda(u_i) = \lambda(u_j)$ for all $i \neq j \in \{0, \dots, m\}$.

Let $R = (sup, con, pro)$ be a region of $\mathcal{R}$ that solves (k, h_1). (Note that R exists, since $\mathcal{R}$ is an admissible set.) We first show now that the set $\mathfrak{S} = \{X \in \mathfrak{U} \mid con(X) > 0\}$ contains at least $\frac{m}{3}$ elements (which thus have all the same location as k): Let $i \in \{0, \dots, m-1\}$ be arbitrary but fixed. By $h_0 \xrightarrow{k}$, we have $con(k) \leq sup(h_0)$, and since R solves (k, h_1), we have $con(k) > sup(h_1)$. On the other hand, by $f_{i,0} \xrightarrow{k}$, we have $con(k) \leq sup(f_{i,0})$, which implies $sup(h_1) < sup(f_{i,0})$, and thus $con(u_i) < pro(u_i)$. By $t_{i,6} \xrightarrow{u_i} t_{i,0}$, and $con(u_i) < pro(u_i)$, we have that $sup(t_{i,0}) > sup(t_{i,6})$. This implies that there is an event $X \in \{X_{i_0}, X_{i_1}, X_{i_2}\}$ on the path $t_{i,0} \xrightarrow{X_{i_0}} t_{i,1} \xrightarrow{X_{i_1}} t_{i,2} \xrightarrow{X_{i_2}} t_{i,6}$ that satisfies $con(X) > 0$. This is due to the fact that $sup(t_{i,6}) = sup(t_{i,0}) - (\sum_{j=0}^{2} con(X_{i_j})) + (\sum_{j=0}^{2} pro(X_{i_j}))$. Since i was arbitrary, this is simultaneously true for all $i \in \{0, \dots, m-1\}$. Hence, as every $X \in \mathfrak{U}$ occurs in exactly three distinct clauses, say M_i, M_j, M_ℓ (corresponding to T_i, T_j, T_ℓ), we have $3|\mathfrak{S}| \geq m$, and thus $|\mathfrak{S}| \geq \frac{m}{3}$. Moreover, for all $X \in \mathfrak{S}$, it holds $\lambda(k) = \lambda(X)$.

Finally, we argue that $\mathfrak{S}$ contains exactly $\frac{m}{3}$ elements: Since λ is a surjective mapping, and $|\{1, \dots, \frac{2m}{3} + 3\} \setminus \{\lambda(k), \lambda(a_0), \lambda(u_0)\}| \geq \frac{2m}{3}$, and $\lambda(a_0) = \dots = \lambda(a_m)$, and $\lambda(u_0) = \dots = \lambda(u_{m-1})$, there have to be $\frac{2m}{3}$ pairwise distinct events left that correspond to the remaining locations, i. e., we have that $|E \setminus (\{k, a_0, \dots, a_m, u_0, \dots, u_{m-1}\} \cup \mathfrak{S})| = |(\mathfrak{U} \setminus \mathfrak{S})| \geq \frac{2m}{3}$. By $|\mathfrak{U}| = m$, and $|\mathfrak{S}| \geq \frac{m}{3}$, this implies $|\mathfrak{S}| = \frac{m}{3}$. In particular, by Lemma 1, we obtain that $\mathfrak{S}$ defines a one-in-three model of $(\mathfrak{U}, M)$. This proves the lemma. $\square$

In order to complete the proof of the adequacy of our reduction, we now show that the existence of a one-in-three model for $(\mathfrak{U}, M)$ implies the existence of a location map $\lambda : E \to \{1, \dots, \frac{2m}{3} + 3\}$ such that there is a λ-localizable admissible set $\mathcal{R}$ of A. So let $\mathfrak{S}$ be a one-in-three model of $(\mathfrak{U}, M)$, and let $\mathfrak{U} \setminus \mathfrak{S} = \{X_{j_1}, \dots, X_{j_{\frac{2m}{3}}}\}$ be set of all variable events, which do not participate at the model. For all $e \in E$, we define λ as follows:

$$\lambda(e) = \begin{cases} 1, & \text{if } e \in \{k\} \cup \mathfrak{S} \\ 2, & \text{if } e \in \{a_0, \dots, a_m\} \\ 3, & \text{if } e \in \{u_0, \dots, u_{m-1}\} \\ \ell + 3, & \text{if } e = X_{j_\ell} \text{ for some } \ell \in \{1, \dots, \frac{2m}{3}\} \end{cases}$$

The following facts show that A's events are solvable by λ-localizable regions. Due to space restrictions, we present regions $R = (sup, con, pro)$ only implicitly by $sup(\iota)$, and con, and pro; it will be easy to chack that the definitions are coherent, i.e., ny support is negative and two different paths to a same state do not lead to different supports. We summarize events by $\mathcal{E}^R_{c,p} = \{e \in E \mid con(e) = c \text{ and } pro(e) = p\}$. If $e \in E$ is not explicitly mentioned in a set $E_{c,p}$, where $c \neq 0$ or $p \neq 0$, then $e \in \mathcal{E}^R_{0,0} = E \setminus \{e \in E \mid con(e) \neq 0 \text{ or } pro(e) \neq 0\}$, and we leave this set implicitly defined in the obvious way.

In order to help the reader understand the regions presented in Fact 1 to Fact 4, and Lemma 3, we gathered in an Appendix several figures illustrating the gadgets $H, T_0, \dots, T_5$ of the TS A that would be the result of the reduction applied on the instance of Example 1. For every figure, the coloring of the states corresponds to the support of the states according to the region addressed by the figure: red colored states have support 1, green colored states have support 2, blue colored states have support 3, and states without color have support 0. These figures are intended to be withdrawed in a ready to publish version, to cope with length constraints.

Fact 1. *The event k is solvable by λ-localizable regions.*

Proof. The following region (see Figure 6) $R_1 = (sup_1, con_1, pro_1)$ solves (k, s) for all $s \in \{h_1\} \cup \bigcup_{i=0}^{m-1}\{f_{i,1}\}$: $sup_1(\iota) = 1$, and $\mathcal{E}^{R_1}_{1,0} = \{k\} \cup \mathfrak{S}$, and $\mathcal{E}^{R_1}_{0,1} = \{u_0, \dots, u_{m-1}\}$.

The following region (see Figure 7) $R_2 = (sup_2, con_2, pro_2)$ solves (k, s) for all states $S \setminus S(H)$: $sup_2(\iota) = 0$, and $\mathcal{E}^{R_2}_{1,1} = \{k\}$, and $\mathcal{E}^{R_2}_{0,1} = \{a_m\}$. □

Fact 2. *If $e \in \{a_0, \dots, a_m\}$, then e is solvable by λ-localizable regions.*

Proof. The following region $R_3 = (sup_3, con_3, pro_3)$ solves (e, s) for all $e \in \{a_0, \dots, a_m\}$ and all $s \in S \setminus \{\iota\}$: $sup_3(\iota) = 1$, and $\mathcal{E}^{R_3}_{1,0} = \{a_0, \dots, a_m\}$. □

Fact 3. *If $e \in \{u_0, \dots, u_{m-1}\}$, then e is solvable by λ-localizable regions.*

Proof. The following region $R_4 = (sup_4, con_4, pro_4)$ solves (u, s) for all $u \in \{u_0, \dots, u_{m-1}\}$, and all $s \in S \setminus (\{h_1\} \cup \bigcup_{i=0}^{m-1}\{t_{i,6}\})$: $sup_4(\iota) = 0$, and $\mathcal{E}^{R_4}_{3,0} = \{u_0, \dots, u_{m-1}\}$, and $\mathcal{E}^{R_4}_{0,2} = \{k\}$, and $\mathcal{E}^{R_4}_{0,1} = \{a_m\} \cup \mathfrak{U}$. See Figure 8, Appendix A.

Let $j \in \{0, \dots, m-1\}$ be arbitrary but fixed. The following region $R^j_5 = (sup^j_5, con^j_5.pro^j_5)$ solves (u_j, s) for all $s \in (\bigcup_{i=0}^{m-1}\{t_{i,6}\}) \setminus \{t_{j,6}\}$: $sup^j_5(\iota) = 0$, and $\mathcal{E}^{R^j_5}_{1,1} = \{u_j\}$, and $\mathcal{E}^{R^j_5}_{0,1} = \{a_j, a_m\}$. See Figure 9, Appendix A. □

Fact 4. *For every $e \in \mathfrak{U}$, the event e is solvable by λ-localizable regions.*

Proof. Let $i \in \{0, \dots, m-1\}$ be arbitrary but fixed, and let $i_0, i_1, i_2 \in \{0, \dots, m-1\}$ be the three pairwise distinct indices such that $X_i \in M_{i_j}$ for all $j \in \{0, 1, 2\}$. The following region $R_6^i = (sup_6^i, con_6^i, pro_6^i)$ solves (X_i, s) for all $s \in S \setminus (\{s \in S \mid s \xrightarrow{X_i}\} \cup S(H))$: $sup_6^i(\iota) = 0$, and $\mathcal{E}_{1,0}^{R_6^i} = \{X_i\}$, and $\mathcal{E}_{0,1}^{R_6^i} = \{a_{i_0}, a_{i_1}, a_{i_2}, u_{i_0}, u_{i_1}, u_{i_2}\}$. See Figure 10, Appendix A.

The following region $R_7^i = (sup_7^i, con_7^i, pro_7^i)$ solves (X_i, s) for all $s \in S(H)$: $sup_7^i(\iota) = 0$, and $\mathcal{E}_{1,1}^{R_7^i} = \{X_i\}$, and $\mathcal{E}_{0,1}^{R_7^i} = \{a_{i_0}, a_{i_1}, a_{i_2}\}$. See Figure 11, Appendix A. Since i was arbitrary, this proves the lemma. $\square$

The following lemma completes the proof of Theorem 4:

Lemma 3. *If there is a one-in-three model for $(\mathfrak{U}, M)$, then there is a location map $\lambda : E \to \{1, \dots, \frac{2m}{3} + 3\}$, and a λ-localizable admissible set $\mathcal{R}$ of A.*

Proof. By Fact 1 to Fact 4, there are enough λ-localizable regions of A that witness the ESSP of A. Moreover, the region R_3 of Fact 2 solves (ι, s) for all $s \in S \setminus \{\iota\}$. Furthermore, if $i \in \{0, \dots, m-1\}$ is arbitrary but fixed, then the following region $R_8^i = (sup_8^i, con_8^i, pro_8^i)$, which is defined by $sup_8^i(\iota) = 0$, and $\mathcal{E}_{0,1}^{R_8^i} = \{a_i\}$, solves (s, s') for all $s \in S(T_i)$ and all $S \setminus S(T_i)$. See Figure 12, Appendix A. Hence, it remains to argue for the solvability of the SSA (s, s') such that s and s' belong to the same gadget of A.

Let $i \in \{0, \dots, m-1\}$ be arbitrary but fixed. One finds out that the regions $R_6^{i_0}$, and $R_6^{i_1}$, and $R_6^{i_2}$ that are defined in Fact 4 in order to solve the events X_{i_0}, and X_{i_1}, and X_{i_2}, respectively, altogether solve all SSA of T_i.

Hence, it remains to consider the SSA of H. Let $i \neq j \in \{0, \dots, m-1\}$ be arbitrary but fixed. The region R_1 of Fact 1 solves (h_0, h_1), and $(f_{i,0}, f_{i,1})$, and the region $R_6^{i_0}$ of Fact 4 solves (s, s') for all $s \in \{h_0, h_1\}$, and all $s' \in \{f_{i,0}, f_{i,1}\}$.

It remains to show that (s, s') is solvable for all $s \in \{f_{i,0}, f_{i,1}\}$, and all $s' \in \{f_{j,0}, f_{j,1}\}$. In order to do that, we observe that there is a $\ell \in \{0, 1, 2\}$, such that $X_{i_\ell} \notin M_j$, since M_i, and M_j would be equal otherwise. Hence, the region $R_6^{i_\ell}$ of Fact 4 solves (s, s'). By the arbitrariness of i, and j, we have finally argued that there is a witness of λ-localizable regions for the SSP of A. $\square$

Combining the various results of this section, we thus get our main result:

Theorem 4. *κ-Distributability is NP-complete.*

5 Conclusion

In this paper, we show that the problem of finding an optimal distributed implementation of a given TS A is a computationally hard problem by showing that the corresponding decision problem is NP-complete. The presented reduction is crucially based on the fact that the transitions of λ-distributed Petri nets may simultaneously consume and produce from the same place. Future work could therefore investigate the complexity of the problem restricted to pure Petri nets. Also, one may investigate whether the parameterized version of the problem is fixed parameter tractable when κ is chosen as the parameter.

References

1. van der Aalst, W.M.P.: Process Mining - Discovery, Conformance and Enhancement of Business Processes. Springer (2011). https://doi.org/10.1007/978-3-642-19345-3
2. Badouel, E., Bernardinello, L., Darondeau, P.: Petri Net Synthesis. Texts in Theoretical Computer Science. An EATCS Series, Springer (2015). https://doi.org/10.1007/978-3-662-47967-4
3. Badouel, E., Caillaud, B., Darondeau, P.: Distributing finite automata through Petri net synthesis. Formal Asp. Comput. **13**(6), 447–470 (2002). https://doi.org/10.1007/s001650200022
4. Best, E., Darondeau, P.: Petri net distributability. In: Ershov Memorial Conference. Lecture Notes in Computer Science, vol. 7162, pp. 1–18. Springer (2011). https://doi.org/10.1007/978-3-642-29709-0_1
5. Cortadella, J., Kishinevsky, M., Kondratyev, A., Lavagno, L., Yakovlev, A.: A region-based theory for state assignment in speed-independent circuits. IEEE Trans. on CAD of Integrated Circuits and Systems **16**(8), 793–812 (1997). https://doi.org/10.1109/43.644602
6. Desel, J., Reisig, W.: The synthesis problem of Petri nets. Acta Inf. **33**(4), 297–315 (1996). https://doi.org/10.1007/s002360050046
7. Devillers, R.: Factorisation of transition systems. Acta Informatica **55**(4), 339–362 (2018). https://doi.org/10.1007/s00236-017-0300-y
8. Holloway, L.E., Krogh, B.H., Giua, A.: A survey of Petri net methods for controlled discrete event systems. Discrete Event Dynamic Systems **7**(2), 151–190 (1997). https://doi.org/10.1023/A:1008271916548
9. Moore, C., Robson, J.M.: Hard tiling problems with simple tiles. Discrete & Computational Geometry **26**(4), 573–590 (2001). https://doi.org/10.1007/s00454-001-0047-6
10. Mukund, M.: Transition system models for concurrency. DAIMI Report Series **21**(399) (Jun 1992). https://doi.org/10.7146/dpb.v21i399.6633, https://tidsskrift.dk/daimipb/article/view/6633
11. Starke, P.H.: Analyse von Petri-Netz-Modellen. Leitfäden und Monographien der Informatik, Teubner (1990)

A Graphical Supports for the Regions of Fact 1 to Fact 4, and Lemma 3

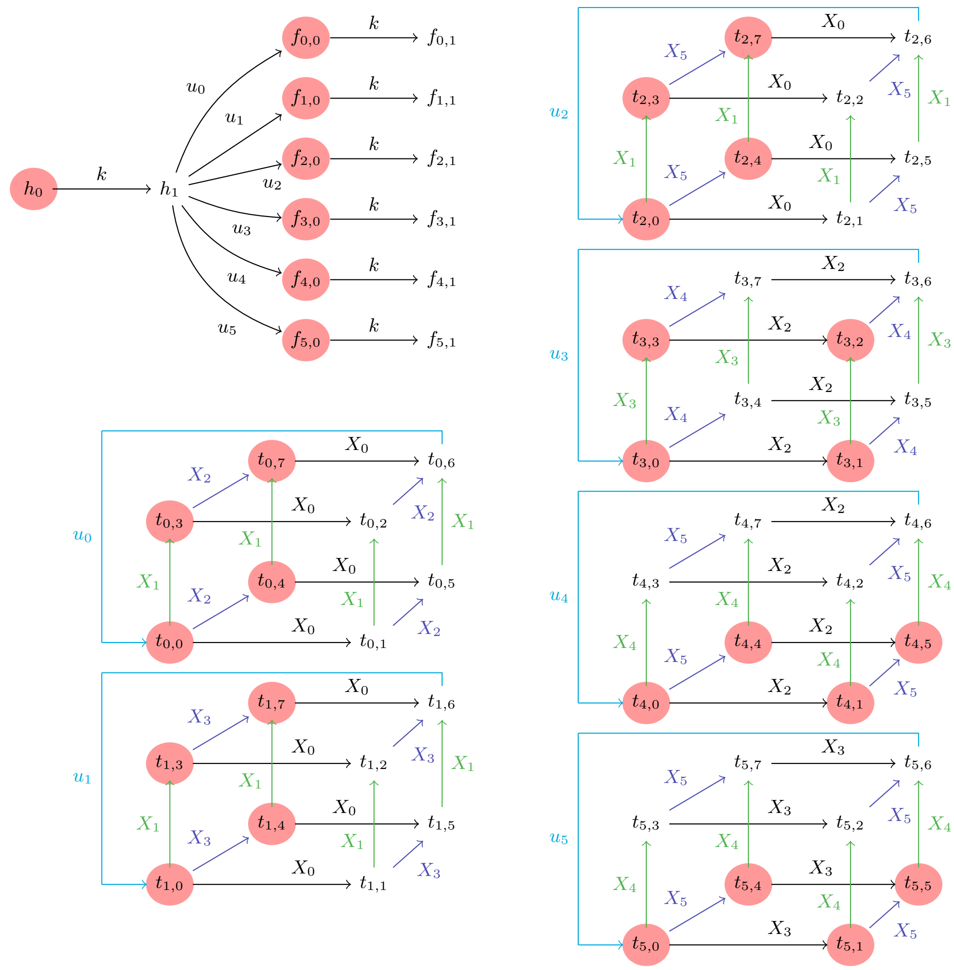

Fig. 6: A sketch of the region R_1 of Fact 1.

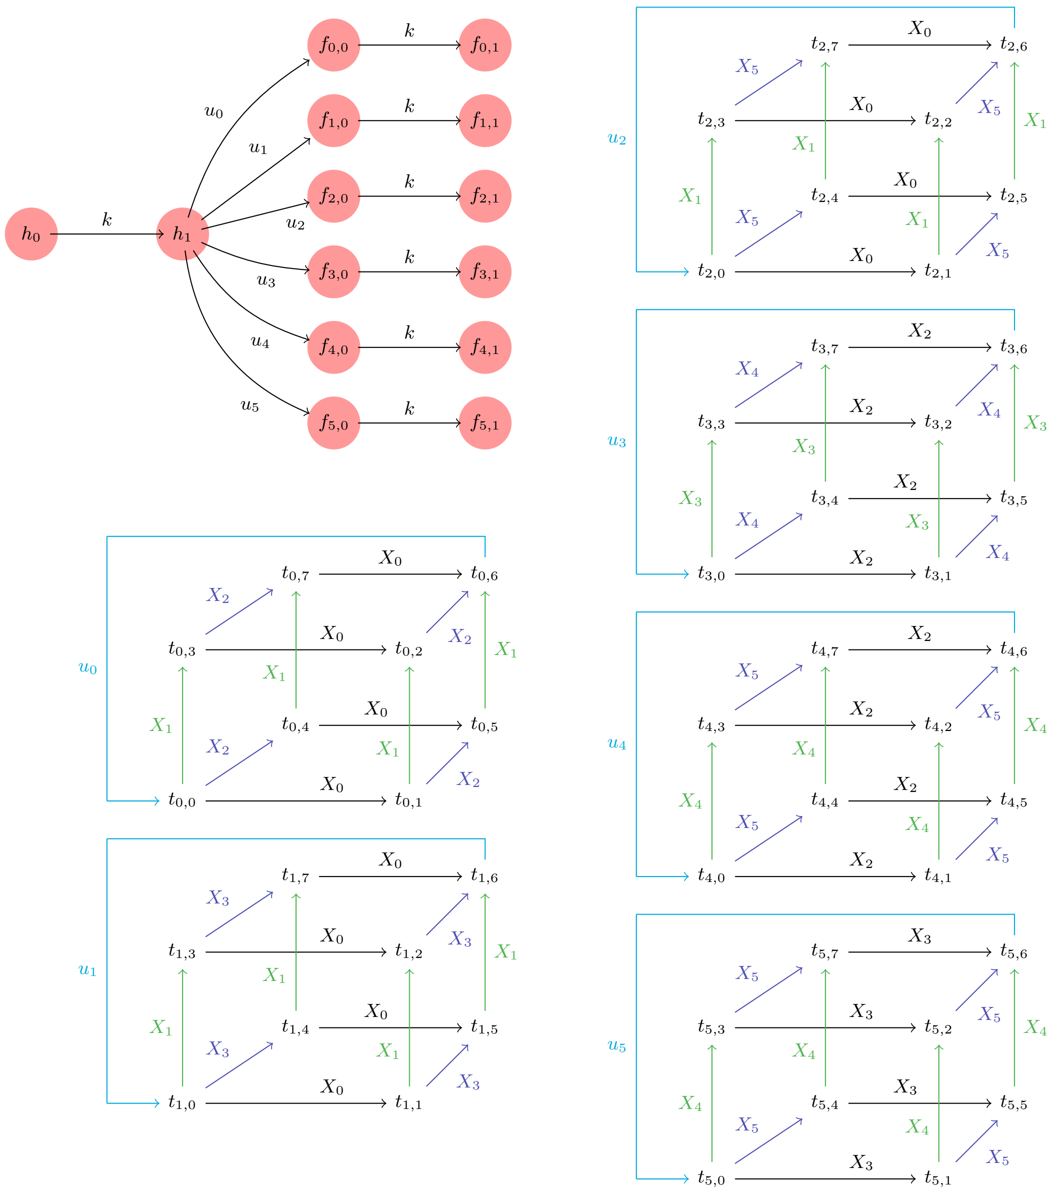

Fig. 7: A sketch of the region R_2 of Fact 1.

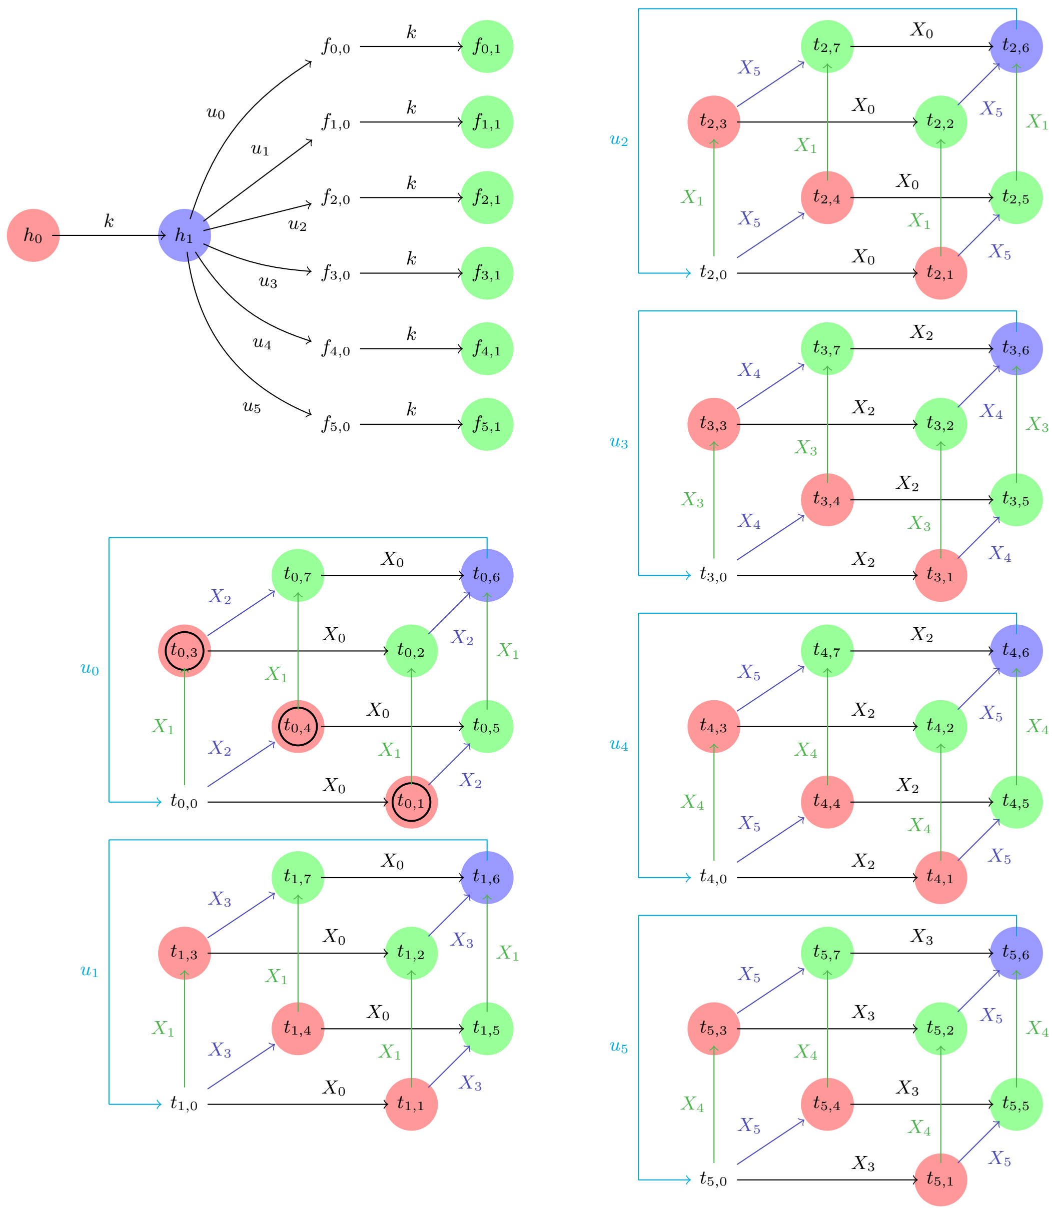

Fig. 8: A sketch of the region R_4 of Fact 3.

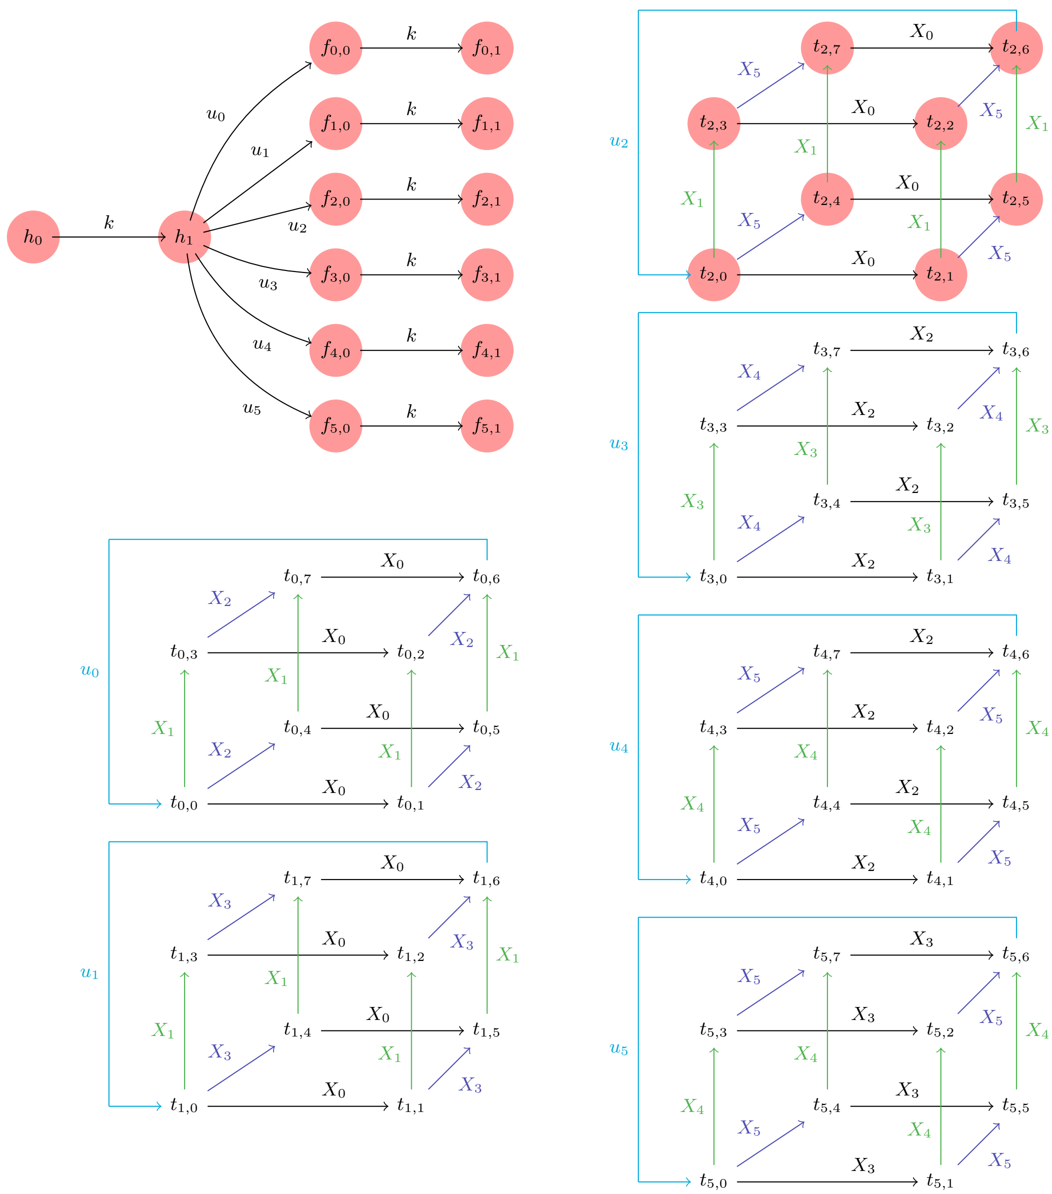

Fig. 9: A sketch of the region R_5^2 of Fact 3.

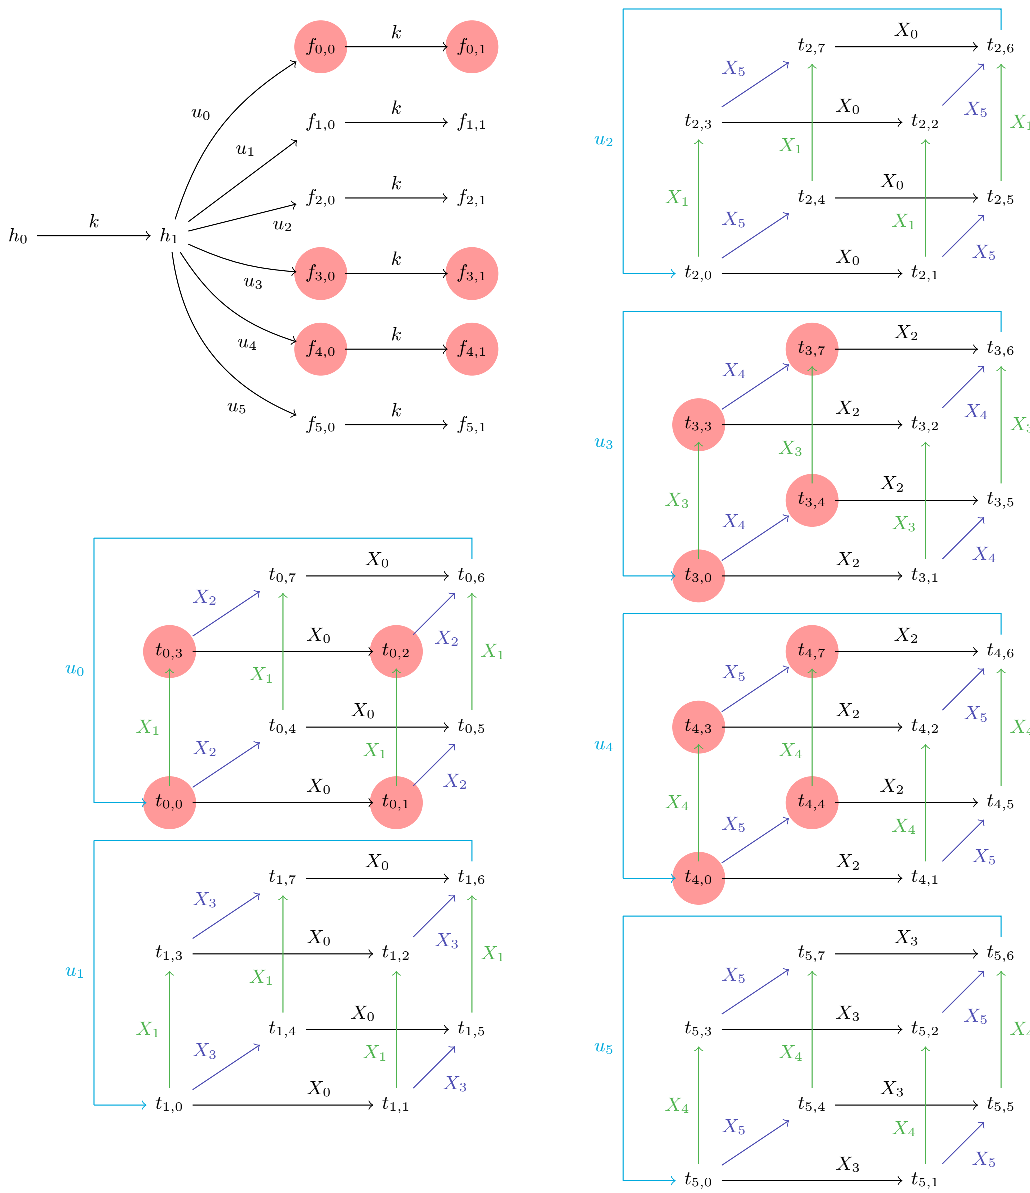

Fig. 10: A sketch of the region R_6^2 of Fact 4.

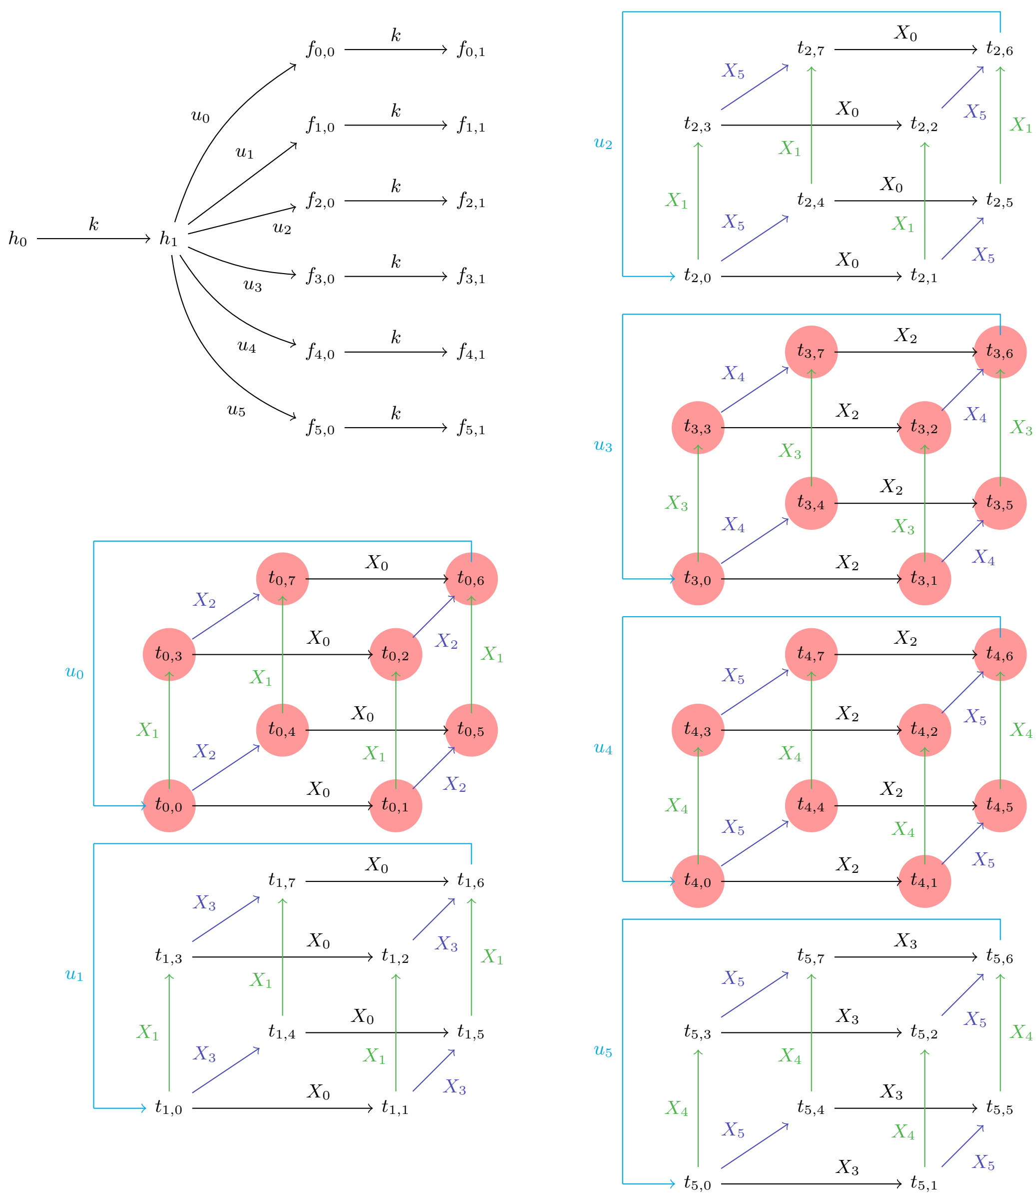

Fig. 11: A sketch of the region R_7^2 of Fact 4.

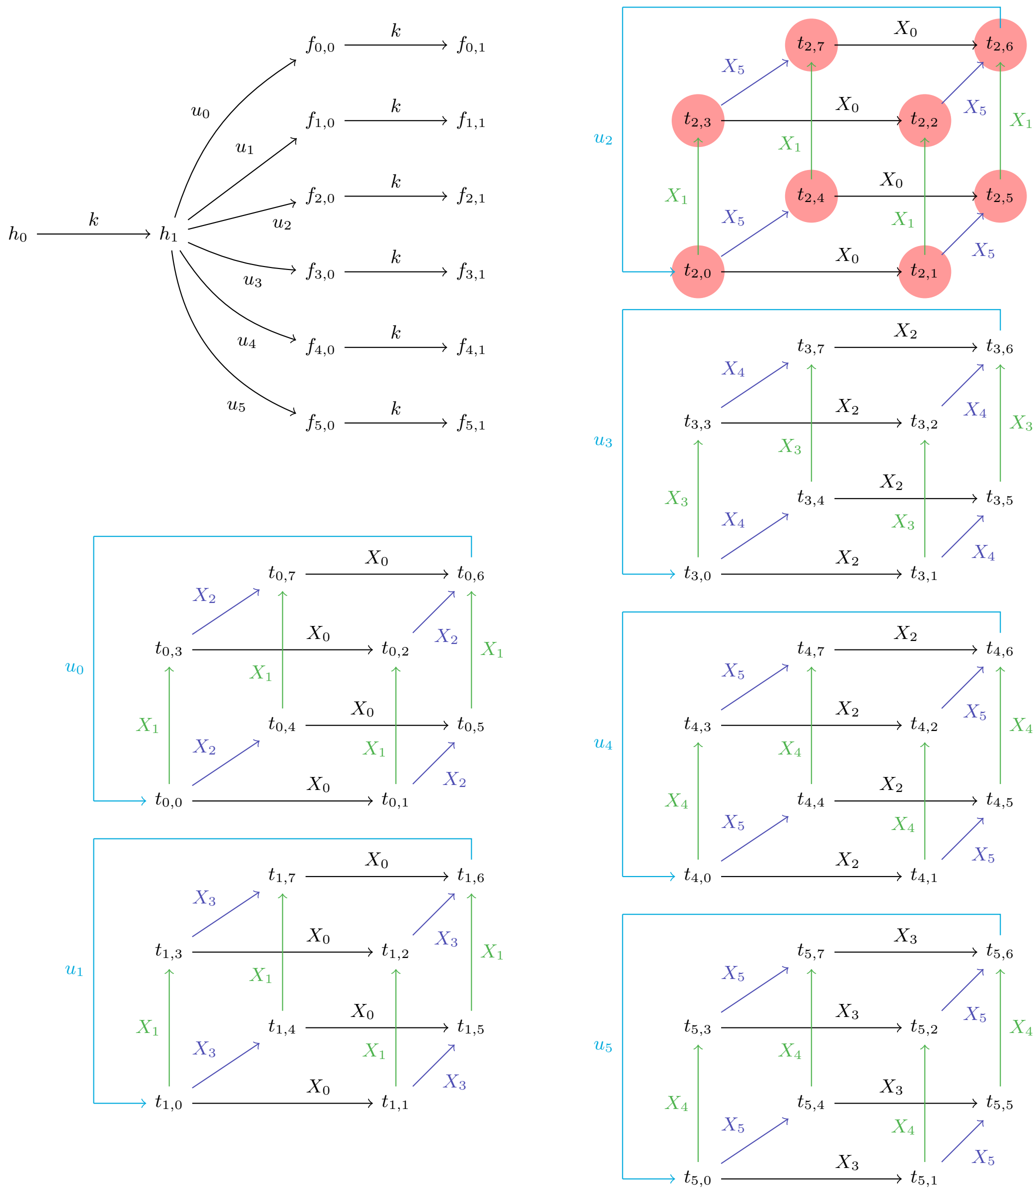

Fig. 12: A sketch of the region R_8^2 of Lemma 3.