



HAL
open science

A Kleene theorem for Petri automata

Paul Brunet

► **To cite this version:**

| Paul Brunet. A Kleene theorem for Petri automata. 2016. hal-01258754v2

HAL Id: hal-01258754

<https://hal.science/hal-01258754v2>

Preprint submitted on 20 Jan 2016

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

A Kleene theorem for Petri automata

Paul Brunet

Plume team – LIP, Université de Lyon, CNRS, ENS de Lyon, Inria, UCBL

paul.brunet@ens-lyon.fr

Abstract

While studying the equational theory of Kleene Allegories (KA), we recently proposed two ways of defining sets of graphs [BP15]: from KA expressions, that is, regular expressions with intersection and converse; and from a new automata model, Petri automata, based on safe Petri nets. To be able to compare the sets of graphs generated by KA expressions, we explained how to construct Petri automata out of arbitrary KA expressions.

In the present paper, we describe a reverse transformation: recovering an expression from an automaton. This has several consequences. First, it generalises Kleene theorem: the graph languages specified by Petri automata are precisely the languages denoted by KA expressions. Second, this entails that decidability of the equational theory of Kleene Allegories is equivalent to that of language equivalence for Petri automata. Third, this transformation may be used to reason syntactically about the occurrence nets of a safe Petri net, provided they are parallel series.

Keywords regular expressions, graph language, Petri nets, finite automata, allegories, intersection, converse, Kleene theorem

1. Introduction

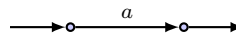
Petri Automata are Petri net-based automata, and are equipped with an operational semantics allowing them to recognise sets of graphs. They were introduced in a recent paper [2] to study Kleene Allegory expressions: terms built from a finite alphabet of variables, with the constants 0 and 1, the unary operators converse and Kleene star, and the binary operators union, composition and intersection. To any such expression, one can associate a set of Two Terminal Series Parallel graphs (called graphs in the sequel) such that two expressions are universally equivalent when interpreted as binary relations if and only if their associated sets of graphs are equal. In [2] we gave a method to build from any Kleene Allegory expression a Petri automaton recognising the appropriate set of series parallel graphs.

It is then natural to wonder what is the class of graph languages recognised by Petri automata. For the usual notion of finite state automata, the well known Kleene Theorem states that the languages recognisable by automata are exactly those specifiable by regular expressions. In this paper we provide a similar theorem for Petri automata and Kleene Allegory expressions.

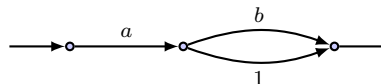
In the remainder of this introduction, we explain informally what are the graphs associated with expressions and Petri automata and we give the outline of the computation of the expression corresponding to a given automaton.

1.1 Expressions

For the bulk of the development, we restrict ourselves to a simpler setting: we consider expressions over the syntax of Kleene lattices, that is to say without the converse operation. To any such expression e , we associate a set $\mathcal{G}(e)$ of labelled directed graphs with two distinguished vertices understood as input and output. A letter a is associated to a singleton set containing the following graph:



The input in this graph is the vertex with the unmarked incoming arrow, and the output is the other vertex, with the outgoing arrow. Similarly the constant 1 yields a singleton set containing the graph where the input is linked to the output by an edge labelled 1. We define the sequential composition of two such graphs by identifying the output of the first graph with the input of the second one. We also build the parallel composition (corresponding to the intersection operation) of two graphs by identifying both their inputs and their outputs. For instance $\mathcal{G}(a \cdot (b \cap 1))$ simply contains the graph:



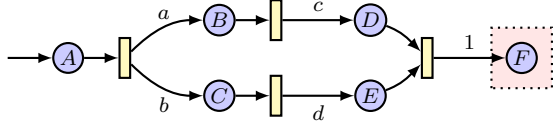
With these constructs we can define the set of graphs associated to an expression by structural induction: $\mathcal{G}(1)$ and $\mathcal{G}(a)$ are defined as before, $\mathcal{G}(e \cup f)$ is the union of $\mathcal{G}(e)$ and $\mathcal{G}(f)$; $\mathcal{G}(0)$ is the empty set; $\mathcal{G}(e \cdot f)$ (respectively $\mathcal{G}(e \cap f)$) is the set of graphs obtained by composing in sequence (resp. in parallel) a graph in $\mathcal{G}(e)$ with a graph in $\mathcal{G}(f)$; and finally the graphs in $\mathcal{G}(e^*)$ are either the graph of 1 or built by sequentially composing any finite number of graphs from $\mathcal{G}(e)$.

1.2 Petri automata

A Petri automaton over the alphabet Σ is defined by:

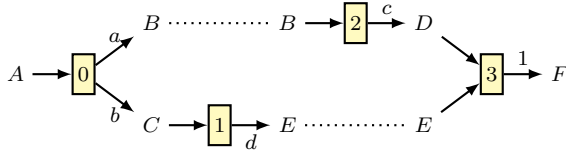
- a finite set of places (denoted in graphical representations by round vertices),
- a finite set \mathcal{T} of transitions (denoted by rectangular vertices). Each transition is equipped with a set of input places (incoming arrows) and a set of outputs. These outputs are pairs of a label in $\Sigma \cup \{1\}$, and a place.
- an initial place and a finite set of final configurations, a configuration being a set of places.

Here is an example of Petri automaton:

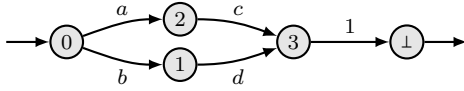


The initial place is A , as denoted by the unmarked incoming arrow. There is a single final configuration, containing only the place F (as indicated by the red dotted rectangle). In this paper, we will restrict ourselves to automata having a single final configuration, composed of one place.

Runs in such an automaton are plays of a token firing game: one starts by placing tokens on certain places, and if every place in the input of a transition holds a token, this transition is *enabled* and may be fired. Firing a transition will then remove one token from each of its inputs, and place one on each output. If the run starts with a single token on the initial place, we call it *initial*, and if it stops with one token on each place of a final configuration it is *final*. An *accepting* run is both initial and final. For instance, here is a graphical representation of an accepting run in the previous automaton:



We further impose that Petri automata should be *safe*, meaning that at any step of any initial run of the automaton at most one token is present on each place. The last step to get the language of an automaton is to generate a graph from any run of the automaton (called the *trace* of the run). We do so by considering a graph whose vertices are the transitions (plus an additional final vertex), and labelled edges are extrapolated from the outputs of the transitions. For instance the run from the previous example yields the graph:



We may thus define the language $\mathcal{G}(\mathcal{A})$ of an automaton \mathcal{A} as the set of the traces of its accepting runs. In this paper we only consider automata whose language contain exclusively parallel series graphs (because we are only interested in such graphs). As a byproduct of the construction presented here we show that this property is decidable.

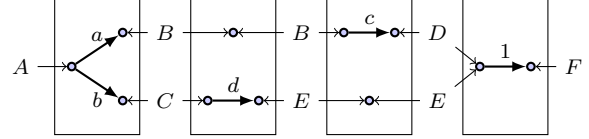
Remark 1. The definitions we give here are sometimes different from those stated in [2], but the results apply nonetheless. We discuss these differences in Section 5.1.

1.3 Main result

The main technical result of this paper is a construction allowing one to build from an automaton \mathcal{A} an expression $\mathcal{E}(\mathcal{A})$ such that:

$$\mathcal{G}(\mathcal{A}) = \mathcal{G}(\mathcal{E}(\mathcal{A})).$$

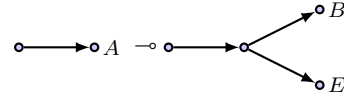
Before properly describing the construction, we need a compositional way to represent a local and partial view of a graph. This will be achieved through the notions of boxes and types. In the same way a sequence of letters may be a factor of some word, boxes are meant to represent “slices” of graphs. Formally, if A and B are sets of places, a box going from A to B is a labelled directed graph G with maps from A to the vertices of G and from B to the vertices of G . For instance, in the run presented before as an example, we may associate a box to each transition. This yields the following sequence of boxes:



We may compose these boxes by merging the outputs of the first box with the inputs of the second box, and reducing the graph by applying the operators \cdot and \cap on edges. For instance the first two boxes above compose to produce the box β , and the whole sequence can be composed as the box γ :

$$\beta = A \begin{array}{c} \xrightarrow{a} B \\ \xrightarrow{b} C \end{array} \quad \gamma = A \xrightarrow{((a \cdot c) \cap (b \cdot d)) \cdot 1} F$$

Types in this setting are trees with leaves labelled with places. We can type a box β with a pair of trees $\sigma \multimap \tau$ if the tree σ composed with the box β yields the tree τ . For instance the box β has the type:



We build a finite state automata whose states are the types, and whose transitions between the states σ and τ are labelled with the boxes built from transitions of the Petri automaton with type $\sigma \multimap \tau$. Using the standard Kleene Theorem, this allows us to produce a regular expression over boxes $e_{\mathcal{A}}$. By changing the set of letters to finite sets of boxes labelled with expressions, we are then able to internalise all of the regular operations. This produces in the end a finite number of boxes with one input (corresponding to the initial place) and one output (corresponding to the final place), linked by an edge labelled with an expression. The union of these expressions is the expression $\mathcal{E}(\mathcal{A})$, that denotes the same set of graphs as \mathcal{A} .

The technical difficulty here is the computation of the star of a set of boxes. For instance, consider the box:

$$\delta = A \begin{array}{c} \xrightarrow{a} A \\ \xrightarrow{b} B \end{array}$$

Its star yields all boxes of the following shape, for $n \in \mathbb{N}$:

$$\delta^n = A \begin{array}{c} \xrightarrow{a^n} A \\ \xrightarrow{b^n} B \end{array}$$

This set of boxes cannot be represented by a finite number of boxes. However, one must realise that because the two branches of the box δ are disconnected, they must correspond to concurrent process in the Petri automaton, and thus can be iterated separately. Intuitively, the net cannot force two independent processes to perform exactly the same number of iterations. Because of this, the following box is the right choice to represent the star of the box δ :

$$\delta^* \equiv A \begin{array}{c} \xrightarrow{a^*} A \\ \xrightarrow{b^*} B \end{array}$$

1.4 Outline

We begin the paper by introducing in Section 2 the expressions and automata we consider here, and the sets of graphs they represent. After a few definitions and results about graphs, we detail in Section 3 the types and boxes we use, and we establish some facts

about them. Section 4 is dedicated to the main contribution of the paper, the construction extracting expressions from Petri automata. Finally we list some consequences of this result, as well as some links with previous work, in Section 5.

2. Petri automata & expressions

Regular expressions and rational languages We recall a few standard definitions and notations about regular expressions. A *regular expression* over the alphabet Σ is a term built from variables taken in Σ , the constants 1 and 0, the binary operators \cdot and \cup , and the unary $_*$. The set is denoted by $\text{Reg}(\Sigma)\Sigma$. The *language denoted by e* , written $\llbracket e \rrbracket$, is a set of words over Σ defined inductively by:

$$\begin{aligned} \llbracket a \rrbracket &:= \{a\}; & \llbracket 1 \rrbracket &:= \{\epsilon\}; & \llbracket 0 \rrbracket &:= \emptyset; & \llbracket e \cup f \rrbracket &:= \llbracket e \rrbracket \cup \llbracket f \rrbracket; \\ \llbracket e \cdot f \rrbracket &:= \{u \cdot v \mid u \in \llbracket e \rrbracket \text{ and } v \in \llbracket f \rrbracket\}; \\ \llbracket e^* \rrbracket &:= \{u_1 \cdot \dots \cdot u_n \mid n \in \mathbb{N}, \forall i, u_i \in \llbracket e \rrbracket\}. \end{aligned}$$

(ϵ being the empty word.)

If $\llbracket e \rrbracket = \llbracket f \rrbracket$, we say that regular laws prove the equality of e and f , written $KA \models e = f$.

2.1 Expressions

Expressions e, f, \dots are built on the signature $\langle 1, \cdot, \cup, \cap, _* \rangle^1$. We denote by Reg_Σ^1 the set of expressions over the finite alphabet Σ . A 2-pointed labelled directed graph G over $\Sigma_\bullet := \Sigma \cup \{1\}$, simply called graph in the sequel, is given by a tuple $\langle V, E, \iota, \circ \rangle$ with a finite set of vertices V , a set of labelled edges $E \subseteq V \times \Sigma_\bullet \times V$, and $\iota, \circ \in V$ two distinguished vertices, respectively called input and output. The sequential composition (written $_ \cdot _$) of two graphs with disjoint sets of vertices can be performed by identifying the output of the first graph and the input of the second one. Their parallel composition (written $_ \cap _$) consists in merging their inputs and merging their outputs. See Figure 1 for a graphical description of these constructions.

We assign to each expression a set of such graphs, called the graph language of the expression.

Definition 2 (Graph language of an expression: $\mathcal{G}(e)$)

The graph language is defined by structural induction as follows:

$$\begin{aligned} \mathcal{G}(a) &:= \left\{ \begin{array}{c} \text{---} \circ \xrightarrow{a} \circ \text{---} \\ \text{---} \circ \xrightarrow{1} \circ \text{---} \end{array} \right\} \\ \mathcal{G}(1) &:= \left\{ \text{---} \circ \xrightarrow{1} \circ \text{---} \right\} \\ \mathcal{G}(0) &:= \emptyset \\ \mathcal{G}(e \cup f) &:= \mathcal{G}(e) \cup \mathcal{G}(f) \\ \mathcal{G}(e \cdot f) &:= \{E \cdot F \mid E \in \mathcal{G}(e) \text{ and } F \in \mathcal{G}(f)\} \\ \mathcal{G}(e \cap f) &:= \{E \cap F \mid E \in \mathcal{G}(e) \text{ and } F \in \mathcal{G}(f)\} \\ \mathcal{G}(e^*) &:= \{E_1 \cdot \dots \cdot E_n \mid n > 0, \forall i, E_i \in \mathcal{G}(e)\} \end{aligned}$$

*

Those graphs were introduced independently by Freyd and Scedrov [3, page 208], and Andr eka and Bredikhin [1] (except for the way 1 is handled). Notice that the graphs produced by this construction are exactly Two Terminal Series Parallel graphs labelled with Σ_\bullet [8].

2.2 Petri automata

A Petri automaton is a Petri net whose transition's outputs are labelled by the set Σ_\bullet .

¹Notice that we don't use the Kleene star, but rather its positive variant. It is more convenient here to have e^* be defined as $1 \cup e^+$.

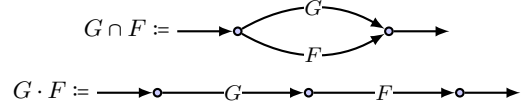


Figure 1: Elementary graph constructions

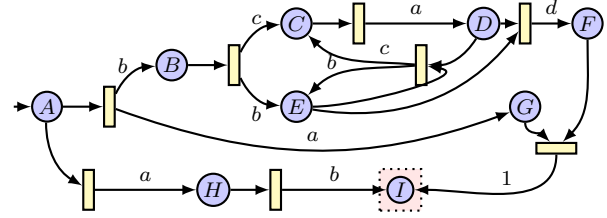


Figure 2: A Petri automaton. The initial place is A , and the final configurations is $\{I\}$.

Definition 3 (Petri Automaton)

A *Petri automaton* \mathcal{A} over the alphabet Σ is a tuple $\langle P, \mathcal{T}, \iota, f \rangle$ where:

- P is a finite set of *places*,
- $\mathcal{T} \subseteq \mathcal{P}(P) \times \mathcal{P}(\Sigma_\bullet \times P)$ is a set of *transitions*,
- $\iota \in P$ is the *initial place* of the automaton,
- $f \in P$ is the *final place* of the automaton.

For each transition $t = (\underline{t}, \bar{t}) \in \mathcal{T}$, \underline{t} and \bar{t} are assumed to be non-empty; $\underline{t} \subseteq P$ is the *input* of t ; and $\bar{t} \subseteq \Sigma_\bullet \times P$ is the *output* of t . *

This definition of Petri automata correspond to what we call *simple Petri automata* in [2]. However, this case is general enough, as discussed in Section 5.1. We use the graphical notation from the introduction to represent Petri automata; the Petri automaton from Figure 2 will be used as a running example.

From a configuration $\xi \subseteq P$, a transition $t = (\underline{t}, \bar{t}) \in \mathcal{T}$ is *enabled* if $\underline{t} \subseteq \xi$. If so, one may *fire* t , which produces a new configuration ξ' defined by

$$\xi' = \xi \setminus \underline{t} \cup \{p \in P \mid \exists x \in \Sigma : (x, p) \in \bar{t}\}.$$

We write $\xi \xrightarrow{t} \xi'$ in this case.

In the sequel, we assume all considered Petri automata to be *safe*. (I.e., in Petri nets terminology, such that any reachable marking has at most one token in each place [6]). Formally, with our definitions: a Petri automaton $\langle P, \mathcal{T}, \iota, f \rangle$ is *safe* if for all configuration $\xi \subseteq P$ reachable from $\{\iota\}$ by firing any number of transitions, if $(\underline{t}, \bar{t}) \in \mathcal{T}$ is enabled from ξ , $p \in \xi$, and $(x, p) \in \bar{t}$, then $p \in \underline{t}$.

Now we recall how to use Petri automata to define languages of graphs. We first define the *runs* of an automaton.

Definition 4 (Run, accepting run)

A *run* ξ in $\mathcal{A} = \langle P, \mathcal{T}, \iota, f \rangle$ is a sequence $\langle (\xi_k)_{0 \leq k \leq n}, (t_k)_{0 \leq k < n} \rangle$ of configurations and transitions, such that $\xi_k \subseteq P$, $t_k \in \mathcal{T}$ and $\forall k < n, \xi_k \xrightarrow{t_k} \xi_{k+1}$. When $\xi_0 = \{\iota\}$ and $\xi_n = \{f\}$, we call ξ an *accepting run*. *

(Note that a run ξ is uniquely determined by ξ_0 and the sequence (t_k) : all subsequent configurations can be computed deterministically.)

Example 5

Consider the following sequence of transitions from the automaton

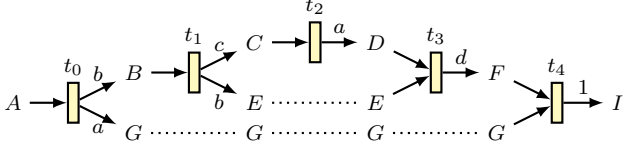


Figure 3: An accepting run in the automaton from Figure 2.




Figure 4: Trace of the run depicted in Figure 3.

in Figure 2:

$$\xi = \langle (\xi_0, \xi_1, \xi_2, \xi_3, \xi_4, \xi_5), (t_0, t_1, t_2, t_3, t_4) \rangle,$$

with

$$\begin{array}{l} \xi_0 = \{A\}, \\ \xi_1 = \{B, G\}, \\ \xi_2 = \{C, E, G\}, \\ \xi_3 = \{D, E, G\}, \\ \xi_4 = \{F, G\}, \\ \xi_5 = \{I\}. \end{array} \quad \left| \begin{array}{l} t_0 = (\{A\}, \{(b, B), (a, G)\}), \\ t_1 = (\{B\}, \{(c, C), (b, E)\}), \\ t_2 = (\{C\}, \{(a, D)\}), \\ t_3 = (\{D, E\}, \{(d, F)\}), \\ t_4 = (\{F, G\}, \{(1, I)\}). \end{array} \right.$$

We have:

$$\{A\} \xrightarrow{t_0} \{B, G\} \xrightarrow{t_1} \{C, E, G\} \xrightarrow{t_2} \{D, E, G\} \xrightarrow{t_3} \{F, G\} \xrightarrow{t_4} \{I\},$$

and since $\{A\}$ is the initial configuration and $\{I\}$ is the final one, this sequence is an accepting run. It can be represented graphically as in Figure 3. ■

In [2], the language of a Petri automaton is defined via an operational semantics, describing the way an automaton “reads” a graph. For the sake of clarity, we use here a different definition, more directly related to the set of accepting runs of the automaton. We associate to each accepting run ξ a 2-graph, written $\mathcal{G}(\xi)$, called the *trace* of ξ . The language of the automaton is then defined as the set of traces of its accepting runs.

The trace is constructed by creating a vertex k for each transition $t_k = (\underline{t}_k, \overline{t}_k)$ of the run, plus a final vertex n . We add an edge (k, x, l) whenever there is some place q such that $(x, q) \in \overline{t}_k$, and t_l is the first transition after t_k in the run with q among its inputs, or $l = n$ if there is no such transition in the run.

Definition 6 (Trace of a run)

Let $\xi = \langle (\xi_k)_{0 \leq k \leq n}, (\underline{t}_k, \overline{t}_k)_{0 \leq k < n} \rangle$ be run. For an index $k \leq n$ and a place q , let $\nu(k, q)$ be either the smallest index l such that $k \leq l$ and $q \in \underline{t}_l$, or n if there is no such index.

The *trace* of ξ is the graph $\mathcal{G}(\xi) := \langle \{0, \dots, n\}, E_{\xi}, 0, n \rangle$ with

$$E_{\xi} := \{ (k, x, \nu(k+1, q)) \mid (x, q) \in \overline{t}_k \}.$$

*

We write $\mathcal{G}(\mathcal{A})$ for the language of a Petri automaton \mathcal{A} , defined as the set of traces of its accepting runs. The trace of the run presented in Figure 3 is depicted in Figure 4.

Remark 7. $\mathcal{G}(\mathcal{A})$ is different from $\mathcal{L}(\mathcal{A})$, as it is defined in [2]. See section Section 5.1 for a more detailed discussion.

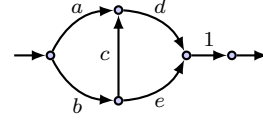


Figure 5: Non series parallel graph.

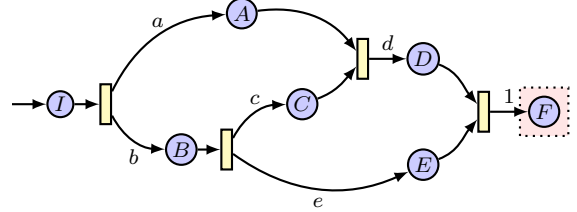


Figure 6: Non-SP Petri automaton.

Theorem 8 ([2, Theorem 17]). *For any expression $e \in \text{Reg}_{\Sigma}^{\square}$, there is a Petri automaton $\mathcal{A}(e)$ such that*

$$\mathcal{G}(e) = \mathcal{G}(\mathcal{A}(e)).$$

2.3 SP-Petri automata

When trying to build expressions out of Petri automata, one immediately encounters a problem: traces of accepting runs are not necessarily series parallel graphs, which means they cannot be generated by any expression from $\text{Reg}_{\Sigma}^{\square}$. For instance the automaton given in Figure 6 has a single accepting run, whose trace is the graph draw in Figure 5. To circumvent this problem, we consider only well-behaved automata: automata whose accepting runs have parallel-series graphs.

Definition 9 (SP-automata)

An automaton \mathcal{A} is said to be an *SP-automaton* if for any accepting run ξ in \mathcal{A} , $\mathcal{G}(\xi)$ is an SP-graph. *

Despite its infinitary formulation, this property is decidable. In fact, the procedure we describe later on would fail (in finite time) if it is given a non-SP automaton: hence it may be used to decide this property.²

3. Boxes

Notations We call *label space* a structure $\langle \mathbb{L}, \cdot, \cap \rangle$ such that \cdot and \cap are two internal associative binary operations, and \cap is commutative. Examples of such spaces include:

\mathbb{W}_X : the set of SP-graphs labelled by X . It is the free label space over a finite set of atomic variables X . We will make no distinctions between this set and the terms over the signature $\langle X, \cdot, \cap \rangle$ quotiented by the label space axioms.

$\mathbb{1}$: the unit label space.

\mathbb{E}_X : the set of expressions $\text{Reg}_{\Sigma}^{\square}$ (quotiented by the congruence generated by label space axioms).

3.1 Graphs

We introduce a few definitions and results about graphs. Unless otherwise stated, all graphs are labelled with some fixed label space \mathbb{L} .

²See Remark 18 for more details.

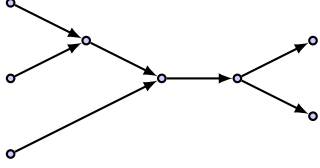


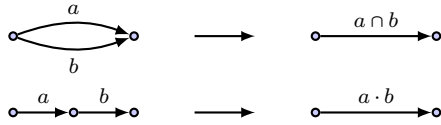
Figure 7: Example of reduced bow tie

DAGs and Series Parallel graphs Given a directed acyclic graph (DAG) $G = \langle V, E \rangle$, we may define $\min G$ to be the set of vertices in V with no incoming edge, and $\max G$ the set of vertices with no outgoing edge. A Two Terminal Parallel Series graph (SP-graph in the following) may be defined as a connected DAG that does not contain the graph from Figure 5 as a minor, and such that $\min G$ and $\max G$ are both singleton sets. A *factor* of $\langle V, E \rangle$ is a graph $\langle V', E' \rangle$ such that:

- $V' \subseteq V$ and $E' \subseteq E$;
- $e \in E'$ if and only if both its extremities are in V' .

A *prefix* of $\langle V, E \rangle$ is a factor $\langle V', E' \rangle$ such that if there is a path from $x \in V$ to $y \in V'$, then $x \in V'$. When considering the graph as a partial order, a prefix is a downward closed set of vertices. Similarly, A *suffix* of $\langle V, E \rangle$ is an upward closed set, meaning a factor $\langle V', E' \rangle$ such that if there is a path from $x \in V'$ to $y \in V$, then $y \in V'$. Finally a *slice* of $\langle V, E \rangle$ is a convex set, i.e. a factor such that if there is a path from $x \in V'$ to $y \in V$ and one from y to $z \in V'$, then $y \in V'$. Any slice may be obtained by removing from the graph a prefix and a suffix. A graph G is called an *SP-slice* if it is a slice of some SP-graph.

We define the SP-rewriting system, enriched with labels taken from a label space:



(This second rule can only be applied if the middle node on the left-hand side has no other adjacent edge.)

Given a DAG G , we write $G \downarrow$ for the unique normal form of G with respect to this system. One may notice here that if G is a non-trivial SP-graph, then $G \downarrow$ has always the shape $\rightarrow \bullet \xrightarrow{e} \bullet \rightarrow^3$, and G is isomorphic to $\mathcal{G}(e)$. This isomorphism means that we will often assimilate SP-graphs and terms built using the operators \cdot and \cap .

Trees and bow ties A *tree* is a DAG $T = \langle V, E \rangle$ such that either $\min T$ or $\max T$ contains a single node called the root, and for any two nodes $x, y \in V$ there exists at most one path from x to y . If $\min T$ is a singleton, then T is a *top-down tree*, otherwise it is a *bottom-up tree*. A tree is said to be proper if it does not contain any node with exactly one incoming edge and one outgoing edge. It is simple enough to check that there is only a finite number of proper trees with leaves chosen from a finite set.

A *bow tie* is a DAG G that does not contain the graph from Figure 5 as a minor and such that there is a node o such that every node in G either can reach o or is reachable from o . By applying the SP-rewriting system to a bow tie, we get a bottom-up tree concatenated with a top-down tree, as illustrated in figure Figure 7 (hence the name). It follows easily that G is an SP-slice if and only if it is a disjoint union of bow ties.

³This is sometimes used as the definition of an SP-graph.

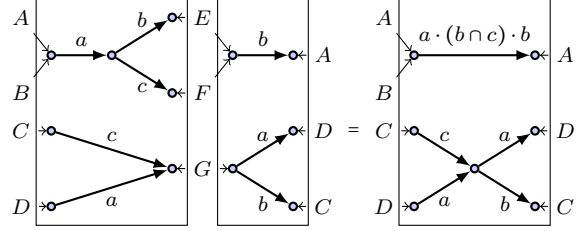


Figure 8: Composition of boxes

3.2 Categories of boxes

Definition 10 (Box)

A *box* from A to B labelled with a label space \mathbb{L} is a tuple $\beta = \langle G, \underline{p}, \bar{p} \rangle$ where G is a DAG labelled with \mathbb{L} , $\underline{p} : A \rightarrow \min G$ is a surjective map and $\bar{p} : B \rightarrow \max G$ is a bijective map, respectively indicating the input ports and output ports. We write $\beta : A \rightarrow B$ if β is a box from A to B . *

Examples of such boxes are given in Figures 10 and 11.

Definition 11 (Type)

A *type* over a set A is a structure $\tau = \langle G_\tau, l_\tau \rangle$ such that G_τ is a proper top-down unlabelled tree and l_τ is a bijection from A to $\max G_\tau$. The set of types over subsets of a finite set P is written \mathbb{T}_P . *

Examples of such types are given in Figure 9.

Remark 12. In the sequel we reason about boxes and types modulo renaming of nodes.

A type τ over A may be seen as a box from $\{\iota\}$ to A labelled with $\mathbb{1}$: we can build the box from $\{\iota\}$ to A $\boxed{\tau} := \langle V_\tau, E, [\iota \mapsto r_\tau], l_\tau \rangle$, with $E = \{(x, 1, y) \mid (x, y) \in E_\tau\}$. Conversely, we may forget about the label information in a box: the erasing of a box β is the box $|\beta|$ where all labels are replaced by 1. It is quite clear that for any type τ , $\boxed{\tau} = |\boxed{\tau}|$.

One of the most interesting features of boxes is their ability to compose. Let $\beta = \langle G_1, \underline{p}_1, \bar{p}_1 \rangle : A \rightarrow B$ and $\gamma = \langle G_2, \underline{p}_2, \bar{p}_2 \rangle : B \rightarrow C$. We define the relation \sim over $G_1 \cup G_2$ as the smallest equivalence relation containing all pairs x, y such that there is some $b \in B$ such that $\bar{p}_1(b) = x$ and $\underline{p}_2(b) = y$. The composition of β and γ , written $\beta \odot \gamma$, is defined as $\langle G \downarrow, \underline{p}, \bar{p} \rangle : A \rightarrow C$, where:

- G is the quotient of the disjoint union of G_1 and G_2 by the relation \sim ;
- the function \underline{p} (respectively \bar{p}) associates to $x \in A$ (resp. $\in B$) the equivalence class of $\underline{p}_1(x)$ (resp. $\bar{p}_2(x)$).

It can be checked that this composition operation is associative and that for any two boxes $\beta : A \rightarrow B$ and $\gamma : B \rightarrow C$, the following holds:

$$|\beta \odot \gamma| = |\beta| \odot |\gamma|. \quad (1)$$

This operation is illustrated in Figure 8.

Definition 13 (Typed boxes)

Let $\beta : A \rightarrow B$ be a box over \mathbb{L} , σ and τ be types respectively over A and B . β has the type $\sigma \rightarrow \tau$ if the following holds:

$$|\boxed{\sigma} \odot \beta| = \boxed{\tau}. \quad (2)$$

We write $\mathfrak{B}_{\sigma \rightarrow \tau}^{\mathbb{L}}$ for the set of boxes over \mathbb{L} of type $\sigma \rightarrow \tau$, and $\mathbb{B}_P^{\mathbb{L}}$ for the set of typed boxes over \mathbb{L} with types from \mathbb{T}_P .

These boxes allow us to define the category $\mathbb{B}^{\mathbb{L}}$ of types and typed boxes over \mathbb{L} where:

- the objects are types;




Figure 9: Some types σ and τ .




Figure 10: A box in $\mathfrak{B}_{\sigma \rightarrow \tau}^{\mathbb{W}_X}$.

Figure 11: The identity box on type σ .

- for any two objects σ, τ , the set of morphisms from σ to τ is $\mathfrak{B}_{\sigma \rightarrow \tau}^{\mathbb{L}}$.

*

(Notice that the type of a box is not unique: a single box may have multiple types.)

For instance the box in Figure 10 has type $\sigma \rightarrow \tau$, with σ and τ as defined in Figure 9.

Let us verify that $\mathbb{B}^{\mathbb{L}}$ satisfies the axioms of a category. We already know that the composition of boxes is associative, but it remains to be shown that the composition of typed boxes is a typed box. In other words: $\mathfrak{B}_{\sigma \rightarrow \tau}^{\mathbb{L}} \circ \mathfrak{B}_{\tau \rightarrow \theta}^{\mathbb{L}} \subseteq \mathfrak{B}_{\sigma \rightarrow \theta}^{\mathbb{L}}$.


Let $\beta, \gamma \in \mathfrak{B}_{\sigma \rightarrow \tau}^{\mathbb{L}} \times \mathfrak{B}_{\tau \rightarrow \theta}^{\mathbb{L}}$. We need to check that $\beta \circ \gamma \in \mathfrak{B}_{\sigma \rightarrow \theta}^{\mathbb{L}}$.

$$\begin{aligned}
 \boxed{\sigma} \circ (\beta \circ \gamma) &= [(\boxed{\sigma} \circ \beta) \circ \gamma] && \text{(Associativity)} \\
 &= \boxed{\sigma} \circ \beta \circ \gamma && \text{(Equation (1))} \\
 &= \boxed{\tau} \circ \gamma && (\beta \in \mathfrak{B}_{\sigma \rightarrow \tau}^{\mathbb{L}}) \\
 &= \boxed{\tau} \circ \gamma && \text{(Equation (1))} \\
 &= \boxed{\theta}. && (\gamma \in \mathfrak{B}_{\tau \rightarrow \theta}^{\mathbb{L}})
 \end{aligned}$$

The last thing we need is the identity over an arbitrary object τ . If L is the set of leaves of τ , then we define $\text{id}_{\tau} = \langle \langle L, \emptyset \rangle, l_{\tau}, l_{\tau} \rangle$. It is immediate to check that $\text{id}_{\tau} \in \mathfrak{B}_{\tau \rightarrow \tau}^{\mathbb{L}}$ and that for any typed boxes β, γ of appropriate types, $\beta \circ \text{id}_{\tau} = \beta$ and $\text{id}_{\tau} \circ \gamma = \gamma$. The identity box on type τ is displayed in Figure 11.

3.3 Boxes over expressions

Consider boxes labelled with a label space \mathbb{E}_X . The elements of \mathbb{E}_X are expressions taken from $\text{Reg}_X^{\mathbb{L}}$, and thus denote languages of SP-graphs, as defined earlier. This means that we can generate from such a box a set of boxes labelled with \mathbb{W}_X , by replacing each edge labelled with e by a graph $G \in \mathcal{G}(e) \subseteq \mathbb{W}_X^4$. For instance, the box:



⁴As mentioned before, we make no distinction between SP-graphs and terms built using letters and the operators \cdot and \cup .





Figure 12: Construction of $\beta_{t,C}$

can generate all boxes of the following shapes, for $n > 0$:



It is useful to notice that for any box $B \in \mathfrak{B}_{\sigma \rightarrow \tau}^{\mathbb{E}_X}$, if B generates the box β , then $\beta \in \mathfrak{B}_{\sigma \rightarrow \tau}^{\mathbb{W}_X}$. In the following, we will denote by $\mathfrak{B}(B)$ the set of boxes generated by a box B labelled with \mathbb{E}_X .

4. Automata and regular sets of boxes

4.1 A regular expression over boxes

It is then quite easy to generate from an SP-automaton \mathcal{A} a regular expression over boxes $e_{\mathcal{A}} \in \text{Reg}(\mathbb{B}_P^{\mathbb{W}_X})$ such that:

$$\beta_1 \cdots \beta_n \in \llbracket e_{\mathcal{A}} \rrbracket \Leftrightarrow \begin{cases} \beta_1 \circ \cdots \circ \beta_n = \langle G, \underline{p}, \underline{p} \rangle \text{ with } G \in \mathcal{G}(\mathcal{A}) \\ \beta_1 \circ \cdots \circ \beta_n \text{ has type } \bullet \rightarrow \bullet \xrightarrow{\iota} \bullet \rightarrow \bullet \circ f. \end{cases}$$

Let us fix an SP-automaton $\mathcal{A} = \langle P, \mathcal{F}, \iota, f \rangle$ over an alphabet Σ . Its transitions are thus labelled with Σ_{\bullet} . Given automata configurations C and C' , and a transition t such that $C \xrightarrow{t} C'$, we can define a box $\beta_{t,C} : C \rightarrow C'$ as follows. Suppose $\bar{t} = \{(a_1, q_1), \dots, (a_m, q_m)\}$ and $C \setminus \bar{t} = \{p_1, \dots, p_n\}$. Recall that because t is a valid transition from C to C' , the sets $\{q_1, \dots, q_m\}$ and $\{p_1, \dots, p_n\}$ are disjoint and $C' = \{p_1, \dots, p_n, q_1, \dots, q_m\}$. Thus $\beta_{t,C}$ is defined as $\langle V, E, \underline{p}, \underline{p} \rangle$ where:

- V contains an is the input node x_0 , a node x_i for every place p_i in $C \setminus \bar{t}$ and a node y_i for every output q_i .
- \underline{p} maps every input place p to x_0 , every every other p_i in $C \setminus \bar{t}$ to x_i .
- \underline{p} maps every output q_i to y_i , and the remaining p_i to x_i .
- for each output (a_i, q_i) , we place an edge (x_0, a_i, y_i) in E .

This construction is illustrated in Figure 12.

We may extend this construction to runs: for every run of length n $\xi = \langle (\xi_k)_{0 \leq k \leq n}, (t_k)_{0 \leq k < n} \rangle$ we define $\beta(\xi)$ by:

$$\beta(\xi) := \beta_{t_0, \xi_0} \circ \cdots \circ \beta_{t_{n-1}, \xi_{n-1}}.$$

It is quite straight-forward to check that for every accepting run ξ ,

$$\beta(\xi) = \iota \rightarrow \mathcal{G}(\xi) \rightarrow f$$

Thus there is no difference between the graphs produced by the automaton and the boxes corresponding to its accepting runs. We call a box *valid* if there is an accepting run ξ with a consecutive sub-run $\xi[i..j]$ such that $\beta = \beta(\xi[i..j])$. We say in that case that the run $\xi[i..j]$ *witnesses* the validity of β . Notice that because \mathcal{A} is an SP-automaton, the graph of a valid box has to be an SP-slice.

The types help us here, as illustrated by the following lemma:

Lemma 14. *If β is a valid box with type $\sigma \rightarrow \tau$ and γ is valid with type $\tau \rightarrow \nu$, then $\beta \odot \gamma$ is a valid box with type $\sigma \rightarrow \nu$.*

We extend this to words and regular expressions over typed boxes.

Definition 15 (Valid expression)

An expression $e \in \text{Reg}(\mathfrak{B}^{\mathbb{W}_{\Sigma^{\bullet}}})$ is valid if for all $\beta_1 \cdots \beta_n \in \llbracket e \rrbracket$, the composition $\beta_1 \odot \cdots \odot \beta_n$ is well defined, and the resulting box is valid. *

Given a word $w = \beta_1 \cdots \beta_n$, we call the box $\beta_1 \odot \cdots \odot \beta_n$ the *compilation* of w .

Definition 16 (Box language of e)

If $e \in \text{Reg}(\mathfrak{B}^{\mathbb{W}_{\Sigma^{\bullet}}})$ is valid, then the box language of e , written $\llbracket e \rrbracket_{\odot}$, is the set of boxes obtained by compiling the words of $\llbracket e \rrbracket$. *

Definition 17 (Typed expression)

A valid expression $e \in \text{Reg}(\mathbb{B}_P^{\mathbb{W}_{\Sigma^{\bullet}}})$ has type $\sigma \rightarrow \tau$ if for all $\beta \in \llbracket e \rrbracket_{\odot}$, β has the type $\sigma \rightarrow \tau$. *

We may now define a deterministic finite-state automaton as follows:

- the set of states is \mathbb{T}_P , the types over the set P of places of \mathcal{A} (as remarked before, this set is finite);
- for every pair of types σ, τ , if C is the set of leaves of σ , if $t \in \mathcal{T}$ can be fired from C and if $\beta_{t,C} \in \mathfrak{B}_{\sigma \rightarrow \tau}^{\mathbb{W}_{\Sigma^{\bullet}}}$ we add a transition $(\sigma, \beta_{t,C}, \tau)$;
- the initial state is the type $\circ \rightarrow \circ \iota$;
- the final state is the type $\circ \rightarrow \circ f$.

Remark 18. This automaton can be computed partially, by enumerating all runs starting from $\{\iota\}$, and computing the corresponding boxes and types. Because of the finite number of states in this automaton this procedure must terminate. If the automaton we were given is not SP, then the computation will reach a step where we have a reachable type σ with leaves C and a transition t enabled at C such that $\llbracket \sigma \rrbracket_{\odot} \beta_{t,C}$ does not reduce to a tree. This gives a decision procedure to check whether \mathcal{A} is indeed an SP-automaton. Using the classic Kleene theorem, we can then compile this automaton into a regular expression $e_{\mathcal{A}} \in \text{Reg}(\mathbb{B}_P^{\mathbb{W}_{\Sigma^{\bullet}}})$ such that:

1. $e_{\mathcal{A}}$ and all its sub-expressions are valid and typed, as defined earlier, meaning that the words in their languages compile to valid boxes and that two boxes in the language of the same sub-expression share a common type⁵;
2. if ξ is an accepting run of the Petri automaton \mathcal{A} , then $\beta(\xi)$ is the compilation of some word in $\llbracket e \rrbracket_{\mathcal{A}}$. Thus $\llbracket e_{\mathcal{A}} \rrbracket_{\odot}$ is the set of all valid boxes of type $\circ \rightarrow \circ \iota \rightarrow \circ \rightarrow \circ f$.

This allow us to state an important property of the expression $e_{\mathcal{A}}$:

Lemma 19. *Let $f, g \in \text{Reg}(\mathbb{B}_P^{\mathbb{W}_{\Sigma^{\bullet}}})$ be valid expressions of type $\sigma \rightarrow \tau$, such that $\llbracket f \rrbracket \subseteq \llbracket g \rrbracket$. If $e_{\mathcal{A}}[f/g]$ is the expression obtained by replacing an occurrence of f with g in $e_{\mathcal{A}}$, then the box language of $e_{\mathcal{A}}$ is equal to that of $e_{\mathcal{A}}[f/g]$.*

Proof. We proceed by mutual inclusion. One direction is straight forward: all regular operations being increasing, we know that $\llbracket f \rrbracket \subseteq \llbracket g \rrbracket$ entails $\llbracket e_{\mathcal{A}} \rrbracket \subseteq \llbracket e_{\mathcal{A}}[f/g] \rrbracket$, thus ensuring $\llbracket e_{\mathcal{A}} \rrbracket_{\odot} \subseteq \llbracket e_{\mathcal{A}}[f/g] \rrbracket_{\odot}$.

⁵This property follows from results about typed Kleene Algebra. See for instance [4, 5, 7].

The property 2 allows us to check the other direction. The typing ensures that for any words u, v, w and w' , if w and w' compile to boxes with the same type then $u \cdot w \cdot v$ is valid if and only if $u \cdot w' \cdot v$ is valid as well, and they have the same type. Thus we know that $\llbracket e_{\mathcal{A}}[f/g] \rrbracket_{\odot}$ only contain valid boxes of the same type as the ones from $\llbracket e_{\mathcal{A}} \rrbracket_{\odot}$. As $\llbracket e_{\mathcal{A}} \rrbracket_{\odot}$ is maximal for its type, it must contain $\llbracket e_{\mathcal{A}}[f/g] \rrbracket_{\odot}$. \square

4.2 Compiling an expression over boxes into $\text{Reg}_{\Sigma}^{\circ}$

What remains to do is to transform this regular expression over boxes into an expression from $\text{Reg}_{\Sigma}^{\circ}$. We do so by moving from typed boxes labelled with $\mathbb{W}_{\Sigma^{\bullet}}$ to finite sets of boxes sharing a common type and labelled with $\mathbb{E}_{\Sigma^{\bullet}} \supseteq \mathbb{W}_{\Sigma^{\bullet}}$, thus internalising the operators. By mapping each box β to $\{\beta\}$, we may translate this way the expression $e_{\mathcal{A}}$ obtained in the previous section to a new expression $e_{\mathcal{A}}$. The operator \odot may be lifted here in a type preserving way, as a pointwise application:

$$S \odot S' := \{B \odot B' \mid B \in S, B' \in S'\}.$$

We also extend the notion of boxes over $\mathbb{W}_{\Sigma^{\bullet}}$ generated by a box over $\mathbb{E}_{\Sigma^{\bullet}}$ to a set S simply by stating that $\mathfrak{B}(S) := \bigcup_{B \in S} \mathfrak{B}(B)$.

Definition 20 (Valid expression)

An expression $e \in \text{Reg}(\mathcal{P}_f(\mathbb{B}_P^{\mathbb{E}_{\Sigma^{\bullet}}}))$ is valid if all $S_1 \cdots S_n \in \llbracket e \rrbracket$, the composition $S_1 \odot \cdots \odot S_n = S$ is defined and $\mathfrak{B}(S)$ only contains valid boxes. *

The definition of $\llbracket _ \rrbracket_{\odot}$ we had for valid expressions with letters in $\mathbb{B}_P^{\mathbb{W}_{\Sigma^{\bullet}}}$ can thus be salvaged for expressions with letters in $\mathcal{P}_f(\mathbb{B}_P^{\mathbb{E}_{\Sigma^{\bullet}}})$. Lemma 14 can be adapted to this setting:

Lemma 21. *Let S_1 and S_2 be finite sets of boxes over $\mathbb{E}_{\Sigma^{\bullet}}$, such that*

1. $\forall \beta \in \mathfrak{B}(S), \beta$ is a valid box with type $\sigma \rightarrow \tau$;
2. $\forall \gamma \in \mathfrak{B}(S'), \gamma$ is a valid box with type $\tau \rightarrow \nu$;

then $\forall \beta \in \mathfrak{B}(S \odot S'), \beta$ is a valid box with type $\sigma \rightarrow \nu$.

A regular expression e with letters in $\mathcal{P}_f(\mathbb{B}_P^{\mathbb{E}_{\Sigma^{\bullet}}})$ represent the set of boxes

$$\llbracket e \rrbracket := \bigcup_{S \in \llbracket e \rrbracket_{\odot}} \mathfrak{B}(S).$$

We extend a definition from the previous section to this setting:

Definition 22 (Typed expression)

An expression valid $e \in \text{Reg}(\mathcal{P}_f(\mathbb{B}_P^{\mathbb{E}_{\Sigma^{\bullet}}}))$ has type $\sigma \rightarrow \tau$ if all $\beta \in \llbracket e \rrbracket$ have the type $\sigma \rightarrow \tau$. *

We can now equate two expressions if they denote the same set of boxes, introducing the equivalence relation \equiv_{β} :

$$e \equiv_{\beta} f := (\llbracket e \rrbracket = \llbracket f \rrbracket)$$

Notice that Lemma 19 can be restated for $e_{\mathcal{A}}$:

Lemma 23. *Let $f, g \in \text{Reg}(\mathcal{P}_f(\mathfrak{B}^{\mathbb{E}_{\Sigma^{\bullet}}}))$ be valid expressions of type $\sigma \rightarrow \tau$, such that $\llbracket f \rrbracket_{\odot} \subseteq \llbracket g \rrbracket_{\odot}$. Then $e_{\mathcal{A}} \equiv_{\beta} e_{\mathcal{A}}[g/f]$.*

The operator \cup and \cdot are easy enough to define on finite sets of boxes labelled over $\mathbb{E}_{\Sigma^{\bullet}}$:

$$S \cdot S' \equiv_{\beta} \bigcup_{B \in S} \bigcup_{B' \in S'} B \odot B'$$

$$S \cup S' \equiv_{\beta} \{B \mid B \in S \cup S'\}$$

The case of $_*$ however is much more interesting. In fact, this is the main technical difficulty of the present paper, and where most of the machinery we have introduced up to this point will prove

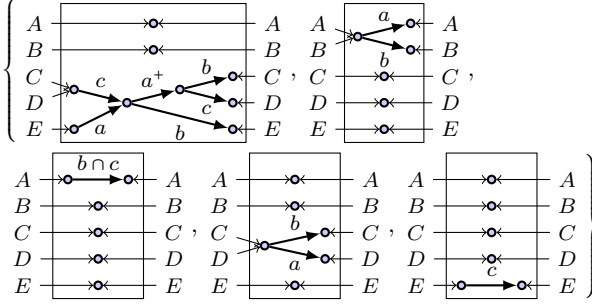


Figure 13: Elementary decomposition of $\{\beta, \gamma\}$.

useful. For typing reasons, notice that we only need to compute the star of a set of boxes of an identity type $\sigma \rightarrow \sigma$. In the next section, we show how to compute from a set S of boxes labelled over \mathbb{E}_Σ , and having type $\sigma \rightarrow \sigma$ a new finite set of boxes S' such that:

1. $\langle S^* \rangle \subseteq \langle S' \rangle$
2. and $\forall \beta \in \langle S' \rangle$, β is valid and has type $\sigma \rightarrow \sigma$.

By Lemma 23 this entails the following equivalence:

$$e_{\mathcal{A}} \equiv_{\beta} e_{\mathcal{A}} [S' / (S^*)].$$

Using these transformation, we now can “compile” the expression $e_{\mathcal{A}}$ into a finite set of boxes of type $\circ \rightarrow \circ \rightarrow \circ \rightarrow \circ \rightarrow f$. Such a set may only have the shape:

$$\left\{ \iota \rightarrow \begin{array}{|c|} \hline e_1 \\ \hline \end{array} \rightarrow f, \dots, \iota \rightarrow \begin{array}{|c|} \hline e_n \\ \hline \end{array} \rightarrow f \right\}.$$

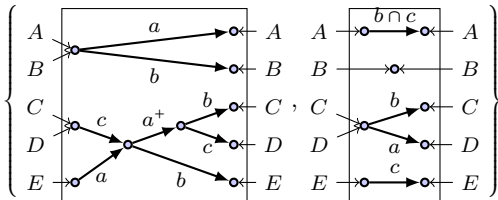
It is then straightforward to see that $\langle e_{\mathcal{A}} \rangle$ is isomorphic to $\mathcal{G}(e_1 \cup \dots \cup e_n)$. If we call $\mathcal{E}(\mathcal{A})$ the expression $e_1 \cup \dots \cup e_n$, the following theorem holds:

Theorem 24 (Main result). *For every SP-automaton \mathcal{A} over Σ , there exists an expression $\mathcal{E}(\mathcal{A}) \in \text{Reg}_{\Sigma}^{\circ}$ such that*

$$\mathcal{G}(\mathcal{A}) = \mathcal{G}(\mathcal{E}(\mathcal{A})).$$

4.3 Staring boxes

Instead of giving a formal description, which would be very involved, we describe the procedure on an example, outlining the important computation steps. Consider the following set of boxes, respectively named β and γ .



Both these boxes have type $\sigma \rightarrow \sigma$, with the σ as shown in Figure 9. First, we decompose each box, by separating disjoint non-trivial connected components. The idea behind this step is that the different connected component correspond to concurrent processes, and as such may be iterated independently. This yields the set of five so-called *elementary* boxes displayed in Figure 13, respectively named η_1 through η_5 .

One can show that the resulting boxes are always valid and have the type $\sigma \rightarrow \sigma$ (a formal proof of the properties of this decomposition is given in Appendix A). Because $\beta = \eta_1 \odot \eta_2$ and $\gamma = \eta_3 \odot \eta_4 \odot \eta_5$, we get that:

$$\langle (\beta \cup \gamma)^* \rangle_{\circ} \subseteq \langle (\eta_1 \cup \eta_2 \cup \eta_3 \cup \eta_4 \cup \eta_5)^* \rangle_{\circ}$$

box	support
η_1	$\{C, D, E\}$
η_2	$\{A, B\}$
η_3	$\{A\}$
η_4	$\{C, D\}$
η_5	$\{E\}$

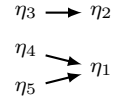


Figure 15: Graph of \leq

Figure 14: Supports of the η_i

On the other hand, because of Lemma 21 we can establish that the set $\langle (\eta_1 \cup \eta_2 \cup \eta_3 \cup \eta_4 \cup \eta_5)^* \rangle$ only contains valid boxes of type $\sigma \rightarrow \sigma$.

By Lemma 23, we know that we may replace in the expression $e_{\mathcal{A}}$ the sub expression $(\beta \cup \gamma)^*$ by $(\eta_1 \cup \eta_2 \cup \eta_3 \cup \eta_4 \cup \eta_5)^*$ without changing the set of produced graphs. This means that we may continue the computation with the set of elementary boxes $H = \{\eta_1, \eta_2, \eta_3, \eta_4, \eta_5\}$. We define the *support* of a box with a single non-trivial connected component as the set of ports that are mapped inside this component. Figure 14 lists the supports of the boxes η_i . Notice that for every two boxes η, η' in H , their supports are either included one in the other, or of empty intersection. This can be shown formally by realising that for a box with a single connected component of type $\sigma \rightarrow \sigma$, the support is the set of leaves of a sub-tree of σ (i.e. the set of all leaves reachable from some node of σ). This means we can endow the set $\{\eta_1, \eta_2, \eta_3, \eta_4, \eta_5\}$ with a pre-order relation \leq defined by $\eta \leq \eta'$ if $\text{support}(\eta) \subseteq \text{support}(\eta')$. The graph of the order on our example set is given in Figure 15.

The order has an interesting property: if η and η' are incompatible, then $\eta \odot \eta' = \eta' \odot \eta$. Thus $(\eta \cup \eta')^* \equiv_{\beta} \eta^* \cdot \eta'^*$. With that in mind we will compute the expression:

$$(\eta_3^* \cup \eta_2^*)^* \cdot ((\eta_4^* \cdot \eta_5^*) \cup \eta_1^*)^*$$

The first step in this computation is the computation of the star of a single box. Notice that we may define easily the star of a “linear” elementary box, i.e. of the shape:

$$\left(\begin{array}{|c|} \hline e \\ \hline \end{array} \right)^* \equiv_{\beta} \left(\begin{array}{|c|} \hline e^+ \\ \hline \end{array} \right)^*$$

This construction gives us η_3^* and η_5^* . Now for general elementary boxes, notice that the non-trivial connected component of an elementary box is always a bow tie. Hence we can always decompose such a box η of type $\sigma \rightarrow \sigma$ as $\delta_1 \odot \delta_2$ such that the non-trivial connected component of δ_1 is a bottom-up tree, and that of δ_2 is a top-down tree, with the middle node o of the bow tie being the root of both trees. See Figure 16 for examples. This means that $\delta_2 \odot \delta_1$ is a linear elementary box. Furthermore, the definition of $\delta_{1,2}$ ensures $\{\eta\}^* \equiv_{\beta} (\{\delta_1\} \cdot \{\delta_2\})^*$. The regular laws entail that for every expressions e and f we have $\langle (e \cdot f)^* \rangle = \langle [1 \cup e \cdot (f \cdot e)^* \cdot f] \rangle$. Thus we can use the previous construction to get a set of boxes representing $(\{\delta_2\} \cdot \{\delta_1\})^* \equiv_{\beta} \{\delta_2 \odot \delta_1\}^*$, and then compute:

$$\{\eta\}^* \equiv_{\beta} \{\text{id}_{\sigma}\} \cup \{\delta_1\} \cdot \{\delta_2 \odot \delta_1\}^* \cdot \{\delta_2\}.$$

The last thing to do is to compute stars of the shape $(S^* \cup \eta)^*$ where:

- S^* is a finite set of boxes computed in an earlier step,
- $\text{support}(\eta)$ contains the “support” of S^* : the set of inputs mapped to a non trivial connected component in some $B \in S^*$.

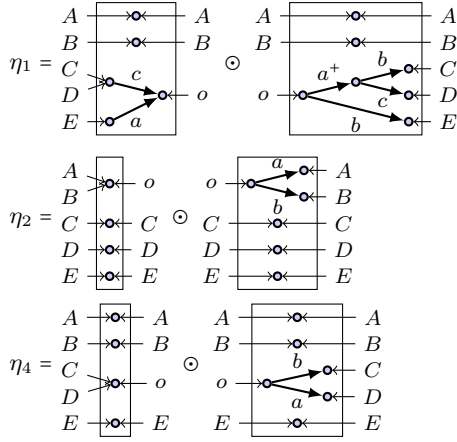
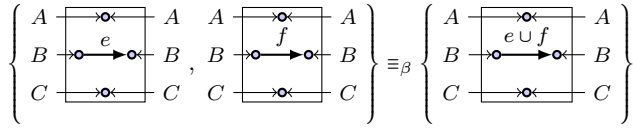


Figure 16: Decomposition as a product of trees of η_1 , η_2 and η_4 .

Again, we decompose η into δ_1 and δ_2 as in the previous construction, and notice that $\delta_2 \odot S^* \odot \delta_1$ is set of a linear boxes, along the same line. We may thus bundle them together in a single box with a line labelled with the union of their labels. For example:



This allow us to produce a set of (two) boxes representing the star of the box $(\delta_2 \odot S^* \odot \delta_1)$. The regular laws allow us to conclude, by stating that:

$$(S^* \cup \delta_1 \cdot \delta_2)^* \equiv_{\beta} S^* \cup \delta_1 \odot (\delta_2 \odot S^* \odot \delta_1)^* \odot \delta_2$$

We can now give the bulk of the computation, in Figure 17. The result is then obtained by performing the composition

$$(\eta_3^* \cup \eta_2)^* \odot ((\eta_4^* \cup \eta_5)^* \cup \eta_1)^*,$$

thus yielding 32 boxes: 8 for $((\eta_4^* \cup \eta_5)^* \cup \eta_1)^*$ times 4 for $(\eta_3^* \cup \eta_2)^*$.

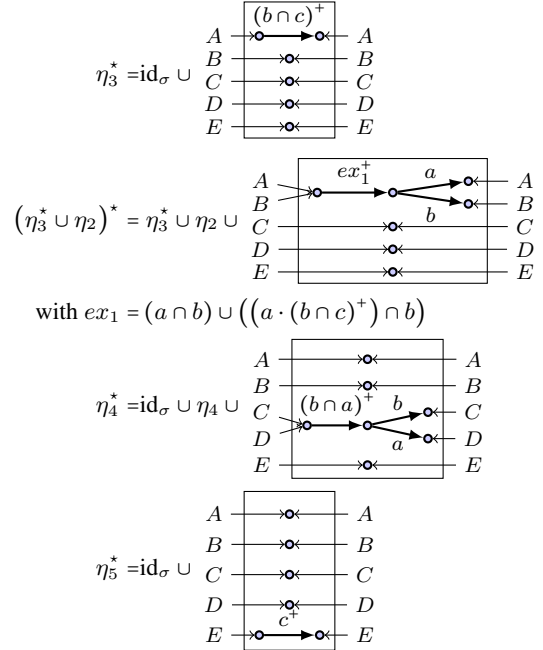
5. Consequences

5.1 Kleene Allegories

The main motivation for this work was the study of the equational theory of Kleene Allegories (KAl). In this line of work we consider binary relations and the operations of union (\cup), intersection (\cap), composition (\cdot), converse ($\bar{}$), transitive closure ($\bar{}^+$), and the constants identity (1) and empty relation (0). Terms made from those operations and some variables $a, b, \dots \in \Sigma$ constitute the set $\text{Reg}_{\Sigma}^{\bar{}}$. A pair $e, f \in \text{Reg}_{\Sigma}^{\bar{}}$ is then *relationally equivalent*, denoted $\text{Rel} \models e = f$, if the corresponding equality holds universally when the expressions are interpreted as binary relations. Notice that we may restrict expressions so that the operator $\bar{}$ is only applied to variables, by using the following rewriting system:

$$\begin{aligned} (a \cup b)^{\bar{}} &\rightarrow a^{\bar{}} \cup b^{\bar{}} & 0^{\bar{}} &\rightarrow 0 & (a^+)^{\bar{}} &\rightarrow (a^{\bar{}})^+ \\ (a \cdot b)^{\bar{}} &\rightarrow b^{\bar{}} \cdot a^{\bar{}} & 1^{\bar{}} &\rightarrow 1 & a^{\bar{}} &\rightarrow a \\ (a \cap b)^{\bar{}} &\rightarrow a^{\bar{}} \cap b^{\bar{}}. \end{aligned}$$

The result of this transformation on an expression e is an expression e' always satisfying $\text{Rel} \models e = e'$. We may associate to these terms graph languages, as in Definition 2.



$$\text{with } ex_1 = (a \cap b) \cup ((a \cdot (bnc)^+) \cap b)$$

$$((\eta_4 \cup \eta_5)^* \cup \eta_1)^* = \eta_4^* \odot \eta_5^* \cup \eta_1 \cup$$

$$\begin{aligned} \text{with } ex_2 &= (a^+ \cdot (bnc) \cdot (bna)^+ \cdot (bna) \cdot c) \cap (b \cdot c^+ \cdot a) \\ &\cup (a^+ \cdot (bnc) \cdot (bna)^+ \cdot (bna) \cdot c) \cap (b \cdot a) \\ &\cup (a^+ \cdot (bnc) \cdot (bna) \cdot c) \cap (b \cdot c^+ \cdot a) \\ &\cup (a^+ \cdot (bnc) \cdot (bna) \cdot c) \cap (b \cdot a) \\ &\cup (a^+ \cdot (bnc) \cdot c) \cap (b \cdot c^+ \cdot a) \end{aligned}$$

Figure 17: Computation steps

Definition 25 (Graph language of a KAl expression: $\mathcal{G}'(e)$)

The graph language is defined by structural induction as follows:

$$\begin{aligned} \mathcal{G}'(a) &:= \left\{ \begin{array}{c} \longrightarrow \circ \xrightarrow{a} \circ \longrightarrow \end{array} \right\}; & \mathcal{G}'(a^{\bar{}}) &:= \left\{ \begin{array}{c} \longrightarrow \circ \xleftarrow{a} \circ \longrightarrow \end{array} \right\}; \\ \mathcal{G}'(0) &:= \emptyset; & \mathcal{G}'(1) &:= \left\{ \begin{array}{c} \longrightarrow \circ \longrightarrow \end{array} \right\}; \\ \mathcal{G}'(e \cup f) &:= \mathcal{G}'(e) \cup \mathcal{G}'(f); \\ \mathcal{G}'(e \cdot f) &:= \{E \cdot F \mid E \in \mathcal{G}'(e) \text{ and } F \in \mathcal{G}'(f)\}; \\ \mathcal{G}'(e \cap f) &:= \{E \cap F \mid E \in \mathcal{G}'(e) \text{ and } F \in \mathcal{G}'(f)\}; \\ \mathcal{G}'(e^+) &:= \{E_1 \cdot \dots \cdot E_n \mid n > 0, \forall i, E_i \in \mathcal{G}'(e)\}. \end{aligned}$$

In [2], we showed the following result:

$$\text{Rel} \models e = f \Leftrightarrow \mathcal{G}'(e) = \mathcal{G}'(f).^6$$

⁶We need not concern ourselves with the definition of $\bar{}$ here, as it plays no role in the present result.

Notice that the only two differences between $\mathcal{G}(_)$ and $\mathcal{G}'(_)$ are the construction for 1 and the case of $_^\sim$. We can mend that gap, by computing $\mathcal{G}'(e)$ from $\mathcal{G}(e)$. The first step is to consider the sub-expressions a^\sim as letters, i.e. to see $e \in \text{Reg}_{\Sigma}^{\sim}$ as an expression in $\text{Reg}_{\mathfrak{S}}^{\sim}$ with $\mathfrak{S} := \Sigma \cup \{a^\sim \mid a \in \Sigma\}$. This means we can apply $\mathcal{G}(_)$ on e , and produce a set of graphs labelled with \mathfrak{S} . We then introduce the graph transformation Φ :

Definition 26 (Φ)

Let $G = \langle V, E, \iota, o \rangle$ be a graph labelled with \mathfrak{S} . Let \equiv_G be the smallest equivalence relation on V containing all pairs (i, j) such that $(i, 1, j) \in E$. Then $\Phi(G)$ is the graph defined by

$$\begin{aligned} \Phi(G) &:= \{ \{ [i]_G \mid i \in V \}, E', [\iota]_G, [o]_G \} \\ [i]_G &:= \{ k \in V \mid i \equiv_G k \} \\ E' &:= \left\{ \left([i]_G, x, [j]_G \right) \left| \begin{array}{l} x \in \Sigma \text{ and} \\ \exists k \in [i]_G, l \in [j]_G : \\ (k, x, l) \in E \text{ or} \\ (l, x^\sim, k) \in E \end{array} \right. \right\} \quad * \end{aligned}$$

It is then a simple matter to check that

$$\forall e \in \text{Reg}_{\Sigma}^{\sim}, \mathcal{G}'(e) = \{ \Phi(G) \mid G \in \mathcal{G}(e) \}.$$

In the same paper we also introduced Petri automata, but again we used a different definition of the language $\mathcal{L}(\mathcal{A})$ of an automaton \mathcal{A} . However, this set can be expressed with our notations as $\mathcal{L}(\mathcal{A}) = \mathbf{*}\Phi(\mathcal{G}(\mathcal{A}))$. A function $\mathcal{A}(_)$ was provided, mapping an expression $e \in \text{Reg}_{\Sigma}^{\sim}$ to a Petri automaton $\mathcal{A}(e)$ such that $\mathcal{L}(\mathcal{A}(e)) = \mathcal{G}(e)$.⁷ Thus:

$$\mathcal{L}(\mathcal{A}(e)) = \mathbf{*}\Phi(\mathcal{G}(e)).$$

Theorem 24 allow us to relate the language of any SP-automaton \mathcal{A} to the graph language of $\mathcal{E}(\mathcal{A})$:

$$\mathcal{L}(\mathcal{A}) = \mathbf{*}\Phi(\mathcal{G}(\mathcal{E}(\mathcal{A}))).$$

This proves that from a decidability stand point, language equivalence of Petri automata and relational equivalence of KAI expressions. Formally:

$$\forall e, f, \text{Rel} \models e = f \Leftrightarrow \mathcal{L}(\mathcal{A}(e)) = \mathcal{L}(\mathcal{A}(f)) \quad (3)$$

$$\forall \mathcal{A}, \mathcal{B}, \mathcal{L}(\mathcal{A}) = \mathcal{L}(\mathcal{B}) \Leftrightarrow \text{Rel} \models \mathcal{E}(\mathcal{A}) = \mathcal{E}(\mathcal{B}) \quad (4)$$

One could also read this result from a graph theoretical viewpoint. What we show here is that the sets of SP-graphs one may express trough regular expressions enhanced with a parallel operator \cap are exactly those generated by SP-automata.

5.2 Petri nets

The construction presented here can also provide a finite regular expression with intersection to represent all finite occurrence nets of a given Petri net N , provided this net is safe and Parallel Series. In terms of concurrent processes, these two conditions roughly amount to requiring that there is a bounded number of processes running at any given time, and that communication between processes is synchronous.

More precisely, suppose we are given a net N , with a finite set P of places and a finite set $T \subseteq \mathcal{P}(P) \times \mathcal{P}(\Sigma \times P)$ of transitions.⁸ If $I, F \subseteq \mathcal{P}(P)$ are respectively a source set and a target set, we

⁷The automaton $\mathcal{A}(e)$ always satisfied the Parallel Series condition we required in Section 2.3.

⁸Here we stick to the setting of this paper by labelling the outputs of transitions rather than the transitions themselves. However, one can start with a net with usual labelled transition, and then simply copy the label of t to each of its outputs. This would lead to graphs where each node would have the same label on each outgoing edge, thus effectively labelling the node rather than the edge.

can produce a Petri automaton \mathcal{A} such that the set of occurrence nets of process going from a source configuration in I to a target configuration in F is isomorphic to $\mathcal{G}(\mathcal{A})$. To achieve this, we add to the net:

- two new places ι and f and a fresh label $\#$;
- for each source configuration $\{p_1, \dots, p_n\} \in I$ a transition $\{\{\iota\}, \{\#\}, \{p_1\}, \dots, \{\#\}, \{p_n\}\}$;
- for each target configuration $\{q_1, \dots, q_n\} \in F$ a transition $\{\{q_1, \dots, q_n\}, \{\#\}, \{f\}\}$.

It is then a simple matter of unfolding definitions to see that $\mathcal{G}(\mathcal{A})$ indeed corresponds to the set of occurrence nets we wanted. Of course, this construction is ill-defined if the net N is not safe for one of the source configurations in I .

If furthermore the produced automaton is SP, which is again a property of the net N , then the construction of this paper yields an expression $\mathcal{E}(\mathcal{A})$ representing faithfully the set of occurrence nets of processes going from any configuration in I to any configuration in F .

References

- [1] H. Andréka and D. Bredikhin. The equational theory of union-free algebras of relations. *Alg. Univ.*, 33(4):516–532, 1995.
- [2] P. Brunet and D. Pous. Petri automata for Kleene allegories. In *Proc. LICS*, 2015. to appear.
- [3] P. J. Freyd and A. Scedrov. *Categories, Allegories*. NH, 1990.
- [4] D. Kozen. Typed kleene algebra. Technical report, Cornell University, 1998.
- [5] B. Möller. Typed kleene algebras. In *Mathematics Of Program Construction, Volume 3125 of LNCS*. Citeseer, 1999.
- [6] T. Murata. Petri nets: Properties, analysis and applications. *Proc. of the IEEE*, 77(4):541–580, Apr 1989.
- [7] D. Pous. Untyping typed algebras and colouring cyclic linear logic. *Logical Methods in Computer Science*, 8(2), 2012.
- [8] J. Valdes, R. E. Tarjan, and E. L. Lawler. The recognition of series parallel digraphs. In *Proceedings of the Eleventh Annual ACM Symposium on Theory of Computing*, STOC '79, pages 1–12, New York, NY, USA, 1979. ACM.

A. Omitted proofs

Definition 27 (Connected component)

A set C of nodes of a graph $G = \langle V, E \rangle$ is a *connected component* if $\forall x \in C, \forall y \in V, y$ is in C if and only if there is a non-oriented path from x to y in G . *

This means any two nodes in C are linked by non-oriented path, and that the set C is maximal for this property (with the containment order).

Definition 28 (Elementary boxes)

A box η is said to be *elementary* if its graph contains a single non-trivial connected component. *

Lemma 29. *If η is a valid elementary box, then its graph is necessarily a bow tie.*

Proof. Validity means the graph has to be an SP-slice, and reduced connected SP-slices are always bow ties. \square

Definition 30 (Support and target of an elementary box)

Let η be an elementary box. The *support* (respectively *target*) of η is the set of input (resp. output) ports of η which are mapped inside its unique non-trivial connected component. *

Lemma 31. *Any valid box β can be decomposed into a set of valid elementary boxes η_1, \dots, η_p such that:*

1. $\beta = \eta_1 \odot \dots \odot \eta_p$;
2. $\forall i \neq j, \text{support}(\eta_i) \cap \text{support}(\eta_j) = \emptyset$;
3. if β has type $\sigma \multimap \sigma$, then each η_i has the same type and $\text{support}(\eta_i) = \text{target}(\eta_i)$.

Proof. We number the connected components of β from C_1 to C_p .

Suppose β is valid, then there is a run $\xi = \xi_0 \xrightarrow{t_1} \dots \xrightarrow{t_n} \xi_n$ such that $\beta = \beta(\xi) = \beta_{t_1, \xi_0} \odot \dots \odot \beta_{t_n, \xi_{n-1}}$. Notice that for any index i , the box $\beta_{t_i, \xi_{i-1}}$ is elementary (see the definition of $\beta_{t, C}$). Thus it can be associated with a single non-trivial connected component in β . We introduce the function $c : \{1, \dots, n\} \rightarrow \{1, \dots, p\}$, such that the transition i is associated with the component $C_{c(i)}$.

The key argument is the following. For any index i , if the set of output places of transition t_i contains a place that appears in the input of the transition t_{i+1} , then the box $\beta_{t_i, \xi_{i-1}} \odot \beta_{t_{i+1}, \xi_i}$ is elementary. This means that $c(i) = c(i+1)$. By contraposition, we may then deduce that if $c(i) \neq c(i+1)$, then the intersection of the outputs of t_i and the inputs of t_{i+1} is empty. This can in turn be used to build a configuration ξ'_i such that:

- $\xi_{i-1} \xrightarrow{t_{i+1}} \xi'_i \xrightarrow{t_i} \xi_{i+1}$;
- and $\beta_{t_i, \xi_{i-1}} \odot \beta_{t_{i+1}, \xi_i} = \beta_{t_{i+1}, \xi_{i-1}} \odot \beta_{t_i, \xi'_i}$.
($\xi'_i = \xi_{i-1} \setminus \underline{t_{i+1}} \cup \{p \mid \exists a : (a, p) \in t_{i+1}\}$)

This means we may group the transitions that contribute to the same connected component. Formally we apply a permutation π to the run ξ , such that:

- $\pi(\xi) = \xi_0 \xrightarrow{t_{\pi(1)}} \dots \xrightarrow{t_{\pi(n)}} \xi_n$
- $\beta(\xi) = \beta(\pi(\xi))$;
- $\pi(\xi)$ can be decomposed as a sequence of runs:

$$\pi(\xi) = \xi_0 \xrightarrow{\chi_1} \xi'_1 \dots \xi'_{p-1} \xrightarrow{\chi_p} \xi_n,$$

such that $t_i \in \chi_j \Leftrightarrow c(i) = j$.

We now define $\eta_j := \beta(\chi_j)$. By construction η_j is valid, it contains only the connected component C_j so it is elementary, and $\beta = \beta(\xi) = \beta(\pi(\xi)) = \eta_1 \odot \dots \odot \eta_p$. Furthermore, for each $p \in \xi_0$,

$p \in \text{support}(\eta_j)$ if and only if $\underline{p}(p) \in C_j$. As the connected components are pairwise disjoint, so are the supports of the η_j . Further analysis of the run $\pi(\xi)$ also shows that $\forall i < j, \text{target}(\eta_i) \cap \text{support}(\eta_j) = \emptyset$.

If β has type $\sigma \multimap \sigma$, that means that in particular $\xi_0 = \xi_n$. This entails the following equality:

$$\bigcup_j \text{support}(\eta_j) = \bigcup_j \text{target}(\eta_j).$$

(The remaining places are those untouched by the execution.) Hence:

$$\begin{aligned} \text{target}(\eta_1) &\subseteq \bigcup_j \text{target}(\eta_j) = \bigcup_j \text{support}(\eta_j) \\ &= \text{support}(\eta_1) \cup \bigcup_{1 < j} \text{support}(\eta_j) \\ \text{support}(\eta_1) &\subseteq \bigcup_j \text{support}(\eta_j) = \bigcup_j \text{target}(\eta_j) \\ &= \text{target}(\eta_1) \cup \bigcup_{1 < j} \text{target}(\eta_j) \end{aligned}$$

We know that $\forall 1 < j, \text{target}(\eta_1) \cap \text{support}(\eta_j) = \emptyset$, hence $\text{target}(\eta_1) \subseteq \text{support}(\eta_1)$. To get the other direction, one has to notice that the run $\chi_2 \dots \chi_p$ may be executed starting from ξ_0 (rather than ξ'_1). This is again because of the possibility to exchange “independent” transitions. Safety of the automaton ensures that

$$\text{support}(\eta_1) \cap \bigcup_{1 < j} \text{target}(\eta_j) = \emptyset.$$

Thus we obtain $\text{support}(\eta_1) \subseteq \text{target}(\eta_1)$, finally proving $\text{support}(\eta_1) = \text{target}(\eta_1)$. This in turn ensures that $\xi'_1 = \xi_0$, and enable us to repeat the proof for η_2, \dots, η_n .

The last thing to do is to “type check” these elementary boxes. Let’s detail the computation of the type β outputs when given a tree σ . Each connected component C_j is a bow tie of centre o_j , with inputs S_j . For each of those, there must be a sub-tree T_j in σ with leaves S_j (otherwise there would be no way of reducing $\sigma \odot \beta$ to a tree). We then remove that sub-tree, and replace it by the top-down tree rooted in o_j . Because all S_j are pairwise disjoint, the trees T_j are also pairwise disjoint (otherwise they would share some leaves). With this description of the type-checking, one realises that 1) each of the η_j contribute to the final type on disjoint subtrees 2) and each of them leaves the rest of the tree alone. Because $\forall j, S_j = \text{support}(\eta_j) = \text{target}(\eta_j)$ we can conclude that η_j replaces the tree with leaves S_j with the tree with leaves S_j , and thus preserves the type $\sigma \multimap \sigma$. \square