

Petri automata for Kleene allegories

Paul Brunet, Damien Pous

► **To cite this version:**

Paul Brunet, Damien Pous. Petri automata for Kleene allegories. Logic in Computer Science, Jul 2015, Kyoto, Japan. IEEE, pp.68-79, 2015, <<http://lics.rwth-aachen.de/lics15/>>. <10.1109/LICS.2015.17>. <hal-01073936v3>

HAL Id: hal-01073936

<https://hal.archives-ouvertes.fr/hal-01073936v3>

Submitted on 16 Mar 2015

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.

Petri automata for Kleene allegories

Paul Brunet and Damien Pous

Plume team – LIP, CNRS, ENS de Lyon, Inria, UCBL, Université de Lyon, UMR 5668
{paul.brunet,damien.pous}@ens-lyon.fr

Abstract—Kleene algebra axioms are complete with respect to both language models and binary relation models. In particular, two regular expressions recognise the same language if and only if they are universally equivalent in the model of binary relations.

We consider Kleene allegories, i.e., Kleene algebras with two additional operations which are natural in binary relation models: intersection and converse. While regular languages are closed under those operations, the above characterisation breaks. Instead, we give a characterisation in terms of languages of directed and labelled graphs. We then design a finite automata model allowing to recognise such graphs, by taking inspiration from Petri nets.

This model allows us to obtain decidability of identity-free relational Kleene lattices, i.e., the equational theory generated by binary relations on the signature of regular expressions with intersection, but where one forbids unit. This restriction is used to ensure that the corresponding graphs are acyclic. The decidability of graph-language equivalence in the full model remains open.

I. INTRODUCTION

We consider binary relations and the operations of union (\cup), intersection (\cap), composition (\cdot), converse ($\check{\cdot}$), transitive closure (\cdot^+), reflexive-transitive closure (\cdot^*), and the constants identity (1) and empty relation (0). This model gives rise to an (in)equational theory: a pair of terms e, f made from those operations and some variables a, b, \dots is a *valid equation*, denoted $\text{Rel} \models e = f$, if the corresponding equality holds universally. Similarly, an inequation $\text{Rel} \models e \leq f$ is valid when the corresponding containment holds universally. Here are valid equations and inequations: they hold whatever the relations we assign to variables a, b , and c .

$$\text{Rel} \models (a \cup b)^* \cdot b \cdot (a \cup b)^* = (a^* \cdot b \cdot a^*)^+ \quad (1)$$

$$\text{Rel} \models a^* \leq 1 \cup a \cdot a^{\check{\cdot}} \cdot a^+ \quad (2)$$

$$\text{Rel} \models a \cdot b \cap c \leq a \cdot (b \cap a^{\check{\cdot}} \cdot c) \quad (3)$$

$$\text{Rel} \models a^+ \cap 1 \leq (a \cdot a)^+ \quad (4)$$

Various fragments of this theory have been studied in the literature:

- *Kleene algebra* [7], where one removes intersection and converse, so that terms are plain regular expressions. The theory is decidable by Kleene’s work [11], and actually PSPACE-complete [14], [15]. Equation (1) lies in this fragment, and one can notice that the two expressions recognise the same language.
- *Kleene algebra with converse*, where one only removes intersection, is also a decidable fragment [3]. It remains PSPACE [5]. Inequation (2) belongs to this fragment.
- (representable, distributive) *allegories* [8], sometimes called positive relation algebras, where transitive and

reflexive-transitive closures are not allowed. They are decidable [8, page 208]; Inequation (3) is known as the *modularity law* in this setting.

To the best of our knowledge, the decidability of the whole theory, *Kleene allegories*, is open. Here we obtain several important steps towards the resolution of this problem:

- 1) we give a characterisation of the full (in)equational theory in terms of graph languages;
- 2) we design an automata model inspired by Petri nets, that makes it possible to recognise such graphs;
- 3) we show how to associate such a graph automaton to any term of Kleene allegories;
- 4) using these graph automata, we give a decision procedure for the fragment where converse and identity are forbidden.

The latter fragment was studied recently by Andréka et al. [1]; its decidability was open as far as we know. The restriction to this fragment allows us to exploit simplifying assumptions about the produced automata, and to obtain a coinductive algorithm for language inclusion (Section V). We actually show that language inclusion for these automata is EXPSpace-complete (Section VI).

The next problem, which remains open, consists in obtaining the decidability of language inclusion in the full automata model: together with the presented results, this would entail decidability of Kleene allegories. We outline some of the difficulties arising with converse or unit in presence of intersection in Section V-D.

We continue this introductory section by an informal description of the graph language characterisation and our automata model.

A. Languages

In the simple case of Kleene algebra, i.e., without converse and intersection, the (in)equational theory generated by relations can be characterised by using regular languages. Write $\llbracket e \rrbracket$ for the language denoted by a regular expression e ; for any two regular expressions e, f , we have

$$\text{Rel} \models e \leq f \text{ if and only if } \llbracket e \rrbracket \subseteq \llbracket f \rrbracket . \quad (5)$$

(This result is easy and folklore; proving that this is also equivalent to provability using Kleene algebra axioms [4], [12], [13] is much harder.)

While regular languages are closed under intersection and converse, the above characterisation does not extend to those operations. For intersection, consider two distinct variables a




Figure 1: Graphs associated to some ground terms.

and b . The extended regular expressions $a \cap b$ and 0 both recognise the empty language, while $\text{Rel} \neq a \cap b = 0$: one can interpret a and b with intersecting relations. For converse, the extended regular expressions a and a^\sim both recognise the singleton language consisting of the single-letter word a . Yet $\text{Rel} \neq a = a^\sim$, as there are non-symmetric relations.

B. Graphs

Freyd and Scedrov’ decision procedure for representable allegories [8, page 208] relies on a notion of directed, labelled, 2-pointed graph. The same notion was proposed independently by Andr eka and Bredikhin [2], in a more comprehensive way.

Call *ground terms* the terms in the syntax of allegories (composition, intersection, converse, and unit). A ground term u can be represented as a labelled directed graph $G(u)$ with two distinguished vertices called the *input* and the *output*. We give some examples in Figure 1, see Definition 1 for a precise definition.

These graphs can be endowed with a preorder relation: we write $G \blacktriangleleft F$ when there exists a graph homomorphism from F to G preserving labels, inputs, and outputs. For instance the graph corresponding to $(a \cdot (b \cap c)) \cap d$ is smaller than the graph of $(a \cdot b) \cap (a \cdot c)$, thanks to the homomorphism depicted in Figure 2 using dotted arrows. Notice that the homomorphism needs not be injective or surjective, so that this preorder has nothing to do with the respective sizes of the graphs: a graph may very well be smaller than another in the sense of \blacktriangleleft , while having more vertices or edges (and vice versa).

The key result from Freyd and Scedrov [8, page 208], or Andr eka and Bredikhin [2, Theorem 1], is that for any two ground terms u, v , we have

$$\text{Rel} \models u \leq v \text{ if and only if } G(u) \blacktriangleleft G(v) . \quad (6)$$

The graphs are finite so that one can search exhaustively for a homomorphism, whence the decidability result.

C. Graph languages

To extend the above graph-theoretical characterisation to Kleene allegories, we need to handle union and (reflexive-) transitive closures. It suffices for that to consider sets of graphs: to each expression e , we associate a set of graphs $\mathcal{G}(e)$. This set is most often infinite when the expression e contains (reflexive-)transitive closures.

Writing \blacktriangleleft^* for the downward closure of a set of graphs X by the preorder \blacktriangleleft on graphs, we obtain the following generalisation of both (5) and (6): for any two expressions e and f ,

$$\text{Rel} \models e \leq f \text{ if and only if } \blacktriangleleft^* \mathcal{G}(e) \subseteq \blacktriangleleft^* \mathcal{G}(f) . \quad (7)$$

This is Theorem 6 in the sequel, and this result is almost there in the work by Andr eka et al. [1], [2]. To the best of our knowledge this explicit formulation is new, as well as its use towards decidability results.

When e and f are ground terms, we recover the characterisation (6) for representable allegories: $\mathcal{G}(e)$ and $\mathcal{G}(f)$ are singleton sets in this case. For plain regular expressions, the graphs are just words and the preorder \blacktriangleleft reduces to isomorphism. We thus recover the characterisation (5) for Kleene algebra. This result also generalises the characterisation provided by  sik et al. [3] for Kleene algebra with converse: graphs of expressions without intersection are just words over a duplicated alphabet, and the corresponding restriction of the preorder \blacktriangleleft precisely corresponds to the word rewriting system they use (see Remark 7).

D. Petri automata

In order to exploit the above characterisation and obtain decidability results, one has first to represent graph languages in a finitary way. We propose for that a new finite automata model, largely based on Petri nets [16]–[18]. We describe this model below, ignoring converse and unit for the sake of clarity.

Recall that a Petri net consists of

- a finite set of *places*, denoted with circles;
- a set of *transitions*, denoted with rectangles;
- for each transition, a set of input places and a set of output places, denoted with arrows;
- an *initial place*, denoted by an entrant arrow;
- a set of *final markings*, denoted by dotted boxes (a marking, or configuration, being a set of places).

The execution model is the following: start by putting a token on the initial place; choose a transition whose input places all contain a token, remove those tokens and put new tokens in the output places of the transition; repeat this process until a final marking is reached. The obtained sequence of transitions is called an *accepting run*. (We actually restrict ourselves to *safe* Petri nets, to ensure that there is always at most one token in a given place when playing this game.)

A *Petri automaton* is just a safe Petri net¹ with variables labelling the outputs of each transition. The automaton depicted

¹For every reachable marking in a *safe* net, there is at most one token in any given place.





Figure 2: A graph homomorphism.

below is the automaton we will construct for the ground term $a \cdot b \cap a \cdot c$. Any run must start by firing the left-most transition, reaching the marking $\{B, C\}$; then we have the choice of firing the upper transition first, reaching the marking $\{D, C\}$, or the lower one, reaching the marking $\{B, E\}$. In both cases we reach the final marking $\{D, E\}$ by firing the remaining transition.



To read a graph in such an automaton, we try to find an accepting run that matches the graph up to homomorphism (Definitions 10 and 11). We do that by using an evolving function from the tokens to the vertices of the graph. We start with the function mapping the unique token, in the initial place, to the input vertex of the graph. To fire a transition, we must check that all its input tokens are mapped to the same vertex in the graph, and that this vertex has several outgoing edges, labelled according to the outputs of the transition. If this is the case, we update the function by removing the mappings corresponding to the deleted tokens, and by adding new mappings for each of the created tokens (using the target vertices of the aforementioned outgoing edges, according to the labels). The graph is accepted by the Petri automaton if we can reach a final marking of the Petri automaton, with all tokens mapped to the output vertex of the graph.


For instance, the previous automaton accepts the graph of $a \cdot b \cap a \cdot c$ (F in Figure 2). We start with the function $\{A \mapsto 0\}$. We can fire the first transition, updating the function into $\{B \mapsto 1, C \mapsto 2\}$ (We could also choose to update the function into $\{B \mapsto 2, C \mapsto 1\}$, or $\{B, C \mapsto 1\}$, or $\{B, C \mapsto 2\}$ but this would lead to a dead-end). Then we can fire the upper transition, evolving the function into $\{D \mapsto 3, C \mapsto 2\}$, and we finish by firing the remaining transition, obtaining the function $\{D, E \mapsto 3\}$.

We call *language* of \mathcal{A} the set of graphs $\mathcal{L}(\mathcal{A})$ accepted by a Petri automaton \mathcal{A} . This language is downward-closed: $\mathcal{L}(\mathcal{A}) = \ulcorner \mathcal{L}(\mathcal{A})$. For instance, the previous automaton also accepts the graph G from Figure 2, which is smaller than F . Indeed, when we fire the first transition, we can associate the two newly created tokens (in places B and C) to the

same vertex (5). This actually corresponds to composing the functions used to accept F with the homomorphism depicted with dotted arrows.


This automata model is expressive enough for Kleene allegories: for any expression e , we can construct a Petri automaton $\mathcal{A}(e)$ such that $\mathcal{L}(\mathcal{A}(e)) = \ulcorner \mathcal{G}(e)$ (Section IV). We give three other examples of Petri automata to give more intuition on their behaviour.

The first transition in the previous Petri automaton splits the initial token into two tokens, which are moved concurrently in the remainder of the run. This corresponds to an intersection in the considered expression. This is to be contrasted with the behaviour of the following automaton, which we would construct for the non-ground expression $a \cdot b \cup a \cdot c$. This automaton has two accepting runs: $\{A\}, \{B\}, \{D\}$ and $\{A\}, \{C\}, \{E\}$, which can be used to accept the (graphs of the) ground terms $a \cdot b$ and $a \cdot c$.




In a sense, two transitions competing for the same tokens represent a non-deterministic choice, i.e., a union in the starting expression.

Still in the first example, the two tokens created by the first transition are later collected in the final marking. Tokens may also be collected and merged by a transition. Consider for instance the following automaton for $(a \cap b \cdot c) \cdot d$. It has only one accepting run, $\{A\}, \{B, C\}, \{B, D\}, \{E\}$, and this run can be used to read the fourth graph from Figure 1.



As a last example, consider the following automaton for the expression $a \cap b^+ \cdot c$. The upper transition introduces a loop, so that there are infinitely many accepting runs. For any $n > 0$, the graph of the ground term $a \cap b^n \cdot c$ is accepted by this automaton.



According to characterisation (7), the next step is to decide whether the containment $\mathcal{L}(\mathcal{A}) \subseteq \mathcal{L}(\mathcal{B})$ holds, for two given Petri automata \mathcal{A} and \mathcal{B} . Several difficulties arise, that do not appear with classical word automata. Our solution nevertheless uses a standard coinductive approach, where we define an appropriate notion of simulation (Section V-C).




Figure 3: Inductive construction of the graph of a ground term.

Standard notations. For any sets A, B , we write $\mathcal{P}(S) := \{P \mid P \subseteq S\}$ for the set of subsets of A , $A \rightarrow B$ for the set of functions from A to B , and $A \multimap B$ for the set of partial maps from A to B . The domain of a partial map f is denoted by $\text{dom}(f)$.

II. GRAPH-THEORETICAL CHARACTERISATION

We consider the signature $(\cap, \cup, \cdot, \cdot^+, \cdot^{\sim}, 0, 1)$ of Kleene allegories, where e^* is an abbreviation for $1 \cup e^+$. We fix a set X of variables, and we denote by Reg_X^{\cap} the set of expressions $e, f \dots$ built from variables in X with these connectives. *Ground terms* are the expressions $u, v, w \dots$ built from X only with the sub-signature $(\cap, \cdot, \cdot^{\sim}, 1)$. If $\sigma : X \rightarrow \mathcal{P}(S \times S)$ is an interpretation of the alphabet X into some space of relations, we write $\bar{\sigma}$ for the unique homomorphism extending σ into a function from Reg_X^{\cap} to $\mathcal{P}(S \times S)$. An inequation between two expressions e and f is *valid*, written $\text{Rel} \models e \leq f$, if for any relational interpretation σ we have $\bar{\sigma}(e) \subseteq \bar{\sigma}(f)$.

We let G range over 2-pointed labelled directed graphs, which we simply call *graphs* in the sequel. Those are tuples $\langle V, E, \iota, o \rangle$ with V a finite set of vertices, $E \subseteq V \times X \times V$ a set of edges labelled with X , and $\iota, o \in V$ the two distinguished vertices, respectively called *input* and *output*.

Definition 1 (Graph of a ground term: $G(w)$)

To each ground term w , we associate a graph $G(w)$, by induction on w . The graph of $a \in X$ has one edge labelled by a linking its input to its output. The graph for 1 has only one vertex, both input and output. The composition 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. The intersection on graphs consists in merging their inputs and merging their outputs. The converse consists simply in exchanging the input and the output of a graph. *

See Figure 3 for a graphical description of this construction. Those graphs were introduced independently by Freyd and Scedrov [8, page 208], and Andr eka and Bredikhin [2].

Definition 2 (Graph homomorphism, preorders \blacktriangleleft and \triangleleft)

A *graph homomorphism* from $\langle V_1, E_1, \iota_1, o_1 \rangle$ to $\langle V_2, E_2, \iota_2, o_2 \rangle$ is a map $\varphi : V_1 \rightarrow V_2$ such that $\varphi(\iota_1) = \iota_2$, $\varphi(o_1) = o_2$, and $(p, x, q) \in E_1$ entails $(\varphi(p), x, \varphi(q)) \in E_2$. We denote by \blacktriangleleft the relation on graphs defined by $G \blacktriangleleft G'$ if there exists a graph homomorphism from G' to G . This

relation gives rise to a preorder on ground terms, written \triangleleft and defined by $u \triangleleft v$ if $G(u) \blacktriangleleft G(v)$. *

Given a set S of graphs, we write $\blacktriangleleft S$ for its downward closure: $\blacktriangleleft S := \{G \mid G \blacktriangleleft G', G' \in S\}$. Similarly, we write $\triangleleft S$ for the downward closure of a set of ground terms w.r.t. \triangleleft .

As explained in the introduction, the above preorder on ground terms precisely characterises inclusion under arbitrary relational interpretations:

Theorem 3 ([2, Theorem 1], or [8, page 208]). *For all ground terms u, v , we have $\text{Rel} \models u \leq v \Leftrightarrow u \triangleleft v$.*

To extend this result to Kleene allegories, we introduce the following generalisation of the language of a regular expression. Sets of words become sets of ground terms.

Definition 4 (Terms and graphs of an expression)

The *set of terms* of an expression $e \in \text{Reg}_X^{\cap}$, written $\llbracket e \rrbracket$, is the set of ground terms defined inductively as follows:

$$\begin{aligned} \llbracket 1 \rrbracket &:= \{1\} & \llbracket 0 \rrbracket &:= \emptyset & \llbracket x \rrbracket &:= \{x\} \\ \llbracket e \cdot f \rrbracket &:= \{w \cdot w' \mid w \in \llbracket e \rrbracket \text{ and } w' \in \llbracket f \rrbracket\} \\ \llbracket e \cap f \rrbracket &:= \{w \cap w' \mid w \in \llbracket e \rrbracket \text{ and } w' \in \llbracket f \rrbracket\} \\ \llbracket e \cup f \rrbracket &:= \llbracket e \rrbracket \cup \llbracket f \rrbracket \\ \llbracket e^* \rrbracket &:= \bigcup_{n \in \mathbb{N}} \{w_1 \cdots w_n \mid \forall i, w_i \in \llbracket e \rrbracket\} \\ \llbracket e^{\sim} \rrbracket &:= \{w^{\sim} \mid w \in \llbracket e \rrbracket\} . \end{aligned}$$

The *set of graphs produced* by an expression e , denoted by $\mathcal{G}(e)$ is the set of graphs associated to the ground terms in $\llbracket e \rrbracket$:

$$\mathcal{G}(e) := \{G(w) \mid w \in \llbracket e \rrbracket\} . *$$

To obtain the characterisation announced in the introduction, we need a slight refinement of a lemma established by Andr eka, Mikul as, and N emeti [1]:

Lemma 5. *For all expression $e \in \text{Reg}_X^{\cap}$, and all relational interpretations $\sigma : X \rightarrow \mathcal{P}(S \times S)$, we have*

$$\bar{\sigma}(e) = \bigcup_{u \in \llbracket e \rrbracket} \bar{\sigma}(u) = \bigcup_{w \in \blacktriangleleft \llbracket e \rrbracket} \bar{\sigma}(w) .$$

Proof. The first equality is exactly [1, Lemma 2.1]; for the second one, we use the fact that $\bar{\sigma}(w) \subseteq \bar{\sigma}(u)$ whenever $w \triangleleft u$, thanks to Theorem 3 (i.e., [2, Theorem 1]). \square

Theorem 6. *The following properties are equivalent, for all expressions $e, f \in \text{Reg}_X^{\cap}$:*

- (i) $\text{Rel} \models e \leq f$,
- (ii) $\llbracket e \rrbracket \subseteq \blacktriangleleft \llbracket f \rrbracket$,
- (iii) $\mathcal{G}(e) \subseteq \blacktriangleleft \mathcal{G}(f)$.

Proof. We give a detailed proof in Appendix A. The implication (ii) \Rightarrow (i) follows easily from Lemma 5, and (iii) \Rightarrow (ii) is a matter of unfolding definitions. For (i) \Rightarrow (iii), we mainly use [2, Lemma 3]. \square

(The exact characterisation announced in the introduction (7) follows: for any sets X, Y , we have $\blacktriangleleft X \subseteq \blacktriangleleft Y$ iff $X \subseteq \blacktriangleleft Y$.)





Figure 4: A Petri automaton. The initial place is A , and the final configurations are $\{I\}$ and $\{F, G\}$.

Also notice that while $\mathcal{G}(f)$ only contains graphs emanating from ground terms, this is not the case for its closure $\mathcal{G}(f)$. For instance, $\mathcal{G}((a \cdot b) \cap (c \cdot d))$ contains the following graph, which is not the graph of any ground term.



Remark 7. The graphs associated to ground terms without intersection are isomorphic to words over a duplicated alphabet. A graph homomorphism between two such graphs is precisely what Ésik et al. call an *admissible map* [3]. Theorem 6 can thus be seen as a generalisation of [3, Theorem 5.3].

III. PETRI AUTOMATA

We extend the set X of variables into a set \bar{X} of labels: $\bar{X} := X \cup \{x^- \mid x \in X\} \cup \{1\}$. A Petri automaton is a Petri net whose transition's outputs are labelled by \bar{X} .

Definition 8 (Petri Automaton)

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

- P is a finite set of *places*,
- $\mathcal{T} \subseteq \mathcal{P}(P) \times \mathcal{P}(\bar{X} \times P)$ is a set of *transitions*,
- $\iota \in P$ is the *initial place* of the automaton,
- $\mathcal{F} \subseteq \mathcal{P}(P)$ is a set of *final configurations*, a *configuration* being a set of places.

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 $\bar{t} \subseteq \bar{X} \times P$ is the *output of t* . *

We use the graphical notation from the introduction to represent Petri automata; the Petri automaton from Figure 4 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 $\xi' = \xi \setminus \underline{t} \cup \{p \in P \mid \exists x \in \bar{X} : (x, p) \in \bar{t}\}$. We write $\xi \xrightarrow{t} \xi'$ in this case.

A set of transitions $T \subseteq \mathcal{T}$ is *statically compatible* (or just compatible) if their inputs are pairwise disjoint. If furthermore all transitions in T are enabled in a configuration ξ , one can observe that the configuration ξ' reached after firing them successively does not depend on the order in which they are fired. In that case we write $\xi \xrightarrow{T} \xi'$.




Figure 5: An accepting run in the automaton from Figure 4.

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 [16]). Formally, with our definitions: a Petri automaton $\langle P, \mathcal{T}, \iota, \mathcal{F} \rangle$ 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 explain how to use Petri automata to define languages of graphs. We first define the *runs* of an automaton.

Definition 9 (Run, accepting run, parallel run)

A *run* is a sequence $\xi = \langle (\xi_k)_{0 \leq k \leq n}, (t_k)_{0 \leq k < n} \rangle$ 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 $\xi_n \in \mathcal{F}$, we call ξ an *accepting run*.

A *parallel run* is defined similarly, as a sequence $\Xi = \langle (\Xi_k)_{0 \leq k \leq n}, (T_k)_{0 \leq k < n} \rangle$, where the $T_k \subseteq \mathcal{T}$ are compatible sets of transitions such that $\Xi_k \xrightarrow{T_k} \Xi_{k+1}$. *

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

Consider the following sequence of transitions from the automaton in Figure 4:

$$\xi = \langle (\xi_0, \xi_1, \xi_2, \xi_3, \xi_4, \xi_5, \xi_6), (t_0, t_1, t_2, t_3, t_4, t_5) \rangle,$$

with

$$\begin{array}{l} \xi_0 = \{A\}, \\ \xi_1 = \{B, G\}, \\ \xi_2 = \xi_4 = \{C, E, G\}, \\ \xi_3 = \xi_5 = \{D, E, G\}, \\ \xi_6 = \{F, G\}. \end{array} \quad \left| \begin{array}{l} t_0 = (\{A\}, \{(b, B), (a, G)\}), \\ t_1 = (\{B\}, \{(c, C), (b, E)\}), \\ t_2 = t_4 = (\{C\}, \{(a, D)\}), \\ t_3 = (\{D, E\}, \{(c, C), (b, E)\}), \\ t_5 = (\{D, E\}, \{(d, F)\}). \end{array} \right.$$

We have

$$\{A\} \xrightarrow{t_0} \{B, G\} \xrightarrow{t_1} \{C, E, G\} \xrightarrow[t_3]{t_2, t_4} \{D, E, G\} \xrightarrow{t_5} \{F, G\},$$

and since $\{A\}$ is the initial configuration and $\{F, G\} \in \mathcal{F}$, this sequence is an accepting run. It can be represented graphically as in Figure 5.

As for standard finite state automata, we now need to specify how to read a graph in an automaton. As explained in the introduction, this is done by linking the intermediate configurations of a run to vertices in the graph, and by imposing conditions to match transitions with labelled edges of the graph.

Definition 10 (Reading, parallel reading, language of a run)

A *reading* of $G = \langle V, E, \iota, o \rangle$ along a run $\xi =$




Figure 6: Graph produced by the run depicted in Figure 5.

$\langle (\xi_k)_{0 \leq k \leq n}, (\underline{t}_k, \bar{t}_k)_{0 \leq k < n} \rangle$ is a sequence $(\rho_k)_{0 \leq k \leq n}$ such that for all k , ρ_k is a map from ξ_k to V , $\rho_0(\xi_0) = \{\iota\}$, $\rho_n(\xi_n) = \{o\}$, and $\forall k < n$, the following holds:

- all tokens in the input of the transition are mapped to the same vertex in the graph: $\forall p, q \in \underline{t}_k, \rho_k(p) = \rho_k(q)$;
- the images of tokens in ξ_k that are not in the input of the transition are unchanged: $\forall p \in \xi_k \setminus \underline{t}_k, \rho_k(p) = \rho_{k+1}(p)$;
- each pair in the output of the transition can be “validated” by the graph: $\forall p \in \bar{t}_k, \forall (x, q) \in \bar{t}_k$,

$$\begin{aligned} x \in X &\Rightarrow (\rho_k(p), x, \rho_{k+1}(q)) \in E, \\ x = y^\sim \text{ and } y \in X &\Rightarrow (\rho_{k+1}(q), y, \rho_k(p)) \in E, \\ x = 1 &\Rightarrow \rho_k(p) = \rho_{k+1}(q). \end{aligned}$$

Similarly, we define a *parallel reading* ρ along some parallel run $\Xi = \langle (\Xi_k)_{0 \leq k \leq n}, (T_k)_{0 \leq k < n} \rangle$ by requiring that: $\rho_0(\Xi_0) = \{\iota\}$, $\rho_n(\Xi_n) = \{o\}$, and $\forall k < n$ the following holds:

- $\forall p \in \Xi_k \setminus \bigcup_{(\underline{t}, \bar{t}) \in T_k} \underline{t}, \rho_{k+1}(p) = \rho_k(p)$;
- $\forall (\underline{t}, \bar{t}) \in T_k, \forall p, q \in \underline{t}, \rho_k(p) = \rho_k(q)$;
- $\forall (\underline{t}, \bar{t}) \in T_k, \forall p \in \underline{t}, \forall (x, q) \in \bar{t}$,

$$\begin{aligned} x \in X &\Rightarrow (\rho_k(p), x, \rho_{k+1}(q)) \in E, \\ x = y^\sim \text{ and } y \in X &\Rightarrow (\rho_{k+1}(q), y, \rho_k(p)) \in E, \\ x = 1 &\Rightarrow \rho_k(p) = \rho_{k+1}(q). \end{aligned}$$

The *language of a run* ξ , denoted by $\mathcal{L}(\xi)$ is the set of graphs that can be read along ξ . *

The language of a Petri automaton is finally obtained by considering all accepting runs.

Definition 11 (Language recognised by a Petri automaton)

The language recognised by \mathcal{A} , written $\mathcal{L}(\mathcal{A})$, is the following set of graphs:

$$\mathcal{L}(\mathcal{A}) := \bigcup_{\xi \text{ accepting in } \mathcal{A}} \mathcal{L}(\xi) . \quad *$$

The language of a run ξ can be characterised by using a single graph which we call the *graph produced* by ξ : graphs are accepted by ξ exactly when they are smaller than the graph produced by ξ according to \blacktriangleleft (Lemma 14 below). For instance, the run presented in Figure 5 produces the graph depicted in Figure 6.

This graph is obtained in two steps, by first considering a notion of *trace*, which is a graph labelled with \bar{X} rather than X , and which actually corresponds to the notion of pomset-trace from standard Petri nets (see Section VII for more details on this correspondence).

The trace is constructed by creating a vertex k for each transition $t_k = (\underline{t}_k, \bar{t}_k)$ of the run, plus a final vertex n . We

add an edge (k, x, l) whenever there is some place q such that $(x, q) \in \bar{t}_k$, 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 12 (Trace of a run)

Let $\xi = \langle (\xi_k)_{0 \leq k \leq n}, (\underline{t}_k, \bar{t}_k)_{0 \leq k < n} \rangle$ 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 $\llbracket \xi \rrbracket := \langle \{0, \dots, n\}, E_\xi, 0, n \rangle$ with $E_\xi := \{(k, x, \nu(k+1, q)) \mid (x, q) \in \bar{t}_k\}$. *

To get the final graph, which is labelled by X , one identifies nodes linked by edges labelled by 1, and one replaces each edge of the form (i, x^\sim, j) by (j, x, i) . Formally:

Definition 13 (Graph produced by a run)

Let $\xi = \langle (\xi_k)_{0 \leq k \leq n}, (\underline{t}_k, \bar{t}_k)_{0 \leq k < n} \rangle$ be run. Let \equiv_ξ be the smallest equivalence relation on $\{0, \dots, n\}$ containing all pairs (i, j) such that $(i, 1, j) \in E_\xi$.

The *graph produced* by ξ , is the graph $\mathcal{G}(\xi)$ defined by

$$\begin{aligned} \mathcal{G}(\xi) &:= \langle \{[i]_\xi \mid 0 \leq i \leq n\}, E'_\xi, [0]_\xi, [n]_\xi \rangle \\ [i]_\xi &:= \{k \in \{0, \dots, n\} \mid i \equiv_\xi k\} \\ E'_\xi &:= \left\langle ([i]_\xi, x, [j]_\xi) \left| \begin{array}{l} x \in X \text{ and} \\ \exists k \in [i]_\xi, l \in [j]_\xi : \\ (k, x, l) \in E_\xi \text{ or} \\ (l, x^\sim, k) \in E_\xi \end{array} \right. \right\rangle \quad * \end{aligned}$$

We write $\mathcal{G}(\mathcal{A})$ for the set of graphs produced by accepting runs of a Petri automaton \mathcal{A} . To avoid confusions with the language $\mathcal{L}(\mathcal{A})$ of \mathcal{A} , we write “ G is produced by \mathcal{A} ” when $G \in \mathcal{G}(\mathcal{A})$, reserving language theoretic terminology like “ G is accepted by \mathcal{A} ” or “ \mathcal{A} recognises G ” to cases where we mean $G \in \mathcal{L}(\mathcal{A})$.

The graph produced by the run presented in Figure 5 happens to be equal to its trace, since it is labelled in X only. A more involved example is given in Figures 7 to 9. Notice that although the trace of a run is acyclic and can be endowed with a partial order structure (simply check that $\forall p, \nu(_, p)$ is increasing), it is not necessarily the case for its produced graph.

Lemma 14. *For any accepting run ξ , we have $G \in \mathcal{L}(\xi)$ if and only if $G \blacktriangleleft \mathcal{G}(\xi)$.*

Proof. Suppose there exists a graph homomorphism φ from $\mathcal{G}(\xi)$ to G . Then we can build a reading by defining $\rho_k(p) = \varphi([\nu(k, p)]_\xi)$ for $0 \leq k \leq n$ and $p \in \xi_k$. On the other hand, if we have a reading $(\rho_k)_{0 \leq k \leq n}$ of G , we can build a homomorphism φ by letting $\varphi([k]_\xi) = \rho_k(p)$ for any $p \in \xi_k$. As $(\rho_k)_k$ is a reading, φ is well defined. The details of this proof can be found in Appendix B. \square

As an immediate corollary, we obtain the following characterisation of the language of a Petri automaton.

Corollary 15. $\mathcal{L}(\mathcal{A}) = \blacktriangleleft \mathcal{G}(\mathcal{A})$.




Figure 7: A run ξ .




Figure 8: The trace of ξ .




Figure 9: The graph produced by ξ .

The left-hand side language is defined through readings along accepting runs, which is a local and incremental notion and which allows us to define *simulations* in Section V-C. By contrast, the right-hand side language is defined globally, which eases the following construction of an automaton recognising the language of an expression.

IV. FROM EXPRESSIONS TO AUTOMATA

We now show how to associate to any expression $e \in \text{Reg}_X^{\sim}$ an automaton $\mathcal{A}(e)$ that recognises the language $\mathbf{G}(e)$. In fact the automaton we obtain has an even stronger connection with e : the graphs in $\mathcal{G}(e)$ are exactly those produced by accepting runs in $\mathcal{A}(e)$. To make the construction simpler, we first modify the expression so that the operator \sim is only applied to variables, by using the following rewriting system:

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

(These rules preserve the set of graphs of the expression; also recall that e^* is a shorthand for $e^+ + 1$, so that we do not need to handle it explicitly.)

The formal construction is inductive; it is given in Definition 16 below. We describe it first in informal terms. For the case of $x \in \bar{X}$, we simply build an automaton with a single transition labelled by x , going from the initial place to a distinct final place². The union consists in putting both automata side by side, and merging their initial places. For the composition of \mathcal{A}_1 and \mathcal{A}_2 on the other hand we put one automaton in front of the other:

- the initial place of the resulting automaton is that of \mathcal{A}_1 ;
- the final configurations are those of \mathcal{A}_2 ;
- for each final configuration f of \mathcal{A}_1 , and for each initial transition (ι_2, \bar{t}) coming out of the initial place of \mathcal{A}_2 , we add a transition with input f and output \bar{t} . This last step amounts to adding epsilon-transitions from the final

²Note that this covers uniformly the case of the unit, of variables, and of conversed variables.

configurations of \mathcal{A}_1 to the initial place of \mathcal{A}_2 , and then apply some epsilon-elimination procedure.

We also put the two automata side by side for the intersection, but we merge their initial places, their initial transitions, and their final configurations:

- for any pair of initial transitions of the two automata $(\{\iota_1\}, \bar{t}_1), (\{\iota_2\}, \bar{t}_2)$, we put in the intersection automaton the transition $(\{\iota_1\}, \bar{t}_1 \cup \bar{t}_2)$;
- a final configuration of this automaton is the union of a final configuration from \mathcal{A}_1 and a final configuration of \mathcal{A}_2 .

For the transitive closure (\cdot^+) , we use the ideas for union and composition, adding loops from the final configurations using the initial transitions of the automaton.

Definition 16

To each expression $e \in \text{Reg}_X^{\sim}$, we associate a Petri automaton $\mathcal{A}(e)$ defined inductively as follows:

- $x \in \bar{X}$, $\mathcal{A}(x) := \langle \{0, 1\}, \{(\{0\}, \{(x, 1)\})\}, 0, \{\{1\}\} \rangle$
- $\mathcal{A}(0) := \langle \{0\}, \emptyset, 0, \emptyset \rangle$
- $\mathcal{A}(e_1 \cup e_2) := \langle P_1 \cup P_2, \mathcal{T}, \iota_1, \mathcal{F}_1 \cup \mathcal{F}_2 \rangle$ with $\mathcal{T} := \mathcal{T}_1 \cup \mathcal{T}_2 \cup \{(\{\iota_1\}, \bar{t}) \mid (\{\iota_2\}, \bar{t}) \in \mathcal{T}_2\}$.
- $\mathcal{A}(e_1 \cdot e_2) := \langle P_1 \cup P_2, \mathcal{T}, \iota_1, \mathcal{F}_2 \rangle$ with $\mathcal{T} := \mathcal{T}_1 \cup \mathcal{T}_2 \cup \{(f, \bar{t}) \mid f \in \mathcal{F}_1 \text{ and } (\{\iota_2\}, \bar{t}) \in \mathcal{T}_2\}$.
- $\mathcal{A}(e_1^+) := \langle P_1, \mathcal{T}, \iota_1, \mathcal{F}_1 \rangle$ with $\mathcal{T} := \mathcal{T}_1 \cup \{(f, \bar{t}) \mid f \in \mathcal{F}_1 \text{ and } (\{\iota_1\}, \bar{t}) \in \mathcal{T}_1\}$.
- $\mathcal{A}(e_1 \cap e_2) := \langle P_1 \cup P_2, \mathcal{T}, \iota_1, \mathcal{F} \rangle$ with

$$\begin{aligned} \mathcal{F} &:= \{f_1 \cup f_2 \mid f_1 \in \mathcal{F}_1, f_2 \in \mathcal{F}_2\} \text{ and} \\ \mathcal{T} &:= \{(\underline{t}, \bar{t}) \mid \exists i \in \{1, 2\} : (\underline{t}, \bar{t}) \in \mathcal{T}_i, \iota_i \notin \underline{t}\} \cup \\ &\quad \{(\{\iota_1\}, \underline{t}_1 \cup \bar{t}_2) \mid \forall i \in \{1, 2\}, (\{\iota_i\}, \bar{t}_i) \in \mathcal{T}_i\}. \end{aligned}$$

(In the inductive cases, we assume $\mathcal{A}(e_i) = \langle P_i, \mathcal{T}_i, \iota_i, \mathcal{F}_i \rangle$ for $i \in \{1, 2\}$, with $P_1 \cap P_2 = \emptyset$.) *

We prove by induction on e that $\mathcal{A}(e)$ is indeed a safe Petri automaton; for the safety requirement, we add to the induction hypothesis the fact that for any configuration ξ accessible in $\mathcal{A}(e)$, if there is a final configuration $f \in \mathcal{F}$ such that $f \subseteq \xi$, then $f = \xi$. Another invariant is that the initial place never appears in a final configuration, nor in the output of any transition. Note that the place ι_2 becomes unreachable by construction in the cases for union, composition and intersection, so that it could safely be removed, together with the associated transitions.

Theorem 17 (Correctness). *For all expression $e \in \text{Reg}_X^{\sim}$, $\mathcal{L}(\mathcal{A}(e)) = \mathbf{G}(e)$.*

Proof. We prove a stronger result: $\mathcal{G}(\mathcal{A}(e)) = \mathcal{G}(e)$ (up to graph isomorphisms—see Appendix C). This allows us to conclude thanks to Corollary 15. \square

Corollary 18. *The (in)equational theory of Kleene allegories is co-recursively enumerable.*

Proof. Construct Petri automata for the two expressions and enumerate all potential counter-examples, i.e., graphs. A graph


is a counter-example if it can be read in one automaton but not in the other, which is a decidable property. \square

Remark 19. If e is an expression without intersection, converse or 1, it can be shown that the transitions in $\mathcal{A}(e)$ are all of the form $(\{p\}, \{(x, q)\})$, with only one input, one output and a label in X . As a consequence, the accessible configurations are singletons, and the resulting Petri automaton has the structure of a non-deterministic finite-state automaton (NFA). Actually, in that case, the construction we described above is just a variation on Thompson's construction [20], with inlined epsilon transition elimination.

V. COMPARING AUTOMATA

A. Simple automata

The above results hold for the whole syntax of regular expressions with converse and intersection. However, in the remainder of the paper, we have to focus on expressions without converse or identity. This is because in combination with intersection, these two operations introduce cycles in the graphs associated to ground terms. Consider for instance the graphs for $a \cap 1$ and $a \cap b^\sim$:



Since reflexive-transitive closure (\cdot^*) implicitly contains an occurrence of the identity, we also have to forbid this operator. Instead, we add the transitive closure (\cdot^+). We thus work with expressions from $\text{Reg}_X^{\cap, \sim, +}$, defined with the following syntax:

$$e, f \in \text{Reg}_X^{\cap, \sim, +} ::= x \in X \mid e \cap f \mid e \cup f \mid e \cdot f \mid e^+ \mid 0.$$

Accordingly, ground terms are restricted to the following syntax: $u, v, w ::= x \in X \mid w \cdot w \mid w \cap w$.

Automata built using Definition 16 from expressions without converse or unit only have transitions labelled with X . This corresponds to the notion of simple automata.

Definition 20 (Simple Petri automaton)


A Petri automaton $\mathcal{A} = \langle P, \mathcal{T}, \iota, \mathcal{F} \rangle$ is called *simple* if $\forall (\underline{t}, \bar{t}) \in \mathcal{T}, \forall (x, p) \in \bar{t}, x \in X$. $*$

For all $e \in \text{Reg}_X^{\cap, \sim, +}$, $\mathcal{A}(e)$ is simple. Moreover for any run ξ of a simple Petri automaton, $\llbracket \xi \rrbracket = \mathcal{G}(\xi)$ (up to isomorphism); in particular, a simple automaton only produces acyclic graphs.

B. Intuitions

In this section, we show how the notion of simulation relation, that allows to compare NFA, can be adapted to handle simple Petri automata. Consider two automata $\mathcal{A}_1 = \langle P_1, \mathcal{T}_1, \iota_1, \mathcal{F}_1 \rangle$ and $\mathcal{A}_2 = \langle P_2, \mathcal{T}_2, \iota_2, \mathcal{F}_2 \rangle$, we try to show that for any graph G accepted by \mathcal{A}_1 , G is recognised by \mathcal{A}_2 . By Lemma 14, this amounts to proving that for any accepting run ξ in \mathcal{A}_1 , $\mathcal{G}(\xi)$ is recognised by some accepting run ξ' in \mathcal{A}_2 . Leaving non-determinism apart, the first idea that comes to mind is to find a relation between the configurations in \mathcal{A}_1

and the configurations in \mathcal{A}_2 , that satisfy some conditions on the initial and final configurations, and such that if $\xi_k \leq \xi'_k$ and $\xi_k \xrightarrow{t} \xi_{k+1}$, then there is a configuration ξ'_{k+1} in \mathcal{A}_2 such that $\xi_{k+1} \leq \xi'_{k+1}$, $\xi'_k \xrightarrow{t'} \xi'_{k+1}$, and these transition steps are compatible in some sense. However, such a definition will not give us the result we are looking for. Consider these two runs:




The graphs produced by the first and the second runs correspond respectively to the ground terms $a \cap (b \cdot c)$ and $(a \cdot c) \cap b$. These two terms are incomparable, but the relation \leq depicted below satisfies the previously stated conditions.

$$\begin{array}{ccccc} \{A\} & \xrightarrow{1} & \{B, C\} & \xrightarrow{2} & \{B, D\} \\ \downarrow & & \downarrow & & \downarrow \\ \{W\} & \xrightarrow{1'} & \{X, Y\} & \xrightarrow{2'} & \{Y, Z\} \end{array}$$

The problem here is that in Petri automata, runs are token firing games. To adequately compare two runs, we need to closely track the tokens. For this reason, we will relate a configuration ξ_k in \mathcal{A}_1 not only to a configuration ξ'_k in \mathcal{A}_2 , but to a map η_k from ξ'_k to ξ_k . This will enable us to associate with each token situated on some place in P_2 another token placed on \mathcal{A}_1 .

We want to find a reading of $\mathcal{G}(\xi)$ in \mathcal{A}_2 , i.e., a run in \mathcal{A}_2 together with a sequence of maps associating places in \mathcal{A}_2 to positions in $\mathcal{G}(\xi)$. Consider the picture below. Since we already have a reading of $\mathcal{G}(\xi)$ along ξ (by defining $\rho_k(p) = \nu(k, p)$, as in the proof of Lemma 14), it suffices to find maps from the places in \mathcal{A}_2 to the places in \mathcal{A}_1 (the maps η_k): the reading of $\mathcal{G}(\xi)$ in \mathcal{A}_2 will be obtained by composing η_k with ρ_k .




We need to impose some constraints on the maps (η_k) to ensure that $(\rho_k \circ \eta_k)_{0 \leq k \leq n}$ is indeed a correct reading in \mathcal{A}_2 . First, we need to ascertain that a transition t'_k in \mathcal{A}_2 may be fired from the reading state $\rho_k \circ \eta_k$ to reach the reading state $\rho_{k+1} \circ \eta_{k+1}$. Furthermore, as for NFA, we want transitions t_k

and t'_k to be related: specifically, we require t'_k to be included (via the homomorphisms η_k and η_{k+1}) in the transition t_k . This is meaningful because transition t_k contains a lot of information about the vertex k of $\mathcal{G}(\xi)$ and about ρ : the labels of the outgoing edges from k are the labels on the output of t_k , and the only places that will ever be mapped to k in the reading ρ are exactly the places in the input of t_k .

This already shows an important difference between the simulations for NFA and Petri automata. For NFA, we relate a transition $p \xrightarrow{a} p'$ to a transition $q \xrightarrow{a} q'$ with the same label a . Here the transitions $\xi_k \xrightarrow{t_k} \xi_{k+1}$ and $\xi'_k \xrightarrow{t'_k} \xi'_{k+1}$ may have different labels. Consider the step represented on the right, corresponding to a square in the above diagram. The output of $\boxed{0}$ has a label b that does not appear in $\boxed{0'}$, and $\boxed{0'}$ has two outputs labelled by a . Nevertheless this satisfies the conditions informally stated above, indeed, $a \cap b \leq a \cap a$ holds.

However this definition is not yet satisfactory. Consider the two runs below:



Their produced graphs correspond respectively to the ground terms $a \cdot (b \cap c)$ and $(a \cdot b) \cap (a \cdot c)$. The problem is that $a \cdot (b \cap c) \leq (a \cdot b) \cap (a \cdot c)$, but with the previous definition, we cannot relate these runs: they do not have the same length. The solution here consists in grouping the transitions $\boxed{1'}$ and $\boxed{2'}$ together, and consider these two steps as a single step in a *parallel run*. This last modification gives us a notion of simulation that suits our needs.

C. Simulations

Definition 21 (Simulation)

A relation $\leq \subseteq \mathcal{P}(P_1) \times \mathcal{P}(P_2 \rightarrow P_1)$ between the configurations of \mathcal{A}_1 and the partial maps from the places of \mathcal{A}_2 to the places of \mathcal{A}_1 is called a *simulation* between \mathcal{A}_1 and \mathcal{A}_2 if:

- if $\xi \leq E$ and $\eta \in E$ then the range of η must be included in ξ ;
- $\{\iota_1\} \leq \{[\iota_2 \mapsto \iota_1]\}$;
- if $\xi \leq E$ and $\xi \xrightarrow{t_i} \xi'$, then $\xi' \leq E'$ where E' is the set of all η' such that there is some $\eta \in E$ and a compatible set of transitions $T \subseteq \mathcal{T}_2$ such that:
 - $\text{dom}(\eta) \xrightarrow{T} \text{dom}(\eta')$;
 - $\forall (\underline{t}, \underline{t}') \in T, \eta(\underline{t}') \subseteq \underline{t}$ and $\forall (x, q) \in \underline{t}', (x, \eta'(q)) \in \underline{t}$;
 - $\forall p \in \text{dom}(\eta), (\forall (\underline{t}', \underline{t}') \in T, p \notin \underline{t}') \Rightarrow \eta(p) = \eta'(p)$.





Figure 10: Embedding of a parallel run into the run from Figure 5.

- if $\xi \leq E$ and $\xi \in \mathcal{F}_1$, then there must be some $\eta \in E$ such that $\text{dom}(\eta) \in \mathcal{F}_2$. *

We will now prove that the language of \mathcal{A}_1 is contained in the language of \mathcal{A}_2 if and only if there exists such a simulation. We first introduce the following notion of embedding.

Definition 22 (Embedding)

Let $\xi = \langle (\xi_k)_{0 \leq k \leq n}, (t_k)_{0 \leq i < n} \rangle$ be a run in \mathcal{A}_1 , and $\Xi = \langle (\Xi_k)_{0 \leq k \leq n}, (T_i)_{0 \leq i < n} \rangle$ a parallel run in \mathcal{A}_2 . An *embedding* of Ξ into ξ is a sequence $(\eta_i)_{0 \leq i \leq n}$ of maps such that for any $i < n$, we have:

- η_i is a map from Ξ_i to ξ_i ;
- the image of T_i by η_i is included in t_i , meaning that for any $(\underline{t}, \underline{t}') \in T_i$, for any $p \in \underline{t}$ and $(x, q) \in \underline{t}'$, $\eta_i(p)$ is contained in the input of t_i and $(x, \eta_{i+1}(q))$ is in the output of t_i ;
- the image of the tokens in Ξ_i that do not appear in the input of T_i are preserved ($\eta_i(p) = \eta_{i+1}(p)$) and their image is not in the input of t_i .

$$\begin{array}{ccc} \xi_i & \xrightarrow{t_i} & \xi_{i+1} \\ \eta_i \uparrow & & \uparrow \eta_{i+1} \\ \Xi_i & \xrightarrow{T_i} & \Xi_{i+1} \end{array} *$$

Figure 10 illustrates the embedding of some parallel run, producing $G(((b \cdot c \cdot a \cdot b) \cap (b \cdot b \cdot c \cdot a)) \cdot d)$, into the run presented in Figure 5. Notice that it is necessary to have a parallel run instead of a simple one: to find something that matches the second transition in the upper run, we need to fire two transitions in parallel in the lower run.

There is a close relationship between simulations and embeddings:

Lemma 23. *Let \mathcal{A}_1 and \mathcal{A}_2 be two Petri automata, the following are equivalent:*

- 1) *there exists a simulation \leq between \mathcal{A}_1 and \mathcal{A}_2 ;*
- 2) *for any accepting run ξ in \mathcal{A}_1 , there is an accepting parallel run Ξ in \mathcal{A}_2 that can be embedded into ξ .*

Proof. If we have a simulation \leq , let $\xi = \langle (\xi_k)_{0 \leq k \leq n}, (t_k)_{0 \leq k < n} \rangle$ be an accepting run in \mathcal{A}_1 . By the definition of simulation, we can find a sequence of sets of maps $(E_k)_{0 \leq k \leq n}$ such that $E_0 = \{[\iota_2 \mapsto \iota_1]\}$ and $\forall k, \xi_k \leq E_k$. Furthermore, we can extract from this a

sequence of maps $(\eta_k)_{0 \leq k \leq n}$ and a sequence of parallel transitions $(T_k)_{0 \leq k < n}$ such that (η_k) is an embedding of $\langle (\text{dom}(\eta_k))_{0 \leq k \leq n}, (T_k)_{0 \leq k < n} \rangle$ (which is accepting) into ξ . This follows directly from the definitions of embedding and simulation.

On the other hand, if we have property 2., then we can define a relation \leq by saying that $\xi \leq E$ if there is an accepting run $\xi' = \langle (\xi'_k)_{0 \leq k \leq n}, (t_k)_{0 \leq k < n} \rangle$ in \mathcal{A}_1 such that there is an index k_0 : $\xi = \xi'_{k_0}$; and the following holds: $\eta \in E$ if there is an accepting parallel run $\Xi = \langle (\Xi_k)_{0 \leq k \leq n}, (T_k)_{0 \leq k < n} \rangle$ and $(\eta'_k)_{0 \leq k \leq n}$ an embedding of Ξ into ξ such that $\eta = \eta'_{k_0}$. It is then immediate to check that \leq is indeed a simulation. \square

If η is an embedding of Ξ into ξ , we can easily check that $(\rho_i \circ \eta_i)_{0 \leq i \leq n}$ is a parallel reading of $\mathcal{G}(\xi)$ along Ξ in \mathcal{A}_2 . Thus, it is clear that once we have such a run Ξ with the sequence of maps η , we have that $\mathcal{G}(\xi)$ is indeed in the language of \mathcal{A}_2 . The more difficult question is the completeness of this approach: if $\mathcal{G}(\xi)$ is recognised by \mathcal{A}_2 , is it always the case that we can find a run Ξ that may be embedded into ξ ? The answer is affirmative, thanks to Lemma 24 below. If $(\rho_j)_{0 \leq j \leq n}$ is a reading of G along $\xi = \langle (\xi_k)_{0 \leq k \leq n}, (t_k, \bar{t}_k)_{0 \leq k < n} \rangle$, we write $\text{active}(j)$ for the only position in $\rho_j(t_j)^3$. We call \leq a *consistent ordering* on $G = \langle V, E, \iota, o \rangle$ if $\langle \bar{V}, \leq \rangle$ is a linear order and $(p, x, q) \in E$ entails $p \leq q$.

Lemma 24. *Let $G \in \mathcal{L}(\mathcal{A}_2)$ and \leq be any consistent ordering on G . Then there exists a run ξ and a reading $(\rho_j)_{0 \leq j \leq n}$ of G along ξ such that $\forall k, \text{active}(k) \leq \text{active}(k+1)$.*

Proof. The proof of this result is achieved by taking any run ξ accepting G , and then exchanging transitions in ξ according to \leq , while preserving the existence of a reading. The details of this proof being a bit technical, we moved them to Appendix D. \square

(Notice that if G contains cycles, this lemma cannot apply because of the lack of consistent ordering.) Lemma 24 enables us to build an embedding from any reading of $\mathcal{G}(\xi)$ in \mathcal{A}_2 .

Lemma 25. *Let ξ a accepting run of \mathcal{A}_1 . Then $\mathcal{G}(\xi)$ is in $\mathcal{L}(\mathcal{A}_2)$ if and only if there is an accepting parallel run in \mathcal{A}_2 that can be embedded into ξ .*

Proof. The detailed proof of this result can be found in Appendix E.

For the if direction, we build a parallel reading from the embedding, as explained above. For the other direction, we consider a reading of $\mathcal{G}(\xi)$ in \mathcal{A}_2 along some run ξ' . Notice that the natural ordering on \mathbb{N} is consistent for $\mathcal{G}(\xi)$; we may thus change the order of the transitions in ξ' (using Lemma 24) and group them adequately to obtain a parallel reading Ξ that embeds in ξ . \square

So we know that the existence of embeddings is equivalent to the inclusion of languages, and we previously established

³Recall that if $(\rho_j)_{0 \leq j \leq n}$ is a reading along ξ then for any $p, q \in t_j$, we have $\rho_j(p) = \rho_j(q)$.




Figure 11: A Petri automaton for $1 \cap a \cdot b^+ \cdot c$.

that it is also equivalent to the existence of a simulation relation. Hence, the following characterisation holds:

Theorem 26. *Let \mathcal{A}_1 and \mathcal{A}_2 be two simple Petri automata. $\mathcal{L}(\mathcal{A}_1) \subseteq \mathcal{L}(\mathcal{A}_2)$ if and only if there exists a simulation relation \leq between \mathcal{A}_1 and \mathcal{A}_2 .*

Proof. By Lemmas 14, 23 and 25. \square

As Petri automata are finite, there are finitely many relations in $\mathcal{P}(\mathcal{P}(P_1) \times \mathcal{P}(P_2 \rightarrow P_1))$. The existence of a simulation thus is decidable, allowing us to prove the main result:

Theorem 27. *Given two expressions $e, f \in \text{Reg}_X^{\square}$, testing whether $\text{Rel} \models e = f$ is decidable.*

Proof. By Theorems 6, 17 and 26, and reasoning by double inclusion. \square

In practice, we can build the simulation on-the-fly, starting from the pair $(\{\iota_1\}, \{[\iota_2 \mapsto \iota_1]\})$ and progressing from there. We have implemented this algorithm in OCAML [6]. Even though its theoretical worst case time complexity is huge⁴, we get a result almost instantaneously on simple one-line examples.

D. The problems with converse or unit

The previous algorithm is not complete in presence of converse or unit. More precisely, Lemma 25 does not hold for general automata. Indeed, it is not possible to compare two runs just by relating the tokens at each step, and checking each transition independently. Consider the automaton from Figure 11. This automaton has in particular an accepting run recognising $1 \cap abc$. Let us try to test if this is smaller than the following runs from another automaton (we represent the transitions simply as arrows, because they only have a single input and a single output):

$$\begin{array}{cccccccccccccccc} x_0 & \xrightarrow{a} & x_1 & \xrightarrow{b} & x_2 & \xrightarrow{c} & x_3 & \xrightarrow{a} & x_4 & \xrightarrow{b} & x_5 & \xrightarrow{c} & x_6 \\ y_0 & \xrightarrow{a} & y_1 & \xrightarrow{b} & y_2 & \xrightarrow{c} & y_3 & \xrightarrow{a} & y_4 & \xrightarrow{b} & y_5 & \xrightarrow{b} & y_6 & \xrightarrow{c} & y_7 \end{array}$$

It stands to reason that we would reach a point where for the first run:

$$\{D, E\} \leq \{[x_3 \mapsto D]\}$$

and for the second run:

$$\{D, E\} \leq \{[y_3 \mapsto D]\}.$$

⁴A quick analysis gives a $\mathcal{O}(2^{n+(n+1)^m})$ complexity bound, where n and m are the numbers of places of the automata.

So if it were possible to relate the end of the runs just with this information, they should both be bigger than $1 \cap abc$ or both smaller or incomparable. But in fact the first run (recognising $abcabc$) is bigger than $1 \cap abc$ but the second (recognising $abcabbc$) is not. This highlights the need for having some memory of previously fired transition when trying to compare runs of general Petri automata, thus preventing our local approach to bear fruits. The same kind of example could be found with the converse operation instead of 1.

VI. COMPLEXITY

The previous notion of simulation actually allows us to decide language inclusion of simple automata in EXPSPACE. We actually obtain that the problem is EXPSPACE-complete.

Lemma 28. *Comparing simple Petri automata is EXPSPACE-easy.*

Proof. Our measure for the size of an automaton here is its number of places (the number of transitions is at most exponential in this number). Here is a non-deterministic semi-algorithm that tries to refute the existence of a simulation relation between \mathcal{A}_1 and \mathcal{A}_2 .

- 1: start with $\xi := \{\iota_1\}$ and $E := \{[\iota_2 \mapsto \iota_1]\}$;
- 2: if $\xi \in \mathcal{F}_1$, check if there is some $\eta \in E$ such that $\text{dom}(\eta) \in \mathcal{F}_2$, if not return FALSE;
- 3: choose non-deterministically a transition $(\underline{t}, \bar{t}) \in \mathcal{T}_1$ such that $\underline{t} \subseteq \xi$;
- 4: fire (\underline{t}, \bar{t}) , which means that

$$\xi := \xi \setminus \underline{t} \cup \{p \in P_1 \mid \exists x \in X : (x, p) \in \bar{t}\};$$

- 5: have E progress along (\underline{t}, \bar{t}) as well, according to the conditions from Definition 21.
- 6: go to step 2.

All these computations can be done in exponential space. In particular as ξ is a set of places in P_1 , it can be stored in space $|P_1| \times \log(|P_1|)$. Similarly, E , being a set of partial functions from P_2 to P_1 , each of which of size $|P_2| \times \log(|P_1| + 1)$, can be stored in space $|P_1 + 1|^{|P_2|} \times |P_2| \times \log(|P_1| + 1)$. This non-deterministic EXPSPACE semi-algorithm can then be turned into an EXPSPACE algorithm by Savitch' theorem [19]. \square

One can check that the number of places in $\mathcal{A}(e)$ is linear in the size of e . (The exponential upper-bound on the number of transitions is asymptotically reached, consider for instance the automaton for $(x_1 \cup y_1) \cap (x_2 \cup y_2) \cap \dots \cap (x_n \cup y_n)$.) Therefore, the previous Lemma gives us a EXPSPACE algorithm for deciding the (in)equational theory of identity-free relational Kleene lattices.

Theorem 29. *Comparing simple Petri automata is EXPSPACE-complete.*

Proof. By Lemma 28, it suffices to show hardness. We perform a reduction from the equality of languages denoted by regular expressions with squaring ($e^2 := e \cdot e$), an EXPSPACE-complete problem [14]. To avoid confusion, the regular language denoted by the expression e will be written $\langle e \rangle$. Any




Figure 12: Squaring of the automaton for e_1 .

word u can be seen uniquely as a linear graph $\lambda(u)$. By extension, the set of graphs of words from $\langle e \rangle$ will be denoted by $\lambda(\langle e \rangle)$.

First, notice that if u and v are just words over X , $\lambda(u) \triangleleft \lambda(v)$ is equivalent to $u = v$. Because of this, it is straightforward to check that for any $e, f \in \text{Reg}_X^2$ the following holds

$$\triangleleft(\lambda(\langle e \rangle)) = \triangleleft(\lambda(\langle f \rangle)) \Leftrightarrow \langle e \rangle = \langle f \rangle. \quad (8)$$

Given an expression e on this signature, we can build in linear time a Petri automaton \mathcal{A} , with a linear number of places and transitions. The closure of the language denoted by e is exactly the language recognised by \mathcal{A} . This automaton is not simple: some outgoing arcs are labelled with 1; therefore, this reduction only ensures that the equivalence of arbitrary Petri automata is EXPSPACE-hard. To get EXPSPACE-hardness for simple Petri automata, we need to refine the construction so that 1 is interpreted as a standard letter. (More details are given in Appendix F.)

The automata we produce here only have one final configuration, consisting in a singleton. The construction is a straightforward adaptation of Thompson's algorithm for NFA [20]. The only interesting case is for computing an automaton for $e = (e_1)^2$. We represent it graphically in Figure 12. The transitions labelled by X are a shorthand for a set of transitions, containing for each letter x in X a transition with one output, labelled by x . This construction is linear: the automaton for e_1 is not copied. Furthermore, a run in this automaton will start by sending one token in \circ and one in ι_1 , the initial state of the automaton for e_1 . Then it will perform a run in this automaton until a single token reaches the final state for e_1 , fn_1 . At this point the tokens from \circ and fn_1 will be sent to \bullet and ι_1 , starting a new run of e_1 . When a token is finally sent to fn_1 , it can be consumed together with the one in \bullet , to reach the final configuration. \square

This proof does not allow us to deduce that also the (in)equational theory of Kleene allegories is EXPSPACE-hard: the Petri automata we construct are not associated to some expressions of polynomial size, a priori.

VII. RELATIONSHIP WITH STANDARD PETRI NET NOTIONS

Our notion of Petri automaton is really close to the standard notion of labelled (safe) Petri net, where the transitions themselves are labelled, rather than their outputs. We motivate this

design choice, and we relate some of the notions we introduced to the standard ones [16].

Any Petri automaton can be translated into a safe Petri net whose transitions are labelled by $\bar{X} \uplus \{\tau\}$, the additional label τ standing for silent actions. For each automaton transition $(\{p_1, \dots, p_n\}, \{(x_1, q_1), \dots, (x_m, q_m)\})$ with $m > 1$, we introduce m fresh places r_1, \dots, r_m and $m + 1$ transitions:

- a silent transition t_0 with preset $\{p_1, \dots, p_n\}$ and postset $\{r_1, \dots, r_m\}$;
- and for each $1 \leq k \leq m$ a transition t_k labelled by x_k , with preset $\{r_k\}$ and postset $\{q_k\}$.

The inductive construction from Section IV is actually simpler to write using labelled Petri nets, as one can freely use τ -labelled transitions to assemble automata into larger ones, one does not need to perform the τ -elimination steps on the fly.

On the other side, we could not define an appropriate notion of simulation for Petri nets: we need to fire several transitions at once in the small net, to provide enough information for the larger net to answer; delimiting which transitions to group and which to separate is non-trivial; similarly, defining a notion of parallel step is delicate in presence of τ -transitions. By switching to our notion of Petri automata, we impose strong constraints about how those τ -transitions should be used, resulting in a more fitted model.

To describe a run in a Petri net N , one may use a *process* $p : K \rightarrow N$, where K is an *occurrence net* (a partially ordered Petri net) [9]. The graphical representation (Figures 5, 7 and 10) we used to describe runs in an automaton are in fact a mere adaptation of this notion to our setting (with labels on arcs rather than on transitions).

Our notion $[[\xi]]$ of trace of a run actually corresponds to the standard notion of *pomset-trace*, via dualisation (see Appendix H). Jategaonkar and Meyer showed that the pomset-trace equivalence problem for safe Petri nets is EXPSpace-complete [10]. However this equivalence is too strong and does not coincide with the one discussed in the present paper, even for simple Petri automata. The graph produced by a run is its trace in this case, so that pomset-trace equivalence for Petri nets corresponds to equivalence of the sets of graphs produced by Petri automata $(\mathcal{G}(\mathcal{A}) = \mathcal{G}(\mathcal{B}))$, up to graph isomorphism). However, for the equational theory we consider, we need to compare the languages, which are downward-closed sets of graphs, $({}^*\mathcal{G}(\mathcal{A}) = {}^*\mathcal{G}(\mathcal{B}))$, i.e., $\mathcal{L}(\mathcal{A}) = \mathcal{L}(\mathcal{B})$) rather than the sets of graphs themselves.

Also note that the class of sets of graphs produced by Petri automata $(\{\mathcal{G}(\mathcal{A}) \mid \mathcal{A} \text{ a Petri automaton}\})$ is not closed under downward closure. Intuitively, the width of any graph in $\mathcal{G}(\mathcal{A})$ is bounded by the number number of places of \mathcal{A} , but ${}^*\mathcal{G}(\mathcal{A})$ usually contains graphs of arbitrary width. As a consequence, one cannot easily reduce our problem to pomset-trace equivalence of safe Petri nets.

VIII. DIRECTIONS FOR FUTURE WORK

The automata model we introduced to recognise downward-closed graph languages allowed us to obtain the decidability

of identity-free relational Kleene lattices, an (in)equational theory studied recently by Andr eka et al. [1]. Thanks to this model, we also obtained that the (in)equational theory of Kleene allegories is co-recursively enumerable. We leave several questions open.

First, is the (in)equational theory of Kleene allegories decidable? Two approaches could provide an affirmative answer: finding an algorithm for comparing arbitrary safe Petri automata, or finding a complete and recursively enumerable axiomatisation.

Second, can we obtain a Kleene-like theorem for Petri automata: is the language of any Petri automaton also the language of a Kleene allegory term? This question can be restricted to simple Petri automata; if one of the answers is negative, is there an algebraic way of representing these automata? Which are the missing operators?

REFERENCES

- [1] H. Andr eka, S. Mikul as, and I. N emeti. The equational theory of Kleene lattices. *Theoretical Computer Science*, 412(52):7099–7108, 2011.
- [2] H. Andr eka and D. Bredikhin. The equational theory of union-free algebras of relations. *Algebra Universalis*, 33(4):516–532, 1995.
- [3] S. L. Bloom, Z.  sik, and G. Stefanescu. Notes on equational theories of relations. *Algebra Universalis*, 33(1):98–126, 1995.
- [4] M. Boffa. Une condition impliquant toutes les identit es rationnelles. *Informatique Th eorique et Applications*, 29(6):515–518, 1995.
- [5] P. Brunet and D. Pous. Kleene algebra with converse. In *Proc. RAMiCS*, volume 8428 of *Lecture Notes in Computer Science*, pages 101–118. Springer Verlag, 2014.
- [6] P. Brunet and D. Pous. Web appendix to this abstract, 2014. <http://perso.ens-lyon.fr/paul.brunet/rklm.html>.
- [7] J. H. Conway. *Regular algebra and finite machines*. Chapman and Hall Mathematics Series, 1971.
- [8] P. J. Freyd and A. Scedrov. *Categories, Allegories*. North Holland, 1990.
- [9] U. Goltz and W. Reisig. The non-sequential behaviour of petri nets. *Information and Control*, 57(2):125–147, 1983.
- [10] L. Jategaonkar and A. R. Meyer. Deciding true concurrency equivalences on safe, finite nets. *Theoretical Computer Science*, 154(1):107 – 143, 1996. Twentieth International Colloquium on Automata, Languages and Programming.
- [11] S. C. Kleene. *Representation of Events in Nerve Nets and Finite Automata*. Memorandum. Rand Corporation, 1951.
- [12] D. Kozen. A completeness theorem for Kleene Algebras and the algebra of regular events. In *Proc. LICS*, pages 214–225. IEEE Computer Society, 1991.
- [13] D. Krob. A Complete System of B-Rational Identities. In *Proc. ICALP*, volume 443 of *Lecture Notes in Computer Science*, pages 60–73. Springer Verlag, 1990.
- [14] A. Meyer and L. Stockmeyer. The equivalence problem for regular expressions with squaring requires exponential space. In *Switching and Automata Theory, 1972., IEEE Conference Record of 13th Annual Symposium on*, pages 125–129. IEEE, 1972.
- [15] A. Meyer and L. J. Stockmeyer. Word problems requiring exponential time. In *Proc. ACM symposium on Theory of computing*, pages 1–9. ACM, 1973.
- [16] T. Murata. Petri nets: Properties, analysis and applications. *Proc. of the IEEE*, 77(4):541–580, Apr 1989.
- [17] C. A. Petri. Fundamentals of a theory of asynchronous information flow. In *IFIP Congress*, pages 386–390, 1962.
- [18] C. A. Petri. *Kommunikation mit Automaten*. PhD thesis, Darmstadt Univ. of Tech., 1962.
- [19] W. J. Savitch. Relationships between nondeterministic and deterministic tape complexities. *Journal of computer and system sciences*, 4(2):177–192, 1970.
- [20] K. Thompson. Regular expression search algorithm. *C. of the ACM*, 11:419–422, 1968.

A. Omitted proof: Theorem 6

We denote by W_X the set of ground terms. We want to establish that for any regular expressions with intersection and converse e and f , the following are equivalent:

- (i) $\text{Rel} \models e \leq f$,
- (ii) $\llbracket e \rrbracket \subseteq \triangleleft \llbracket f \rrbracket$,
- (iii) $\mathcal{G}(e) \subseteq \triangleleft \mathcal{G}(f)$.

a) (iii) \Rightarrow (ii): Suppose $\mathcal{G}(e) \subseteq \triangleleft \mathcal{G}(f)$, let $u \in \llbracket e \rrbracket$ be a ground term. Necessarily there is some graph $G \in \mathcal{G}(f)$ such that $G(u) \blacktriangleleft G$. By definition of $\mathcal{G}(f)$, there is some term $v \in \llbracket f \rrbracket$ such that $G = G(v)$. This means that $u \triangleleft v$, thus proving that $u \in \triangleleft \llbracket f \rrbracket$.

b) (ii) \Rightarrow (i): Let $e, f \in \text{Reg}_X^{\cap}$ two expressions such that $\llbracket e \rrbracket \subseteq \triangleleft \llbracket f \rrbracket$, and $\sigma : X \rightarrow \mathcal{P}(S \times S)$ some relational interpretation.

$$\begin{aligned} \widehat{\sigma}(e) &= \bigcup_{w \in \llbracket e \rrbracket} \widehat{\sigma}(w) && \text{(Lemma 5)} \\ &\subseteq \bigcup_{w \in \triangleleft \llbracket f \rrbracket} \widehat{\sigma}(w) && (\llbracket e \rrbracket \subseteq \triangleleft \llbracket f \rrbracket) \\ &= \widehat{\sigma}(f) && \text{(Lemma 5)} \end{aligned}$$

c) (i) \Rightarrow (iii): In order to prove this last implication, we need the following lemma, which is due to Andr eka and Bredekhn.

Lemma 30 ([2, Lemma 3]). *Let S be a base set, $i, j \in S$, $v \in W_X$, $G(v) = \langle V_v, E_v, \iota_v, o_v \rangle$ and $\sigma : X \rightarrow \mathcal{P}(S \times S)$. The following are equivalent:*

- 1) $(i, j) \in \widehat{\sigma}(v)$;
- 2) $\exists \varphi : V_v \rightarrow S$ such that $\varphi(\iota_v) = i$; $\varphi(o_v) = j$ and $(p, a, q) \in E_v \Rightarrow (\varphi(p), \varphi(q)) \in \sigma(a)$.

Let $e, f \in \text{Reg}_X^{\cap}$ two expressions such that $\text{Rel} \models e \leq f$, and $u \in \llbracket e \rrbracket$ such that $G(u) = \langle V_u, E_u, \iota_u, o_u \rangle$; we can build an interpretation $\sigma : X \rightarrow \mathcal{P}(V_u \times V_u)$ by specifying:

$$\sigma(a) := \{(p, q) \mid (p, a, q) \in E_u\}.$$

It is quite simple to check that $\widehat{\sigma}(u) = \{(\iota_u, o_u)\}$. By Lemma 5 and $\text{Rel} \models e \leq f$, we know that

$$\widehat{\sigma}(u) \subseteq \widehat{\sigma}(f) = \bigcup_{v \in \llbracket f \rrbracket} \widehat{\sigma}(v).$$

Thus there is some $v \in \llbracket f \rrbracket$ such that $(\iota_u, o_u) \in \widehat{\sigma}(v)$. By Lemma 30 we get that there is a map $\varphi : V_v \rightarrow V_u$ such that $\varphi(\iota_v) = \iota_u$; $\varphi(o_v) = o_u$ and

$$(p, a, q) \in E_v \Rightarrow (\varphi(p), \varphi(q)) \in \sigma(a).$$

Using the definition of σ , we rewrite this last condition as

$$(p, a, q) \in E_v \Rightarrow (\varphi(p), a, \varphi(q)) \in E_u.$$

Thus φ is a graph homomorphism from $G(v)$ to $G(u)$, proving that $G(u) \blacktriangleleft G(v)$, hence $G(u) \in \triangleleft \mathcal{G}(f)$.

B. Omitted proof: Lemma 14

Let $G = \langle V, E, \iota, o \rangle$, $\llbracket \xi \rrbracket = \langle \{0, \dots, n\}, E_\xi, 0, n \rangle$, and $\mathcal{G}(\xi) = \langle \{[i]_\xi \mid 0 \leq i \leq j\}, E'_\xi, [0]_\xi, [n]_\xi \rangle$. Suppose there exists a graph homomorphism φ from $\mathcal{G}(\xi)$ to G . We build a reading $(\rho_k)_k$ of G along ξ by letting $\rho_k(p) := \varphi([\nu(k, p)]_\xi)$ for $0 \leq k \leq n$ and $p \in \xi_k$. We now have to check that ρ is truly a reading of G in \mathcal{A} :

- for the initialisation and conclusion of the reading:

$$\begin{aligned} \rho_0(\iota_{\mathcal{A}}) &= \varphi([\nu(0, \iota_{\mathcal{A}})]_\xi) && \text{(by definition)} \\ &= \varphi([0]_\xi) = \{\iota\} && (\varphi \text{ is a homomorphism}) \end{aligned}$$

$$\begin{aligned} p \in \xi_n, \rho_n(p) &= \varphi([\nu(n, p)]_\xi) \\ &= \varphi([n]_\xi) && (\forall k, p, k \leq \nu(k, p) \leq n) \\ &= \{o\}. && (\varphi \text{ is a homomorphism}) \end{aligned}$$

- for all $p \in \underline{t}_k$, $\rho_k(p) = \varphi([\nu(k, p)]_\xi) = \varphi([k]_\xi)$ which does not depend on p .
- for all $p \in \xi_k \setminus \underline{t}_k$, we have $\nu(k, p) = \nu(k+1, p)$ (since $p \notin \underline{t}_k$). Hence

$$\begin{aligned} \rho_k(p) &= \varphi([\nu(k, p)]_\xi) = \varphi([\nu(k+1, p)]_\xi) \\ &= \rho_{k+1}(p). \end{aligned}$$

- for all $p \in \underline{t}_k$ and $(x, q) \in \overline{t}_k$, we know that $\rho_k(p) = \varphi([k]_\xi)$ and that $(k, x, \nu(k+1, q)) \in E_\xi$.
 - If $x \in X$, we also have $([k]_\xi, x, [\nu(k+1, q)]_\xi) \in E'_\xi$. Because φ is a homomorphism we can deduce that:

$$(\varphi([k]_\xi), x, \varphi([\nu(k+1, q)]_\xi)) \in E,$$

which can be rewritten $(\rho_k(p), x, \rho_{k+1}(q)) \in E$.

- If $x = y^\sim$, $y \in X$, we also have $([\nu(k+1, q)]_\xi, y, [k]_\xi) \in E'_\xi$. Because φ is a homomorphism we can get like before $(\rho_{k+1}(q), y, \rho_k(p)) \in E$.
- If finally $x = 1$, then we know that $k \equiv_\xi \nu(k+1, q)$, thus proving that

$$\rho_k(p) = \varphi([k]_\xi) = \varphi([\nu(k+1, q)]_\xi) = \rho_{k+1}(q).$$

If on the other hand we have a reading $(\rho_k)_{0 \leq k \leq n}$ of G , we define $\varphi : \{0, \dots, n\} \rightarrow V$ by $\varphi([k]_\xi) := \rho_k(p)$ for any $p \in \underline{t}_k$. As $(\rho_k)_k$ is a reading, φ is well defined⁵. Let us check that φ is a homomorphism from $\mathcal{G}(\xi)$ to G :

- $\varphi([0]_\xi) = \rho_0(\iota_{\mathcal{A}}) = \iota$;
- $\varphi([n]_\xi) = \rho_n(p)$ with $p \in \xi_n$, and since $(\rho_k)_k$ is a reading $\rho_n(p) = o$.
- if $([k]_\xi, x, [l]_\xi) \in E'_\xi$ is an edge of $\mathcal{G}(\xi)$, then it was produced from some edge $(i, y, j) \in E_\xi$, with either $x = y$ and $i, j \in [k]_\xi \times [l]_\xi$ or $y = x^\sim$ and $i, j \in [l]_\xi \times [k]_\xi$. There is some $p \in \underline{t}_i$ and q such that $(y, q) \in \underline{t}_j$ and $j = \nu(i+1, q)$.

By definition of ν we know that $\forall i+1 \leq m < j, q \notin \underline{t}_m$. Thus, because $(\rho_k)_k$ is a reading, $\rho_{i+1}(q) = \rho_j(q)$ and either $(\rho_i(p), x, \rho_{i+1}(q)) \in E$ or $(\rho_{i+1}(q), x, \rho_i(p)) \in E$, and thus $(\varphi([k]_\xi), x, \varphi([l]_\xi)) \in E$.

⁵It is not difficult to check that $k \equiv_\xi l \Rightarrow \forall p, q \in \underline{t}_k \times \underline{t}_l, \rho_k(p) = \rho_l(q)$.

C. Omitted proof: Lemma 17

We have to prove that $\mathcal{G}(\mathcal{A}(e)) \equiv \mathcal{G}(e)$, where \equiv denotes equality of sets of graphs, up to graph isomorphisms. We proceed by induction on $e \in \text{Reg}_X^{\cup}$ (with the restrictions introduced in Section IV: $\cdot^+ \cup 1$ instead of \cdot^* and the converse only on variables); we recall the inductive definition of $\mathcal{A}(e)$ (Definition 16), giving the arguments for the correctness for each inductive case.

- $\mathcal{A}(x) := \langle \{0, 1\}, \{(\{0\}, \{(x, 1)\})\}, 0, \{1\} \rangle$



Proof. The only accepting run here is

$$\xi = (\{0\} \xrightarrow{\{(\{0\}, \{(x, 1)\})\}} \{1\}),$$

. If $x \in X$, then $\mathcal{G}(\xi) = \bullet \xrightarrow{x} \bullet \rightarrow$. If $x = y^{\cup}$ then $\mathcal{G}(\xi) = \bullet \xrightarrow{y} \bullet \rightarrow$. If finally $x = 1$ then $\mathcal{G}(\xi) = \bullet \rightarrow$. In every case this is the graph of the only ground term in $\llbracket x \rrbracket$. \square

- $\mathcal{A}(0) := \langle \{0\}, \emptyset, 0, \emptyset \rangle$

Proof. As there are no final configuration, no run is accepting in this automaton. Hence $\mathcal{G}(\mathcal{A}(0)) = \emptyset = \mathcal{G}(0)$. \square

- $\mathcal{A}(e_1 \cup e_2) := \langle P_1 \cup P_2, \mathcal{T}, \iota_1, \mathcal{F}_1 \cup \mathcal{F}_2 \rangle$ with $\mathcal{T} = \mathcal{T}_1 \cup \mathcal{T}_2 \cup \{(\{t_1\}, \bar{t}) \mid (\{t_2\}, \bar{t}) \in \mathcal{T}_2\}$.

Proof. If $\xi = \langle (\xi_k)_{0 \leq k \leq n}, (t_k)_{0 \leq k < n} \rangle$ is an accepting run in $\mathcal{A}(e_1 \cup e_2)$, then $\xi_0 = \{t_1\}$ and either $\forall k, t_k \in \mathcal{T}_1$ and ξ is accepting in \mathcal{A}_1 , or $t_0 = (\{t_1\}, \bar{t})$ and $\langle (\{t_2\}, \xi_1, \dots, \xi_n), (\{t_2\}, \bar{t}), t_1, \dots, t_{n-1} \rangle$ is an accepting run in \mathcal{A}_2 . Hence

$$\mathcal{G}(\mathcal{A}(e_1 \cup e_2)) = \mathcal{G}(\mathcal{A}(e_1)) \cup \mathcal{G}(\mathcal{A}(e_2)),$$

thus proving that:

$$\begin{aligned} \mathcal{G}(e_1 \cup e_2) &= \mathcal{G}(e_1) \cup \mathcal{G}(e_2) \\ &\equiv \mathcal{G}(\mathcal{A}(e_1)) \cup \mathcal{G}(\mathcal{A}(e_2)) \\ &\equiv \mathcal{G}(\mathcal{A}(e_1 \cup e_2)) \end{aligned} \quad \square$$

- $\mathcal{A}(e \cdot f) := \langle P_1 \cup P_2, \mathcal{T}, \iota_1, \mathcal{F}_2 \rangle$ with $\mathcal{T} = \mathcal{T}_1 \cup \mathcal{T}_2 \cup \{(f, \bar{t}) \mid f \in \mathcal{F}_1 \text{ and } (\{t_2\}, \bar{t}) \in \mathcal{T}_2\}$.

Proof. If $\xi = \langle (\xi_k)_{0 \leq k \leq n}, (t_k)_{0 \leq k < n} \rangle$ is an accepting run in $\mathcal{A}(e_1 \cdot e_2)$, then $\xi_0 = \{t_1\}$. There must also be some m such that:

- $\xi^1 = \langle (\xi_k)_{0 \leq k \leq m}, (t_k)_{0 \leq k < m} \rangle$ is accepting in $\mathcal{A}(e_1)$, in particular $\xi_m \in \mathcal{F}_1$;
- $t_m = (\xi_m, \bar{t})$, with $(\{t_2\}, \bar{t}) \in \mathcal{T}_2$;
- the run $\xi^2 = \langle (\{t_2\}, \xi_{m+1}, \dots, \xi_n), (\{t_2\}, \bar{t}), t_{m+1}, \dots, t_{n-1} \rangle$ is accepting in $\mathcal{A}(e_2)$.

Furthermore, it can be shown that $\mathcal{G}(\xi) \equiv \mathcal{G}(\xi^1) \cdot \mathcal{G}(\xi^2)$. The isomorphism is:

$$\begin{aligned} \pi : \mathcal{G}(\xi) &\rightarrow \mathcal{G}(\xi^1) \cdot \mathcal{G}(\xi^2) \\ [k]_{\xi} &\mapsto [k]_{\xi^1} \in V_{\xi^1} \text{ if } k < m \\ [k]_{\xi} &\mapsto [k-m]_{\xi^2} \in V_{\xi^2} \text{ if } k \geq m \end{aligned}$$

Thus we can deduce that:

$$\begin{aligned} \mathcal{G}(e_1 \cdot e_2) &= \mathcal{G}(e_1) \cdot \mathcal{G}(e_2) \\ &\equiv \mathcal{G}(\mathcal{A}(e_1)) \cdot \mathcal{G}(\mathcal{A}(e_2)) \\ &\equiv \mathcal{G}(\mathcal{A}(e_1 \cdot e_2)) \end{aligned} \quad \square$$

- $\mathcal{A}(e_1^+) := \langle P_1, \mathcal{T}, \iota_1, \mathcal{F}_1 \rangle$ with $\mathcal{T} = \mathcal{T}_1 \cup \{(f, \bar{t}) \mid f \in \mathcal{F}_1 \text{ and } (\{t_1\}, \bar{t}) \in \mathcal{T}_1\}$.

Proof. The proof is extremely similar to the case of $e_1 \cdot e_2$, simply with multiple accepting runs separated by transitions in (f, \bar{t}) with f final and $(\{t_1\}, \bar{t}) \in \mathcal{T}_1$. Note that the construction does not add any transitions ending in ι_1 . \square

- $\mathcal{A}(e_1 \cap e_2) := \langle P_1 \cup P_2, \mathcal{T}, \iota_1, \mathcal{F} \rangle$ with $\mathcal{F} := \{f_1 \cup f_2 \mid f_1 \in \mathcal{F}_1, f_2 \in \mathcal{F}_2\}$ and

$$\begin{aligned} \mathcal{T} &:= \{(\underline{t}, \bar{t}) \mid i \in \{1, 2\}, (t, \bar{t}) \in \mathcal{T}_i, \iota_i \notin \underline{t}\} \\ &\cup \{(\{t_1\}, \bar{t}_1 \cup \bar{t}_2) \mid (\{t_i\}, \bar{t}_i) \in \mathcal{T}_i, \forall i \in \{1, 2\}\} \end{aligned}$$

Proof. Let $\xi = \langle (\xi_k)_{0 \leq k \leq n}, (t_k)_{0 \leq k < n} \rangle$ be an accepting run in $\mathcal{A}(e_1 \cap e_2)$, then $(\xi_k)_{0 \leq k \leq n} = \{t_1\}$ and $t_0 = (\{t_1\}, \bar{t}_1 \cup \bar{t}_2)$ with $(\{t_i\}, \bar{t}_i) \in \mathcal{T}_i, i \in \{1, 2\}$. For $i \in \{1, 2\}$, define $\pi_i(\xi) := \xi \setminus P_{2-i}$, and $\delta_i(t) := \{t\}$ if $t \in \mathcal{T}_i$, and \emptyset otherwise.

We check that the parallel run ξ^i defined by the sequence of configurations $(\{t_i\}, \pi_i(\xi_1), \dots, \pi_i(\xi_n))$ and the sequence of transitions $(\{t_i\}, \bar{t}_i, \delta_i(t_1), \dots, \delta_i(t_{n-1}))$ is an accepting parallel run in $\mathcal{A}(e_i)$, and that

$$\mathcal{G}(\xi) \equiv \mathcal{G}(\xi^1) \cap \mathcal{G}(\xi^2),$$

which allows us to conclude.

(Although we did not properly introduce the graph produced by a parallel run, this does not hold any difficulty. As stated when we introduced the notion of parallel transition, such a transition may be unfolded by sequentially firing the underlying transitions, in any order. Accordingly, all of the unfoldings of a parallel run produce the same graph (up to isomorphisms). We may thus take for the graph produced by a parallel run any of those graphs.) \square

D. Omitted proof: Lemma 24

In the following, let $\mathcal{A} = \langle P, \mathcal{T}, \iota, \mathcal{F} \rangle$ be a Petri automaton, and $\xi = \langle (\xi_k)_{0 \leq k \leq n}, (t_k)_{0 \leq k < n} \rangle$ a run of \mathcal{A} with $t_k = (\underline{t}_k, \bar{t}_k)$ for every k .

Definition 31 (Exchangeable transitions)

Two transitions t_k and t_{k+1} are *exchangeable* in ξ if for all $p \in \xi_k$, p is in \underline{t}_{k+1} implies that there is no $x \in X$ such that $(x, p) \in \bar{t}_k$. \ast

Lemma 32. Suppose t_k and t_{k+1} are exchangeable for some $0 \leq k < n-1$. We write $C' = \xi_k \setminus \overline{t_{k+1}} \cup \{p \mid \exists x : (x, p) \in \overline{t_{k+1}}\}$. Then $\xi_k \xrightarrow{t_{k+1}}_{\mathcal{A}} C'$ and $C' \xrightarrow{t_k}_{\mathcal{A}} \xi_{k+2}$. Furthermore, for any graph G , if $G \in \mathcal{L}(\xi)$, then $G \in \mathcal{L}(\xi[k \leftrightarrow k+1])$, where:

$$\xi[k \leftrightarrow k+1] := \langle (\xi_0, \xi_1, \dots, \xi_k, C', \xi_{k+2}, \dots, \xi_n), \\ (t_0, t_1, \dots, t_{k+1}, t_k, \dots, t_{n-1}) \rangle$$

Proof. The fact that $\xi_k \xrightarrow{t_{k+1}}_{\mathcal{A}} C'$ and $C' \xrightarrow{t_k}_{\mathcal{A}} \xi_{k+2}$ is trivial to check, with the definition of exchangeable.

Let $(\rho_j)_{0 \leq j \leq n}$ be a reading of G along ξ . If $(\rho'_j)_{0 \leq j \leq n}$ is defined by:

$$\rho'_j(p) = \begin{cases} \rho_j(p) & \text{if } j \neq k+1, \\ \rho_{k+2}(p) & \text{if } j = k+1 \text{ and } (x, p) \in \overline{t_{k+1}} \text{ for some } x, \\ \rho_k(p) & \text{otherwise.} \end{cases}$$

Then $(\rho'_j)_{0 \leq j \leq n}$ is a reading of G along $\xi[k \leftrightarrow k+1]$. \square

Recall that if $(\rho_j)_{0 \leq j \leq n}$ is a reading of G along ξ we write $active(j)$ for the only position in $\rho_j(\underline{t}_j)$.

Lemma 33. Let \leq be any consistent ordering (a linear ordering compatible with the directed edges of the graph) on G . If $(\rho_j)_{0 \leq j \leq n}$ is a reading of G along ξ , and if $active(k+1) \leq active(k)$ for some k , then t_k and t_{k+1} are exchangeable.

Proof. Let $G = \langle V, E, \iota, o \rangle$. As $(\rho_i)_{0 \leq i \leq n}$ is a reading, for any $(x, p) \in \overline{t_k}$, $(active(k), x, \rho_{k+1}(p)) \in E$, thus

$$active(k) < \rho_{k+1}(p).$$

We know that $active(k+1) \leq active(k)$, meaning by transitivity that $active(k+1) < \rho_{k+1}(p)$. Hence

$$active(k+1) \neq \rho_{k+1}(p)$$

and because $(\rho_i)_{0 \leq i \leq n}$ is a reading we can infer that $p \notin \overline{t_{k+1}}$, thus proving that t_k and t_{k+1} are exchangeable. \square

Proof of Lemma 24. Because G is in $\mathcal{L}(\mathcal{A})$, we can find a reading ρ' along some run ξ . If that reading is not in the correct order, then by Lemma 33 we can exchange two transitions and Lemma 32 ensures that we can find a corresponding reading. We repeat this process until we get a reading in the correct order. \square

E. Omitted proof: Lemma 25

Let $\xi = \langle (\xi_k)_{0 \leq k \leq n}, (t_k)_{0 \leq k < n} \rangle$ be a run, with $t_k = (\underline{t}_k, \overline{t}_k)$ for all k . Define $\rho_k^1(p) := \nu(k, p)$ for all index k place $p \in \xi_k$. We have that $(\rho_k^1)_{0 \leq k \leq n}$ is a reading of $\mathcal{G}(\xi)$ along ξ .

- Assume an embedding $(\eta_k)_{0 \leq k \leq n}$ of an accepting parallel run Ξ into ξ . We define a parallel reading (ρ_k^2) of $\mathcal{G}(\xi)$ in \mathcal{A}_2 by letting $\rho_k^2(p) := \rho_k^1(\eta_k(p))$.
- On the other hand, notice that the natural ordering on \mathbb{N} is consistent for $\mathcal{G}(\xi)$, and that $\forall 0 \leq k < n$, $\rho_k^1(t_k) = \{k\}$. By Lemma 24 we gather that $\mathcal{G}(\xi)$ is in $\mathcal{L}(\mathcal{A}_2)$ if and only if there exists a reading $(\rho_j^2)_{0 \leq j \leq n'}$ of $\mathcal{G}(\xi)$ along some run $\xi' = \langle (\xi'_j)_{0 \leq j \leq n'}, (t'_j)_{0 \leq j < n'} \rangle$ such

that $\forall j, active(j) \leq active(j+1)$ (with $active(j)$ the only position in $\rho_j^2(\underline{t}_j)$).

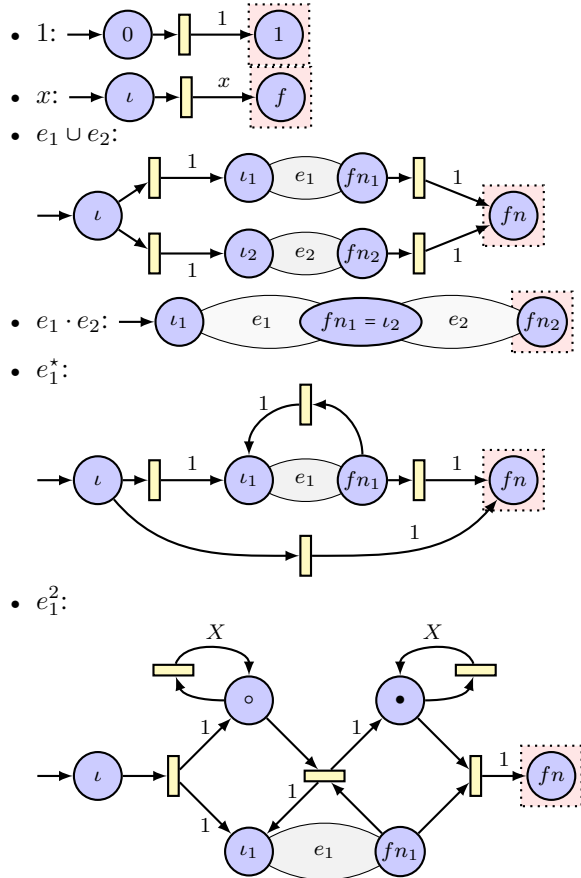
Now, suppose we have such a reading; we can build an embedding $(\eta_k)_{0 \leq k \leq n}$ as follows. For $k < n$, define $T_k := \{(\underline{t}'_j, \overline{t}'_j) \mid active(j) = k\}$. We describe the construction incrementally:

- $\eta_0 = [\iota_2 \mapsto \iota_1]$.
- For all $p \in \text{dom}(\eta_k) \setminus \cup_{(\underline{t}, \overline{t}) \in T_k} \underline{t}$ we simply set $\eta_k(p) = \eta_{k-1}(p)$.
- Otherwise, $\forall (\underline{t}'_j, \overline{t}'_j) \in T_k$, let $q \in \underline{t}'_j$. Then, for all (x, p) in \underline{t}'_j , because ρ^2 is a reading and by construction of $\mathcal{G}(\xi)$ we also know that there is some $p' \in \xi_{k+1}$ that satisfies $(x, p') \in \overline{t}_k$ and $\rho_{k+1}^1(p') = \rho_{j+1}^2(p)$. That p' is a good choice for p , hence we define $\eta_k(p) = p'$.

It is then administrative to check that $(\eta_k)_{0 \leq k \leq n}$ is indeed an embedding.

F. Omitted proof: Lemma 29

We define the automaton of a regular expression with squaring with a single final state. We describe this construction graphically.



This construction can be performed in linear time and space. By further analysing this construction, one can prove that

- 1) $\lambda(|e|) \in \mathcal{L}(\mathcal{A}(e))$;
- 2) for every $G \in \mathcal{G}(\mathcal{A}(e))$, there is a word $w \in (|e|)$ such that $G \blacktriangleleft \lambda(w)$.

By combining Items 1 and 2 we establish that

$$\mathcal{L}(\mathcal{A}(e)) = \blacktriangleleft \lambda(|e|).$$

Thus by Equation (8), we get that $(|e|) = (|f|)$ is equivalent to $\mathcal{L}(\mathcal{A}(e)) = \mathcal{L}(\mathcal{A}(f))$. This gives us the EXPSPACE-hardness of the equivalence of Petri automata.

G. Comparing simple automata is EXPSPACE-hard

We perform a reduction from the problem $\text{RSQ}(X)$: deciding whether a regular expressions with squaring e ($e^2 := e \cdot e$) denotes the universal language X^* . This problem was shown in [14] to be EXPSPACE-complete. The proof will follow the method used in [10] to obtain a complexity lower bound on the trace equivalence of finite nets without hidden transitions. To avoid confusion, the regular language denoted by the expression e will be written $(|e|)$. Any word u can be seen uniquely as a linear graph $\lambda(u)$. By extension, the set of graphs of words from $(|e|)$ will be denoted by $\lambda(|e|)$.

First, notice that if u and v are just words over X , $\lambda(u) \blacktriangleleft \lambda(v)$ is equivalent to $u = v$. Because of this, it is straightforward to check that for any $e, f \in \text{Reg}_X^2$ the following holds

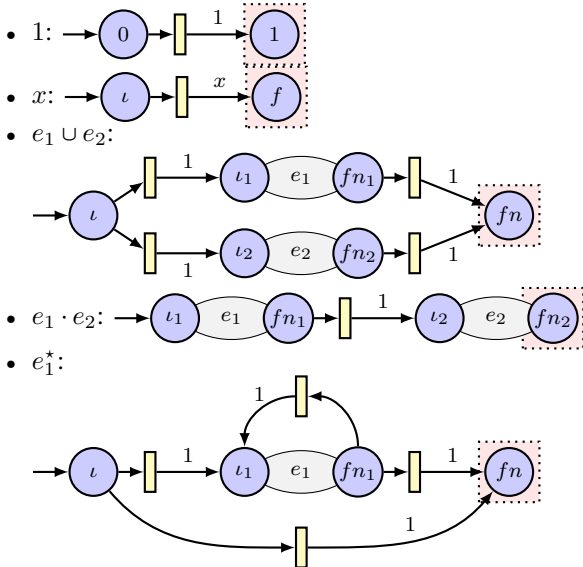
$$\blacktriangleleft (\lambda(|e|)) = \blacktriangleleft (\lambda(|f|)) \Leftrightarrow (|e|) = (|f|). \quad (9)$$

We first reduce to the containment of Petri automata with 1.


Lemma 34. *The problem of deciding whether the language of a regular expression with squaring is X^* is polynomial-time reducible to the containment of languages recognised by Petri automata.*

Proof. Given an expression e on this signature, we can build in linear time a Petri automaton \mathcal{A} , with a linear number of places and transitions. The closure of the language denoted by e is exactly the language recognised by \mathcal{A} .

The automata we produce here only have one final configuration, consisting in a singleton. The construction is a straightforward adaptation of Thompson's algorithm for NFA [20]. We describe this construction graphically.



• e_1^2 :



The automaton for e has at most $4 \times |e|$ places, and at most $(2 \times |X| + 4) \times |e|$ transitions. It is quite obvious that this construction can be performed in linear time and space. By further analysing this construction, one can prove that

- 1) $\lambda(|e|) \subseteq \mathcal{L}(\mathcal{A}(e))$;
- 2) for every $G \in \mathcal{G}(\mathcal{A}(e))$, there is a word $w \in (|e|)$ such that $G \blacktriangleleft \lambda(w)$.

By combining Items 1 and 2 we establish that

$$\mathcal{L}(\mathcal{A}(e)) = \blacktriangleleft \lambda(|e|).$$

Thus by Equation (9), we get that $(|e|) = (|f|)$ is equivalent to $\mathcal{L}(\mathcal{A}(e)) = \mathcal{L}(\mathcal{A}(f))$. Hence $\mathcal{L}(\mathcal{A}(X^*)) \subseteq \mathcal{L}(\mathcal{A}(e))$ is equivalent to $(|e|) = X^*$. \square

The previous automata are not simple, as they use 1 as a label. We can now get rid of this label, to obtain EXPSPACE-hardness for simple Petri automata.

Lemma 35. *The problem of deciding whether the language of a regular expression with squaring is X^* is polynomial-time reducible to the containment of languages recognised by simple Petri automata.*

Proof. We use a similar trick as in [10]: we reuse the previous automaton $\mathcal{A}(e)$ by considering it as a simple automaton over the alphabet $X \cup \{1\}$ (thus forgetting about the special semantics of 1, and seeing it simply as a standard letter). Notice that $w \in (|e|)$ is equivalent to the existence of a word $u \in (X \cup 1)^*$ such that $\lambda(u) \in \mathcal{L}(\mathcal{A}(e))$ and w can be obtained by erasing from u the occurrences of the letter 1. By carefully analysing the automaton $\mathcal{A}(e)$, we may actually impose that u does not contain more than $n_e := (2 \times |X| + 4) \times |e|$ consecutive 1. By adding to each place of $\mathcal{A}(e)$ a transition looping on that place and labelled by 1, one obtains an automaton $\mathcal{A}'(e)$ such that $w \in (|e|)$ iff the word $w' \in (1^{n_e} X)^* 1^{n_e}$ obtained from w by inserting 1s as necessary is accepted by $\mathcal{A}'(e)$. As a consequence $(|e|) = X^*$ iff $\mathcal{L}(\mathcal{B}) \subseteq \mathcal{L}(\mathcal{A}'(e))$, where \mathcal{B} is a simple Petri automaton recognising the regular language $(1^{n_e} X)^* 1^{n_e}$ (there are automata of linear size for this language). \square

H. Traces as pomsets

Let ξ be a run in a simple Petri automaton, and let ξ' be the corresponding run in the corresponding labelled Petri net. Let $\mathcal{G}(\xi) = \langle V, E, \iota, o \rangle$. It is not difficult to check that the pomset-trace of ξ' is isomorphic to $\langle E, <_E \rangle$, where $<_E$ is the

transitive closure of the relation $<$ defined on E by $\forall x, y, z \in V, a, b \in X, (x, a, y) < (y, b, z)$.

The correspondence is even stronger: two graphs produced by accepting runs in (possibly different) automata $\mathcal{G}(\xi_1) = \langle V_1, E_1, \iota_1, o_1 \rangle$ and $\mathcal{G}(\xi_2) = \langle V_2, E_2, \iota_2, o_2 \rangle$ are isomorphic if and only if their pomset-traces $(E_1, <_{E_1})$ and $(E_2, <_{E_2})$ are isomorphic. The proof of this relies on the fact that accepting runs produce graphs satisfying the following properties:

$$\begin{aligned} \forall (x, a, y) \in E, x \neq \iota &\Leftrightarrow (\exists z, b \in V \times X : (z, b, x) \in E) \quad , \\ \forall (x, a, y) \in E, y \neq o &\Leftrightarrow (\exists z, b \in V \times X : (y, b, z) \in E) \quad . \end{aligned}$$

Hence, pomset-trace language equivalence corresponds exactly to equivalence of the sets of produced graphs in our setting (up to isomorphism).