

EQUIVALENCE CHECKING FOR WEAK BI-KLEENE ALGEBRA *

TOBIAS KAPPÉ^a, PAUL BRUNET^b, BAS LUTTIK^c, ALEXANDRA SILVA^b, AND FABIO ZANASI^b

^a Cornell University, Ithaca, New York, USA
e-mail address: tkappe@cornell.edu

^b University College London, United Kingdom

^c Eindhoven University of Technology, The Netherlands

ABSTRACT. Pomset automata are an operational model of *weak bi-Kleene algebra*, which describes programs that can *fork* an execution into parallel threads, upon completion of which execution can *join* to resume as a single thread. We characterize a fragment of pomset automata that admits a decision procedure for language equivalence. Furthermore, we prove that this fragment corresponds precisely to *series-rational expressions*, i.e., rational expressions with an additional operator for bounded parallelism. As a consequence, we obtain a new proof that equivalence of series-rational expressions is decidable.

1. INTRODUCTION

Kleene’s theorem states the correspondence between the operational world of automata and the denotational world of expressions, on formal languages [Kle56]. This famous discovery has proven pivotal to establish later results, such as Kozen’s axiomatisation of equivalence of rational expressions [Koz94], as well as to transpose the application of algorithms from a denotational to an operational setting — for instance, one can leverage Hopcroft and Karp’s algorithm for finite automata [HK71] to decide equivalence of rational expressions.

In spite of their simplicity, finite automata and rational expressions provide valuable tools in analyzing the behaviour of sequential programs [Koz96]. The behavioural patterns of present-day programs, however, are not limited to sequential scenarios, where each event either strictly precedes or succeeds all others. Indeed, reasoning about programs that run on multi-core processors requires us to adapt our descriptions such that two events need not be strictly ordered, but instead may occur *in parallel*. The study of *concurrent Kleene algebra* [HMSW09], in the broadest sense, is concerned with extending techniques from rational expressions and finite automata to reason about systems that include parallelism.

Key words and phrases: sr-expressions, pomset automata, Kleene theorem, language equivalence.

* An earlier version of this paper was published at CONCUR’17 [KBL⁺17].

T. Kappé was partially supported by the ERC Starting Grant 679127 (ProFoundNet) and DARPA grant HR001120C0107 (Pronto). A. Silva was partially supported by the ERC Starting Grant 679127 (ProFoundNet) and a Leverhulme Prize (PLP-2016-129). P. Brunet acknowledges support from EPSRC grant n. EP/R006865/1. F. Zanasi acknowledges support from EPSRC grant n. EP/R020604/1.

We propose *pomset automata* as an operational model for a fragment of concurrent Kleene algebra known as *weak bi-Kleene algebra*, which describes programs where an execution may *fork* into parallel computations, to *join* the results of those computations later on, resuming execution. The language semantics of these automata is given by sets of *partially ordered multisets*, or *pomsets*. The first main contribution is a proof that language equivalence of states is decidable for the class of *fork-acyclic* finite pomset automata. The second main contribution is a Kleene theorem, which shows that this same fragment corresponds precisely to the denotational model of bi-Kleene algebra, known as *series-rational expressions*, or *sr-expressions* for short [LW00, LS14] — that is, rational expressions extended with parallel composition. This correspondence then yields a decision procedure for deciding equivalence of series-rational expressions, via pomset automata.

In Section 2, we discuss related work; we go over the necessary background regarding pomsets in Section 3. In Section 4, we introduce pomset automata and their semantics. Additionally, we introduce a (structural) restriction on pomset automata, defining the class of *fork-acyclic* pomset automata. In Section 5, we develop an algorithm for checking language equivalence of a subclass of fork-acyclic pomset automata; in Section 6, we extend this procedure to fork-acyclic pomset automata in general. In Section 7, we show how to obtain a pomset automaton that recognizes the language of a series-rational expression; conversely, in Section 8, we show how to obtain an equivalent series-rational expression from a finite and fork-acyclic pomset automaton. We list directions for further work in Section 9.

For the sake of self-containment, we include proofs of formal claims that are not cited. Routine proofs are delegated to the appendices to ensure brevity.

2. RELATED WORK

There exist three pomset-based operational models for sr-expressions in the literature. *Branching automata* were pioneered by Lodaya and Weil [LW00]. These are non-deterministic finite automata enriched with two additional types of transition to mediate forking and joining of computation. No decision procedure for language equivalence of branching automata is known. Branching automata also come equipped with a reverse construction, which shows how to obtain an equivalent sr-expression from a particular class of branching automata. The difference is that the description of this class involves a *semantic* condition, i.e., makes a statement about the pomsets that can be accepted by a state in the branching automaton, whereas fork-acyclic pomset automata are described in purely structural terms.

Another model, proposed by Jipsen and Moshier [JM16] based on [LW00] and also called *branching automata*, is given by non-deterministic finite automata, enriched with a relation that specifies where computation may be joined, provided all threads can be traced back to a given common state. In some sense, pomset automata are a dual to this model, in that they specify where execution can be forked, after which all threads can be joined at a given state after termination. As far as we can tell, there is no known decision procedure for language equivalence of branching automata. Branching automata in the style of Jipsen and Moshier also come with a translation back to sr-expressions; being based on [LW00], this construction inherits the semantic description of automata to which it can be applied.

Petri nets, specifically safe Petri nets, can also be used to describe the behaviours modelled by sr-expressions [BPS17, LRR03]. The advantage of this approach is that it allows one to use results from Petri net theory to study sr-expressions. Furthermore, particularly in the case of [BPS17], one can leverage the encoding of sr-expressions into Petri nets to

develop a decision procedure for equivalence of sr-expressions, as well as a more general type of equivalence that allows threads to sequentialise (corresponding to concurrent Kleene algebra proper). However, it should be noted that safe Petri nets can express a form of concurrency that is strictly more general than the type of concurrency that can be described by concurrent Kleene algebra [Gra81]. As such, converting a safe Petri net to an equivalent sr-expression necessarily discards some behaviour [LRR03].

The operational models discussed above associate an automaton or Petri net to an sr-expression by induction on the structure of the expression, using a Thompson-style translation [Tho68]. In contrast, our expressions-to-automata translation generalizes Antimirov’s construction [Ant96], and thus allows the operational representation to be constructed lazily. This is particularly beneficial for algorithms that explore the state space of automata step-by-step, as it prevents them from computing the entire state space.

In [BLMvT16], Baeten et al. give an operational semantics to sr-expressions in terms of (non-deterministic) transitions systems, by interpreting the parallel composition as interleaving, which obviates the need for a parallel thread construction. They show that there exist transition systems that are not bisimilar to any sr-expression, and characterise the fragment of systems for which such an sr-expression does exist. A full Kleene theorem is recovered when they extend sr-expressions with interaction.

Our algorithm to check language equivalence in pomset automata was inspired by the work of Laurence and Struth [LS14] on a more general form of sr-expressions; specifically, the idea of checking language equality by computing the *atoms* of languages is due to them.

This paper is an extension of a CONCUR’17 paper [KBL⁺17]; in comparison, the present work contains a more generalised presentation of pomset automata that allows for non-determinism, which enables us to define the class of *well-structured* pomset automata. We then use this structural restriction to guide the construction of a decision procedure for language equivalence of fork-acyclic pomset automata. Moreover, the syntactic derivatives are now presented in the style of Antimirov [Ant96], rather than Brzozowski [Brz64].

Another closely related paper is [KBL⁺19], in which we generalised the Kleene theorem of [KBL⁺17] to *series-parallel regular expressions*, which include the parallel variant of the Kleene star known as *parallel star*. For this generalisation to work, one needs to loosen the notion of fork-acyclicity to the strictly more liberal *well-nestedness*; crucially, the definition of well-nested pomset automata relies on the presentation of pomset automata found in [KBL⁺17]. Furthermore, in [KBL⁺19], we show that language equivalence of pomset automata in general is undecidable — this justifies the fact that the decision procedure presented in this paper requires fork-acyclicity.

3. PRELIMINARIES

We fix a finite set of symbols Σ , referred to as the *alphabet*. When S and T are sets, we write T^S for the set of functions from S to T . We write 2^S for the set of subsets of S , which can be identified with the functions from S to the two-element set $2 = \{0, 1\}$.

A *multiset* over a set S is a “subset” of S where elements may occur more than once; more formally, it is a function $\phi : S \rightarrow \mathbb{N}$. Finite multisets are denoted using double braces, e.g., $\phi = \{1, 1\}$ is the multiset over \mathbb{N} where $\phi(n) = 2$ if $n = 1$, and $\phi(n) = 0$ otherwise. In the following, we fix multisets ϕ, ψ over S . When $s \in S$, we use $s \in \phi$ as a shorthand for $\phi(s) \neq 0$. We say that ϕ is *finite* if there are finitely many $s \in S$ such that $s \in \phi$. If ϕ is finite, then the *size* of ϕ , denoted $|\phi|$, is given by $\sum_{s \in S} \phi(s)$. We write $\mathbb{M}(S)$ for the *set of*

finite multisets over S . We denote the empty multiset by \square , and write $\phi \sqcup \psi$ for the *disjoint union* of ϕ and ψ , where $(\phi \sqcup \psi)(s) = \phi(s) + \psi(s)$. For instance, $\{0\} \sqcup \{0, 1\} = \{0, 0, 1\}$.

A *non-deterministic finite automaton (NFA)* is a tuple $A = \langle Q, \delta, F \rangle$ where Q is a finite set of *states*, with $F \subseteq Q$ the *accepting states*, and $\delta : Q \times \Sigma \rightarrow 2^Q$ is a function. The *language* of $q \in Q$ in A , denoted $L_A(q)$, is the set of words $w = \mathbf{a}_1 \cdots \mathbf{a}_n$ such that there exist $q = q_0, \dots, q_n \in Q$ where $q_{i+1} \in \delta(q_i, \mathbf{a}_{i+1})$ for $0 \leq i < n$, and $q_n \in F$.

3.1. Pomsets. We commonly represent an execution of a program as a *word* over some finite alphabet Σ . In such a word, each *position* corresponds to an *event* in the execution whose *name* is given by the symbol on that position; events are ordered according to their positions. For instance, if $\Sigma = \{\mathbf{a}, \mathbf{b}, \mathbf{c}\}$, then the word \mathbf{abca} represents an execution where an event of type \mathbf{a} occurs, followed by events of type \mathbf{b} , \mathbf{c} and \mathbf{a} , in that order.

To represent an execution of a program with concurrency, we need to relax this model to allow a partial order on events. For instance, a concurrent program may execute an event of type \mathbf{a} before forking into two threads that execute events of type \mathbf{b} and \mathbf{c} respectively, after which the threads join to perform an event of type \mathbf{a} . Note now, in this execution, there is no ordering of the events labelled by \mathbf{b} and \mathbf{c} — they are *concurrent*.

The model most commonly found in the literature to account for such executions was proposed independently by Winkowski [Win77] and Pratt [Pra82], and studied extensively by Winkowski [Win79] and Grabowski [Gra81]; we adopt Pratt's terminology. Defining the model requires some patience, as the indirection between *events* and their *names*, which is given implicitly in words by *positions* and their *symbols* is slightly tricky to generalize.

Definition 3.1. A *labelled partially ordered set (labelled poset)* is a tuple $\mathbf{u} = \langle S_{\mathbf{u}}, \leq_{\mathbf{u}}, \lambda_{\mathbf{u}} \rangle$ consisting of a *carrier set* $S_{\mathbf{u}}$, a partial order $\leq_{\mathbf{u}}$ on $S_{\mathbf{u}}$ and a *labelling function* $\lambda_{\mathbf{u}} : S_{\mathbf{u}} \rightarrow \Sigma$.

For technical reasons, we adopt the convention that the carrier of a labelled poset is a subset of \mathbb{N} ; under this convention, the collection of labelled posets is a proper set.

The definition above gets us close to where we need to be. For instance, the example execution above can be represented by the labelled poset $\langle S_{\mathbf{u}}, \leq_{\mathbf{u}}, \lambda_{\mathbf{u}} \rangle$, in which

$$S_{\mathbf{u}} = \{1, 2, 3, 4\} \quad 1 \leq_{\mathbf{u}} 2 \leq_{\mathbf{u}} 4 \quad 1 \leq_{\mathbf{u}} 3 \leq_{\mathbf{u}} 4 \quad \lambda_{\mathbf{u}} = \{1 \mapsto \mathbf{a}, 2 \mapsto \mathbf{b}, 3 \mapsto \mathbf{c}, 4 \mapsto \mathbf{a}\}$$

However, if the event labelled by \mathbf{b} were represented by 5 instead of 2 (adjusting $\leq_{\mathbf{u}}$ and $\lambda_{\mathbf{u}}$ accordingly), then this new labelled poset would still represent the same execution; in some sense, we care only about the labels of the events, and their order. Hence, we should abstract from the exact contents of the carrier. This is done as follows.

Definition 3.2 (Pomsets). Let $\mathbf{u} = \langle S_{\mathbf{u}}, \leq_{\mathbf{u}}, \lambda_{\mathbf{u}} \rangle$ and $\mathbf{v} = \langle S_{\mathbf{v}}, \leq_{\mathbf{v}}, \lambda_{\mathbf{v}} \rangle$ be labelled posets. A *labelled poset isomorphism* from \mathbf{u} to \mathbf{v} is a bijection $h : S_{\mathbf{u}} \rightarrow S_{\mathbf{v}}$, such that $\lambda_{\mathbf{v}} \circ h = \lambda_{\mathbf{u}}$, and for $s, s' \in S_{\mathbf{u}}$ we have $s \leq_{\mathbf{u}} s'$ if and only if $h(s) \leq_{\mathbf{v}} h(s')$. We write $\mathbf{u} \cong \mathbf{v}$ if such an isomorphism exists between \mathbf{u} and \mathbf{v} , and note that \cong is an equivalence.

A *partially ordered multiset*, or *pomset* for short, is an equivalence class of labelled posets; we write $[S_{\mathbf{u}}, \leq_{\mathbf{u}}, \lambda_{\mathbf{u}}]$ for the equivalence class of the labelled poset $\langle S_{\mathbf{u}}, \leq_{\mathbf{u}}, \lambda_{\mathbf{u}} \rangle$.

Since the collection of labelled posets forms a proper set, \cong is a proper relation. Consequently, a pomset is also a proper set, as is the set of pomsets.

We write 1 for the *empty pomset*, i.e., the isomorphism class consisting of the unique empty labelled poset. When $\mathbf{a} \in \Sigma$ we may write \mathbf{a} to denote the unique pomset containing one event, which is labelled with \mathbf{a} ; such a pomset is called *primitive*.

Definition 3.3 (Pomset composition). Let $U = [S_{\mathbf{u}}, \leq_{\mathbf{u}}, \lambda_{\mathbf{u}}]$ and $V = [S_{\mathbf{v}}, \leq_{\mathbf{v}}, \lambda_{\mathbf{v}}]$ be pomsets. Without loss of generality, we may assume that $S_{\mathbf{u}}$ and $S_{\mathbf{v}}$ are disjoint.

We define two types of composition. The *sequential composition* of U and V , denoted $U \cdot V$, is the pomset $[S_{\mathbf{u}} \cup S_{\mathbf{v}}, \leq_{\mathbf{u}} \cup \leq_{\mathbf{v}} \cup (S_{\mathbf{u}} \times S_{\mathbf{v}}), \lambda_{\mathbf{u}} \cup \lambda_{\mathbf{v}}]$. The *parallel composition* of U and V , denoted $U \parallel V$, is the pomset $[S_{\mathbf{u}} \cup S_{\mathbf{v}}, \leq_{\mathbf{u}} \cup \leq_{\mathbf{v}}, \lambda_{\mathbf{u}} \cup \lambda_{\mathbf{v}}]$.

In the above, $\lambda_{\mathbf{u}} \cup \lambda_{\mathbf{v}} : S_{\mathbf{u}} \cup S_{\mathbf{v}} \rightarrow \Sigma$ is the function that agrees with $\lambda_{\mathbf{u}}$ on $S_{\mathbf{u}}$, and with $\lambda_{\mathbf{v}}$ on $S_{\mathbf{v}}$ — note that $\lambda_{\mathbf{u}} \cup \lambda_{\mathbf{v}}$ is well-defined, given that $S_{\mathbf{u}}$ and $S_{\mathbf{v}}$ are disjoint. The order relations of $U \cdot V$ and $U \parallel V$ are partial orders for the same reason.

Of course, one should check that the operators given above are well-defined, i.e., that $U \cdot V$ and $U \parallel V$ are the same (as in, given by isomorphic labelled posets) regardless of the disjoint labelled posets that are chosen as representatives. This turns out to be the case.

It is easy to see that both operators are associative, and that \parallel is commutative. Furthermore, 1 is the unit of both sequential and parallel composition, i.e., $U \cdot 1 = 1 \cdot U = U \parallel 1 = U$ for all pomsets U [Gis88]. For the remainder of this paper, we adopt the convention that \cdot binds more tightly than \parallel , i.e., $U \cdot V \parallel W$ should be read as $(U \cdot V) \parallel W$.

Definition 3.4 (Pomset types). Let U be a pomset. We say that U is *sequential* if there exist non-empty pomsets U_1, U_2 such that $U = U_1 \cdot U_2$. Also, U is a *sequential prime* if it is non-empty, and for all pomsets V and W such that $U = V \cdot W$, we have $V = 1$ or $W = 1$.

Similarly, we say that U is *parallel* if there exist non-empty pomsets U_1, U_2 such that $U = U_1 \parallel U_2$. Also, U is a *parallel prime* if it is non-empty, and for all pomsets V and W such that $U = V \parallel W$, it holds that $V = 1$ or $W = 1$.

For our type of concurrency, we study a specific type of pomset, as follows.

Definition 3.5. The set of *series-parallel pomsets*, or *sp-pomsets*, denoted $\text{SP}(\Sigma)$, is the smallest set that contains the empty and primitive pomsets, and is closed under sequential and parallel composition. In other words, $\text{SP}(\Sigma)$ is the smallest set satisfying the rules

$$\frac{}{1 \in \text{SP}(\Sigma)} \quad \frac{\mathbf{a} \in \Sigma}{\mathbf{a} \in \text{SP}(\Sigma)} \quad \frac{U, V \in \text{SP}(\Sigma)}{U \cdot V \in \text{SP}(\Sigma)} \quad \frac{U, V \in \text{SP}(\Sigma)}{U \parallel V \in \text{SP}(\Sigma)}$$

It should be clear that the pomsets in $\text{SP}(\Sigma)$ built without parallel composition have a total order on their nodes, and hence correspond exactly to words. We will make this identification throughout this paper, writing Σ^* for the set of words over Σ .

Series-parallel pomsets have the convenient property that they may be partitioned into empty, primitive, sequential and parallel pomsets, in the following way.

Lemma 3.6 [Gis88, Theorem 3.1]. *Let $U \in \text{SP}(\Sigma)$. Exactly one of the following is true: (i) U is empty, or (ii) U is primitive, or (iii) U is sequential, or (iv) U is parallel.*

Another useful tool in dissecting pomsets comes from *factorisation*. In a sense, factorising a pomset is analogous to writing a word as the sequence of its letters; the difference here is that whereas words can be composed only sequentially, pomsets can also be composed in parallel — hence, we obtain two types of factorisation.

Definition 3.7 (Factorisation). Let U be a pomset. When $U = U_1 \cdots U_n$ with U_1, \dots, U_n sequential primes, we refer to the sequence U_1, \dots, U_n as a *sequential factorisation* of U .

Similarly, when $U = U_1 \parallel \cdots \parallel U_n$ such that U_1, \dots, U_n are parallel primes, we refer to the multiset $\{U_1, \dots, U_n\}$ as a *parallel factorisation* of U .

Just like a word can be uniquely written as a concatenation of its symbols, so does sequential factorisation give rise to such a unique concatenation.

Lemma 3.8 ([Gis88, Lemma 3.2] and [Gra81, Proposition 2]; see also [GKM72]). *Sequential factorisations exist uniquely for sp-pomsets.*

We conclude with a similar statement about parallel factorisation.

Lemma 3.9. *Parallel factorisations exist uniquely for sp-pomsets.*

3.2. Pomset languages. All possible executions of a sequential program (described as words) can be collected in a set to form a *language* describing the behaviour of a program. Analogously, we can collect pomsets in a *pomset language* to describe the behaviour of a concurrent program, as follows.

Definition 3.10. A *pomset language* is a set of pomsets; a pomset language made up of sp-pomsets is referred to as a *series-parallel language*, or *sp-language* for short.

The composition operators of pomsets lift in a pointwise manner; concretely

$$L \cdot L' = \{U \cdot U' : U \in L, U' \in L'\} \quad L \parallel L' = \{U \parallel U' : U \in L, U' \in L'\}$$

The Kleene closure also applies to pomset languages, as follows

$$L^* = \bigcup_{n \in \mathbb{N}} L^n \quad \text{in which} \quad L^0 = \{1\} \quad \text{and} \quad L^{n+1} = L \cdot L^n$$

To relate pomset languages over different alphabets, the notion of *substitution* is useful; a substitution allows us to translate a pomset language over one alphabet into a pomset language over another alphabet, just by substituting the letters.

Definition 3.11. Let Δ be an alphabet. A *substitution* is a function $\zeta : \Sigma \rightarrow 2^{\text{SP}(\Delta)}$. We can lift the domain ζ to $\text{SP}(\Sigma)$ inductively, as follows:

$$\zeta(1) = \{1\} \quad \zeta(U \cdot V) = \zeta(U) \cdot \zeta(V) \quad \zeta(U \parallel V) = \zeta(U) \parallel \zeta(V)$$

When L is a pomset language, we write $\zeta(L)$ for the set $\bigcup_{U \in L} \zeta(U)$.

Finally, we call ζ *atomic* if all of the following hold:

- (i) for $\mathbf{a} \in \Sigma$ we have that $\zeta(\mathbf{a})$ consists of sequential primes exclusively, and
- (ii) for $\mathbf{a}, \mathbf{b} \in \Sigma$, we have that $\zeta(\mathbf{a}) \cap \zeta(\mathbf{b}) \neq \emptyset$ if and only if $\mathbf{a} = \mathbf{b}$.

Atomic substitutions have the following useful properties.

Lemma 3.12. *Let $L, L' \subseteq \Sigma^*$, and let ζ be a substitution. If ζ is atomic, then*

$$\zeta(L \cap L') = \zeta(L) \cap \zeta(L') \quad \zeta(L \setminus L') = \zeta(L) \setminus \zeta(L') \quad \zeta(L) = \emptyset \iff L = \emptyset$$

3.3. Series-rational expressions. Rational expressions can denote languages, which in turn can describe the behaviour of a program. Continuing our analogy, *series-rational expressions* [LW00] can denote the behaviour of a program with fork/join-style concurrency. Essentially, these are rational expressions extended with parallel composition.

Definition 3.13. The set of *series-rational expressions*, or *sr-expressions* for short, denoted \mathcal{T} , is the smallest set generated by the grammar

$$e, f ::= 0 \mid 1 \mid \mathbf{a} \in \Sigma \mid e + f \mid e \cdot f \mid e \parallel f \mid e^*$$

Like rational expressions, series-rational expressions have a straightforward semantics in terms of sp-languages, where each operator is interpreted as an operator on sp-languages.

Definition 3.14 (Semantics). We define $\llbracket - \rrbracket : \mathcal{T} \rightarrow 2^{\text{SP}(\Sigma)}$ inductively, as follows.

$$\begin{aligned} \llbracket 0 \rrbracket &= \emptyset & \llbracket 1 \rrbracket &= \{1\} & \llbracket \mathbf{a} \rrbracket &= \{\mathbf{a}\} & \llbracket e + f \rrbracket &= \llbracket e \rrbracket \cup \llbracket f \rrbracket \\ \llbracket e \cdot f \rrbracket &= \llbracket e \rrbracket \cdot \llbracket f \rrbracket & \llbracket e \parallel f \rrbracket &= \llbracket e \rrbracket \parallel \llbracket f \rrbracket & \llbracket e^* \rrbracket &= \llbracket e \rrbracket^* \end{aligned}$$

When L is a pomset language and there exists an $e \in \mathcal{T}$ such that $L = \llbracket e \rrbracket$, we say that L is a *series-rational language*, or *sr-language* for short.

An important property of sr-expressions is whether their semantics contains the empty pomset. We syntactically characterise such sr-expressions.

Definition 3.15. \mathcal{F} is the smallest subset of \mathcal{T} that satisfies the following for all $e, f \in \mathcal{T}$:

$$\frac{}{1 \in \mathcal{F}} \quad \frac{e \in \mathcal{F}}{e + f \in \mathcal{F}} \quad \frac{f \in \mathcal{F}}{e + f \in \mathcal{F}} \quad \frac{e, f \in \mathcal{F}}{e \cdot f \in \mathcal{F}} \quad \frac{e, f \in \mathcal{F}}{e \parallel f \in \mathcal{F}} \quad \frac{}{e^* \in \mathcal{F}}$$

To see that \mathcal{F} indeed characterizes the empty pomset property, we have

Lemma 3.16. *Let $e \in \mathcal{T}$. Now $e \in \mathcal{F}$ if and only if $1 \in \llbracket e \rrbracket$.*

4. POMSET AUTOMATA

We now turn our attention to pomset automata, which are intended as an operational model for bi-Kleene algebra. Intuitively, a pomset automaton is a *non-deterministic* automaton enriched with an additional type of transition. Instances of this type of transition tell us where execution may fork into several “threads”, as well as where execution resumes when each of these threads has reached an accepting state. More formally, we have the following.

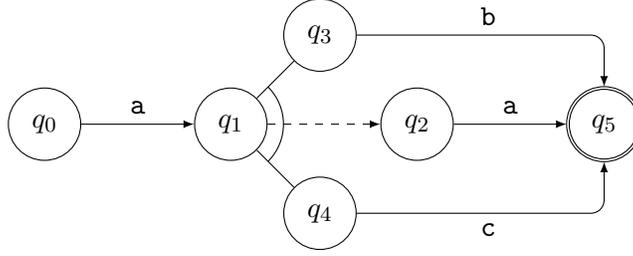
Definition 4.1. A *pomset automaton* (PA) is a tuple $\langle Q, F, \delta, \gamma \rangle$ where

- Q is a set of *states*, with $F \subseteq Q$ the *accepting states*, and
- $\delta : Q \times \Sigma \rightarrow 2^Q$ is the *sequential transition function*, and
- $\gamma : Q \times \mathbb{M}(Q) \rightarrow 2^Q$ is the *parallel transition function*.

Lastly, for all $q \in Q$, there are only finitely many $\phi \in \mathbb{M}(Q)$ such that $\gamma(q, \phi) \neq \emptyset$.

In the above, $q' \in \gamma(q, \{r_1, \dots, r_n\})$ should be interpreted to mean that, in state q , the automaton may fork into states r_1, \dots, r_n , and when all of these have reached an accepting state, may resume computation from q' . The final requirement ensures that only finitely many fork transitions of a state can be meaningful, i.e., lead to an accepting state.

Visually, we can represent pomset automata in a style similar to finite automata, as in Figure 1. There, a state is represented by a vertex, which is doubly circled when the state is

FIGURE 1. PA accepting $\mathbf{a} \cdot (\mathbf{b} \parallel \mathbf{c}) \cdot \mathbf{a}$.

accepting. Transitions are represented by edges: the edge between q_0 and q_1 encodes that $q_1 \in \delta(q_0, \mathbf{a})$, and the multi-ended edge that connects q_1 to q_2 and q_3 as well as q_4 signifies that $q_2 \in \gamma(q_1, \{\!\{q_3, q_4\}\!\})$. To avoid visual clutter, we draw neither the sequential transitions where $\delta(q, \mathbf{a}) = \emptyset$, nor the parallel transitions where $\gamma(q, \{\!\{r_1, \dots, r_n\}\!\}) = \emptyset$. Our convention that for $q \in Q$ there are only finitely many $\phi \in \mathbb{M}(Q)$ with $\gamma(q, \phi) \neq \emptyset$ ensures that a PA with finitely many states also has only finitely many transitions, and can therefore be drawn.

For the remainder of this section, we fix a PA $A = \langle Q, F, \delta, \gamma \rangle$. We proceed to define how one can “read” an sp-pomset U while transitioning from a state q to a state q' , in a way that matches the intuition given to the parallel transition function. This information is encoded in a ternary relation between states, pomsets, and states, called the *run relation*. If a state q is related to a pomset U and another state q' by this relation, it means that, starting in state q , we can read the pomset U to end up in state q' . The pomsets that we can read to reach an accepting state form the language of a state.

Definition 4.2 (Runs). We define $\rightarrow_A \subseteq Q \times \text{SP}(\Sigma) \times Q$ as the smallest relation satisfying

$$\frac{}{q \xrightarrow{1}_A q} \quad \frac{q' \in \delta(q, \mathbf{a})}{q \xrightarrow{\mathbf{a}}_A q'} \quad \frac{q \xrightarrow{U}_A q'' \quad q'' \xrightarrow{V}_A q'}{q \xrightarrow{U \cdot V}_A q'} \quad \frac{\forall 1 \leq i \leq n. q_i \xrightarrow{U_i}_A q'_i \in F \quad q' \in \gamma(q, \{\!\{q_1, \dots, q_n\}\!\})}{q \xrightarrow{U_1 \parallel \dots \parallel U_n}_A q'}$$

The *language* of $q \in Q$, denoted $L_A(q)$, is the set $\{U \in \text{SP}(\Sigma) : q \xrightarrow{U}_A q' \in F\}$.

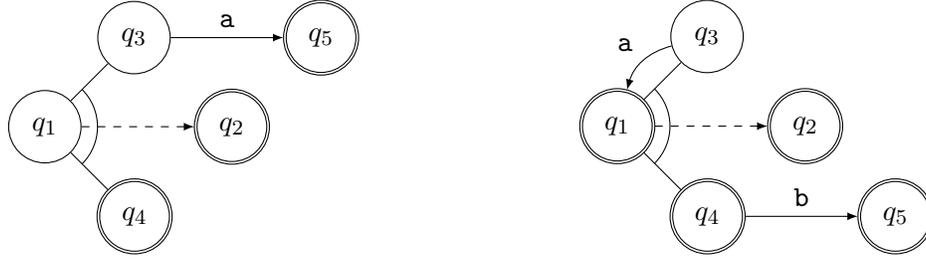
Given an element $q \xrightarrow{U}_A q'$ of the run relation, a proof tree (using the inference rules above) witnessing that this triple occurs in \rightarrow_A contains all structural information about how the pomset automaton reads the pomset, i.e., the order in which the states were visited, which fork transitions were used, et cetera. Such a proof tree can therefore be called a *run* of the pomset automaton; we shall often abuse this nomenclature and refer to individual elements of the run relation as runs, with their underlying proof tree implicitly present.

Example 4.3. Suppose A is the PA in Figure 1. Then, we have that

$$q_0 \xrightarrow{\mathbf{a}}_A q_1 \quad q_2 \xrightarrow{\mathbf{a}}_A q_5 \quad q_3 \xrightarrow{\mathbf{b}}_A q_5 \quad q_4 \xrightarrow{\mathbf{c}}_A q_5$$

From the latter two runs and the fact that $q_2 \in \gamma(q_1, \{\!\{q_3, q_4\}\!\})$, it follows that $q_1 \xrightarrow{\mathbf{b} \parallel \mathbf{c}}_A q_2$ by the last rule above. By applying the third rule of run composition to this run and the first two runs above, we find that $q_0 \xrightarrow{\mathbf{a} \cdot (\mathbf{b} \parallel \mathbf{c}) \cdot \mathbf{a}}_A q_5$. Since $q_5 \in F$, we have that $\mathbf{a} \cdot (\mathbf{b} \parallel \mathbf{c}) \cdot \mathbf{a} \in L_A(q_0)$, that is to say, q_0 accepts the pomset $\mathbf{a} \cdot (\mathbf{b} \parallel \mathbf{c}) \cdot \mathbf{a}$.

It is useful to distinguish runs based on the rules that induce them. To this end, we establish the following terminology for $q, q' \in Q$ and $U \in \text{SP}(\Sigma)$. If $q \xrightarrow{U}_A q'$ follows by an



(a) A PA with a parallel unit run, from q_1 to q_2 , which is labelled by a primitive pomset.

(b) A PA with unbounded parallelism as a result of a fork-cycle.

FIGURE 2. Some problematic pomset automata.

application of the first rule, we speak of a *trivial run*. Also, if $q \xrightarrow{U}_A q'$ has a derivation in which the second rule is applied last, this run is a *sequential unit run*. Furthermore, if $q \xrightarrow{U}_A q'$ is a consequence of the last rule, this run is a *parallel unit run*. The sequential and parallel unit runs are collectively referred to as *unit runs*. Lastly, if $q \xrightarrow{U}_A q'$ is a result of applying the third rule, i.e., there exist $U_1, U_2 \in \text{SP}(\Sigma)$ and $q'' \in Q$ such that $q \xrightarrow{U_1}_A q''$ as well as $q'' \xrightarrow{U_2}_A q'$, and neither of these is trivial, then $q \xrightarrow{U}_A q'$ is known as a *composite run*. By definition of \rightarrow_A , each run falls into at least one of these categories.

Example 4.4. Returning to Example 4.3 above, we find that $q_0 \xrightarrow{a}_A q_1$ is a sequential unit run, and $q_1 \xrightarrow{b||c}_A q_2$ is a parallel unit run. An example of a trivial run is $q_5 \xrightarrow{1}_A q_5$. Lastly, the run $q_0 \xrightarrow{a \cdot (b||c) \cdot a}_A q_5$ is a composite run.

Remark 4.5. The type of a run is not uniquely determined by the kind of pomset that labels it. For example, even though most parallel unit runs are labelled by a parallel pomset, this is not true in general: if A is the PA in Figure 2a, then we can construct the parallel unit run $q_1 \xrightarrow{a}_A q_2$, even though a is not parallel. Similarly, not every run labelled by the empty pomset is trivial. For instance, if $q, q' \in Q$ such that $q' \in \gamma(q, \emptyset)$, then $q \xrightarrow{1}_A q'$. We deal with this kind of confusion between run types in Section 5.

We reap the benefits of our new vocabulary with the following useful lemma.

Lemma 4.6. *Let $q \xrightarrow{U}_A q'$. There exist $q = q_0, \dots, q_\ell = q' \in Q$ and $U_1, \dots, U_\ell \in \text{SP}(\Sigma)$, such that $U = U_1 \cdots U_\ell$, and for all $1 \leq i \leq \ell$ we have that $q_{i-1} \xrightarrow{U_i}_A q_i$ is a unit run.*

The minimal ℓ for a given run as obtained above is known as the *length* of the run.

4.1. Fork-acyclicity. As it turns out, pomset automata are a rather powerful model of computation. For instance, they can be used to recognize context-free languages; consequently, language equivalence of states is undecidable in general [KBL⁺19]. The heart of the argument from op. cit. is that we can use the forks of a pomset automaton to simulate something that resembles the call stack of a recursive program. To prevent this excessive amount of expressive power, we need to put a structural restriction on PAs; specifically, what we will do is ensure that the level of (nested) forks in constructing any run is bounded from above.¹

¹Earlier automata models for sr-expressions had to apply similar restrictions [LW00, JM16].

Example 4.7. Suppose A is the PA in Figure 2b. We then find that $q_4 \xrightarrow{b}_A q_5$ and $q_3 \xrightarrow{a}_A q_1$. Since $q_2 \in \gamma(q_1, \{q_3, q_4\})$, we know that $q_1 \xrightarrow{a\|b}_A q_2$. However, because $q_3 \xrightarrow{a}_A q_1$, it follows that $q_3 \xrightarrow{a \cdot (a\|b)}_A q_2$, and hence we also find that $q_1 \xrightarrow{a \cdot (a\|b)\|b}_A q_2$. This pattern can be repeated indefinitely, leading to an unbounded number of forks in the construction of runs originating from q_1 . In addition, we note that $L_A(q_1)$ is not series-rational [LS14].

The problem in the example above is that q_1 can fork into a state q_3 , which can itself reach q_1 and fork again, giving rise to an unbounded number of nested forks. More generally, this behaviour can occur if q_3 can construct a run that somehow involves q_1 , through a series of transitions and forks. Hence, the first step to inhibit this kind of recursion is to get a handle on the states that can be involved in constructing a run that originates in a given state, by transitioning to that state, by forking into it, or some combination of the two.

Definition 4.8 (Support). We define \preceq_A as the smallest preorder on Q satisfying for $q \in Q$

$$\frac{a \in \Sigma \quad q' \in \delta(q, a)}{q' \preceq_A q} \quad \frac{\phi \in \mathbb{M}(Q) \quad q' \in \gamma(q, \phi)}{q' \preceq_A q} \quad \frac{r \in \phi \in \mathbb{M}(Q) \quad \gamma(q, \phi) \neq \emptyset}{r \preceq_A q}$$

We refer to \preceq_A as the *support relation* of A . This relation in turn gives rise to the *strict support relation* \prec_A , which is the strict order in which $q' \prec_A q$ holds if $q' \preceq_A q$ and $q \not\preceq_A q'$.

Example 4.9. Returning to the PA in Figure 2b, we see that $q_5 \preceq_A q_4$, since $q_5 \in \delta(q_4, b)$ (by the first rule). Because $q_2 \in \gamma(q_1, \{q_3, q_4\})$, it follows that $q_2 \preceq_A q_1$ (by the second rule) as well as $q_3, q_4 \preceq_A q_1$ (by the third rule). By transitivity, it then follows that $q_5 \preceq_A q_1$.

Intuitively, if q' is necessary to establish some run originating from q , then $q' \preceq_A q$; hence, we say that q' *supports* q . In particular, if $r \preceq_A q$ because there exists a $\phi \in \mathbb{M}(Q)$ with $r \in \phi$ and $\gamma(q, \phi) \neq \emptyset$, then r serves as the start of one or more threads that q may fork into, and we say that r is a *fork target* of q . Support can be mutual, for instance when $q' \in \delta(q, a)$ and $q \in \delta(q', a)$; consequently, \preceq_A need not be antisymmetric.

To break fork cycles, we can define the restriction necessary to avoid infinitely nested forks, by stipulating that a state cannot be supported by any of its fork targets.

Definition 4.10 (Fork-acyclicity). We say that A is *fork-acyclic* if for $q, r \in Q$ such that there exists a $\phi \in \mathbb{M}(Q)$ with $r \in \phi$ and $\gamma(q, \phi) \neq \emptyset$, we have that $r \prec_A q$, i.e., $q \not\preceq_A r$.

Example 4.11. Returning to the PA in Figure 2b, we see that $\gamma(q_1, \{q_3, q_4\}) \neq \emptyset$, while $q_1 \preceq_A q_3$. Hence, this PA is not fork-acyclic. On the other hand, the PA in Figure 2a is fork-acyclic, because neither q_3 nor q_4 is supported by q_1 .

4.2. Boundedness. Our primary interest is pomset automata with finitely many states, known as *finite pomset automata*. It is sometimes convenient to relax this property and speak of PAs with infinitely many states, but where each state relies on only finitely many states to establish its runs. To formalize this, we introduce the following notions.

Definition 4.12 (Support and boundedness). We say that $Q' \subseteq Q$ is *support-closed* if for all $q \in Q'$ with $q' \preceq_A q$ we have $q' \in Q'$. The *support* of $q \in Q$, denoted $\pi_A(q)$, is the smallest support-closed set containing q . When $\pi_A(q)$ is finite for all $q \in Q$, we say that A is *bounded*.

Example 4.13. In Figure 2b, the set $\{q_2, q_4, q_5\}$ is support-closed, while the set $\{q_3\}$ is not, since $q_1 \preceq_A q_3$. The support of q_1 is given by the set of all states, while the support of q_4 is given by $\{q_4, q_5\}$. Like all finite PAs, this PA is bounded.

When working with bounded and fork-acyclic automata, it is also useful to reason based on how many nested forks can occur in a single run. To this end, the following is useful.

Definition 4.14 (Depth). If q is a state in A , we write $D_A(q)$ for the *depth of q in A* , which is the maximum n such that there exist $q_1, \dots, q_n \in Q$ with $q_1 \prec_A q_1 \prec_A \dots \prec_A q_n = q$, if such an n exists; otherwise, $D_A(q)$ is undefined. If Q is finite and $D_A(q)$ is defined for every $q \in Q$, we write D_A for the *depth of A* , i.e., the maximum of $D_A(q)$ for all $q \in Q$.

If A is bounded, then one can use the definition of \prec_A to argue that $D_A(q)$ is defined for every state q . Moreover, if A is finite, then D_A is defined.

Example 4.15. The PA in Figure 2a has depth 2, since $q_3 \prec_A q_1$. The PA in Figure 2b has depth 3, as $q_5 \prec q_4 \prec q_1$.

4.3. Implementation. We will perform a number of transformations on automata to enforce desirable properties. To ensure correctness, we require these constructions to transform an automaton A into an automaton A' implementing A , in the following sense:

Definition 4.16. Let $A = \langle Q, F, \delta, \gamma \rangle$ and $A' = \langle Q', F', \delta', \gamma' \rangle$ be pomset automata. We say that A' *implements* A if the following hold:

- (i) $Q \subseteq Q'$, and for every $q \in Q$ it holds that $L_A(q) = L_{A'}(q)$, and
- (ii) if A is fork-acyclic, then so is A' .

Any PA A can be restricted to a support-closed set of states, such that A implements the restricted PA. In particular, if A is bounded, this means that the language of any particular state q can also be described by a finite PA, simply by restricting A to the support of q ; moreover, if A is bounded, then so is this new PA. Formally, we have the following.

Lemma 4.17. *If $Q' \subseteq Q$ is support-closed, then we can construct a PA $A[Q']$ with states Q' such that A implements $A[Q']$. Moreover, if A is bounded (resp. finite), then so is $A[Q']$.*

Proof. Since Q' is support-closed, we can regard δ as a function of type $Q' \times \Sigma \rightarrow 2^{Q'}$, and γ as a function of type $Q' \times \mathbb{M}(Q') \rightarrow 2^{Q'}$. We choose $A[Q'] = \langle Q', Q' \cap F, \delta, \gamma \rangle$.

We should show that, for all $q \in Q'$, it holds that $L_{A[Q']}(q) = L_A(q)$. The inclusion from left to right should be clear: any (accepting) run in $A[Q']$ can be replayed in A by construction. For the other inclusion, we prove more generally that if $q \in Q'$ and $q \xrightarrow{U}_A q'$ then $q' \in Q'$ and $q \xrightarrow{U}_{A[Q']} q'$, by induction on \rightarrow_A . In the base, there are two cases.

- If $U = 1$ and $q = q'$, then $q' \in Q'$ and $q \xrightarrow{U}_{A[Q']} q'$ immediately.
- If $U = \mathbf{a}$ for some $\mathbf{a} \in \Sigma$, and $q' \in \delta(q, \mathbf{a})$, then $q' \preceq_A q$, and hence $q' \in Q'$. Furthermore, we find that $q \xrightarrow{U}_{A[Q']} q'$.

For the inductive step, there are two cases to consider.

- Suppose that $q \xrightarrow{U}_A q'$ because $U = V \cdot W$, and there exists a $q'' \in Q$ such that $q \xrightarrow{V}_A q''$ and $q'' \xrightarrow{W}_A q'$. It follows that $q'' \in Q'$ and $q \xrightarrow{V}_{A[Q']} q''$ by induction. Similarly, we find $q' \in Q'$ and $q'' \xrightarrow{W}_{A[Q']} q'$, again by induction. We then conclude that $q \xrightarrow{U}_{A[Q']} q'$.
- Suppose that $q \xrightarrow{U}_A q'$ because $U = U_1 \parallel \dots \parallel U_n$, and there exist $r_1, \dots, r_n \in Q$ and $q'_1, \dots, q'_n \in F'$ such that for $1 \leq i \leq n$ we have $q_i \xrightarrow{U_i}_A q'_i$, as well as $q' \in \gamma(q, \{\!\{r_1, \dots, r_n\}\!\})$. We have $r_1, \dots, r_n \preceq_A q$, and thus $r_1, \dots, r_n \in Q'$, as well as $q' \in Q'$. By induction, we then find for $1 \leq i \leq n$ that $r'_i \in Q'$ and $r_i \xrightarrow{U_i}_{A[Q']} r'_i$. We conclude that $q \xrightarrow{U}_{A[Q']} q'$.

A variation of the above argument shows that for $q, q' \in Q$, we have $q \preceq_{A[Q']} q'$ if and only if $q \preceq_A q'$ and $q, q' \in Q'$; hence, if A is fork-acyclic, then so is $A[Q']$. This also shows that, for $q \in Q'$, we have that $\pi_{A[Q']}(q) = \pi_A(q)$; hence, if A is bounded, then so is $A[Q']$. Finally, if A is finite, then clearly Q' is finite, and hence $A[Q']$ has finitely many states. \square

5. LANGUAGE EQUIVALENCE

We consider the following decision problem: given a PA $A = \langle Q, F, \delta, \gamma \rangle$ and states $q, q' \in Q$, do q and q' accept the same language? As stated previously, this problem is undecidable for finite PAs that need not be fork-acyclic [KBL⁺19]. In this section, we investigate whether there exists a general decision procedure for fork-acyclic finite PAs.

To fully appreciate the complexity of this problem, we start by illustrating the intricacies of PAs through a series of examples. These show how PAs with very different structures accept the same language. Any procedure to decide language equivalence for fork-acyclic PAs must take such cases into account.

Example 5.1 (Run confusion). In Figure 2a, we find that both q_1 and q_3 accept the singleton language $\{\mathbf{a}\}$. However, q_1 and q_3 have a very different transition structure, since $\delta(q_1, \mathbf{a}) = \emptyset$ and $\delta(q_3, \mathbf{a}) = \{q_5\}$, with $q_5 \in F$. Indeed, q_1 accepts \mathbf{a} by means of a parallel unit run (forking to q_3 and q_4), while q_3 accepts \mathbf{a} by means of a sequential unit run.

More generally, a state could accept a more complicated pomset by means of a composite run (i.e., of length greater than one), while another state accepts the same pomset with a parallel unit run (i.e., of length one) forking into one or more accepting states. A similar phenomenon occurs when a state forks into a multiset of size one.

Example 5.2 (Empty forks). The definition of pomset automata does not prohibit forks into the empty multiset; in a sense, these are analogous to ϵ -transitions in NFAs, since the parallel composition of zero pomsets is 1. This could allow a non-accepting state to accept the empty pomset, or a state without parallel transitions to accept a parallel pomset.

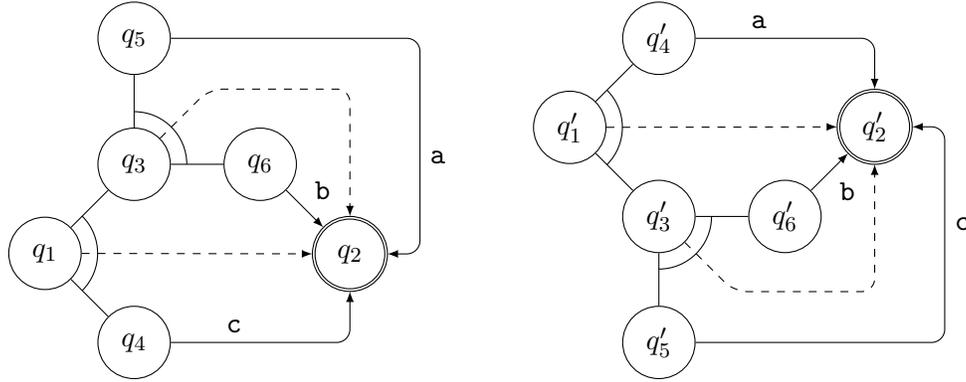
Hence, different types of runs may be labelled by the same pomset, especially when fork targets can be accepting, or when unary or nullary forks are allowed. The next example is about how nested forks may encode the same behaviour differently.

Example 5.3 (Associativity). Consider the PA in Figure 3a, where both q_1 and q'_1 accept the language $\{\mathbf{a} \parallel \mathbf{b} \parallel \mathbf{c}\}$. In the transition from q_1 to q_2 , the pomset $\mathbf{a} \parallel \mathbf{b}$ is contributed by q_3 , and \mathbf{c} comes from q_4 , while in the transition from q'_1 to q'_2 , we obtain \mathbf{a} from q'_4 , and $\mathbf{b} \parallel \mathbf{c}$ from q'_3 . The language of q_3 is distinct from, and in fact incomparable with, the languages of q'_3 and q'_4 . Hence, to compare the language of q_1 to that of q'_1 , we need not only consider the states they may fork into, but also the states where those states may fork into, provided that the second-level forks land in an accepting state.

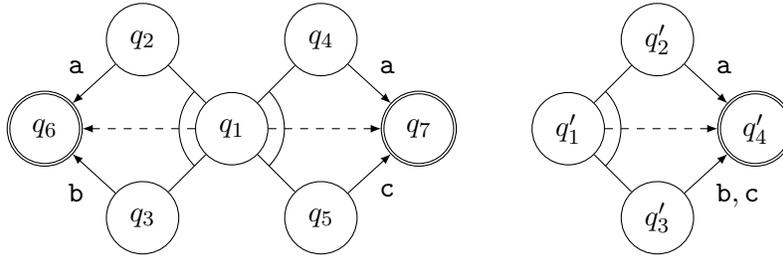
We can counteract the phenomena exhibited in Examples 5.1 to 5.3, by preventing the structures that enable this kind of behaviour. This leads to the following definition.

Definition 5.4 (Well-structured). A pomset automaton $A = \langle Q, F, \delta, \gamma \rangle$ is *well-structured* if for $q, q' \in Q$, $\phi \in \mathbb{M}(Q)$ with $q' \in \phi$ and $\gamma(q, \phi) \neq \emptyset$, all of the following hold:

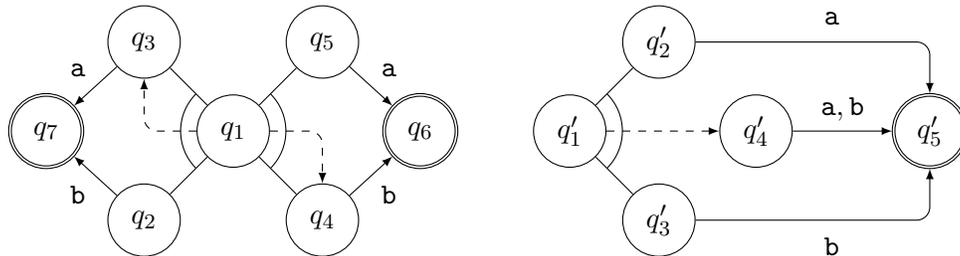
$$|\phi| \geq 2, \quad q' \notin F, \quad \text{and} \quad \forall \psi \in \mathbb{M}(Q). \gamma(q', \psi) \cap F = \emptyset.$$



(a) Associativity of parallelism.



(b) Distributivity of parallelism over union.



(c) Distributivity of sequential composition over union.

FIGURE 3. Examples of PAs where q_1 and q'_1 accept the same language, while having different transition structures.

Example 5.5. The PAs in Figures 1 and 2b are well-structured. On the other hand, a PA with a state q such that $\gamma(q, \square) \neq \emptyset$, or $\gamma(q, \{r\}) \neq \emptyset$ for some state r , is not well-structured (by the first condition). The PA in Figure 2a is not well-structured (by the second condition) because $\gamma(q_1, \{q_3, q_4\}) \neq \emptyset$, while $q_4 \in F$. Finally, the PA in Figure 3a is not well-structured (by the third condition) because $\gamma(q_1, \{q_3, q_4\}) \neq \emptyset$, while $q_2 \in \gamma(q_3, \{q_5, q_6\}) \cap F$.

Well-structuredness does not diminish the expressive power of PAs, as we shall formally prove in Section 6. For instance, the behaviour of the PA discussed in Example 5.3 (see Figure 3a) can be expressed using a ternary fork, and the behaviour of the PA discussed in Example 5.1 (see Figure 2a) can be obtained by a sequential transition. For now, we focus our decision procedure on the fragment of well-structured PAs.

One of the benefits of well-structured PAs is that the type of a run is uniquely determined by the kind of its pomset, as spelled out by the following lemma.

Lemma 5.6. *If A is well-structured and $q \xrightarrow{U}_A q'$, then the following hold:*

- (i) $q \xrightarrow{U}_A q'$ is trivial if and only if U is empty.
- (ii) $q \xrightarrow{U}_A q'$ is a sequential unit run if and only if U is primitive.
- (iii) $q \xrightarrow{U}_A q'$ is a composite run if and only if U is sequential.
- (iv) $q \xrightarrow{U}_A q'$ is a parallel unit run if and only if U is parallel.

Proof. We start by treating (i) in detail. Here, we note that the implication from left to right holds by definition. For the proof from right to left, we proceed by induction on the construction of $q \xrightarrow{U}_A q'$. In the base, we need only consider the case where $q \xrightarrow{U}_A q'$ is already trivial. For the inductive step, there are two cases to consider.

- Suppose that $U = V \cdot W$ and there exists a $q'' \in Q$ such that $q \xrightarrow{V}_A q''$ and $q'' \xrightarrow{W}_A q'$, then $V = W = 1$. By induction, we then know that $q \xrightarrow{V}_A q''$ and $q'' \xrightarrow{W}_A q'$ are trivial, and hence $q = q'' = q'$. We conclude that $q \xrightarrow{U}_A q'$ must also be trivial.
- Suppose that $U = U_1 \parallel \dots \parallel U_n$ and there exist $r_1, \dots, r_n \in Q$ as well as $r'_1, \dots, r'_n \in F$ such that for $1 \leq i \leq n$ we have $r_i \xrightarrow{U_i}_A r'_i$, and furthermore $q' \in \gamma(q, \{r_1, \dots, r_n\})$. Then necessarily $U_1, \dots, U_n = 1$. Since A is well-structured, we also know that $r_1, \dots, r_n \notin F$, and furthermore that $n \geq 2$. However, by induction, we know that for $1 \leq i \leq n$ it holds that $r_i \xrightarrow{U_i}_A r'_i$ is trivial, and hence it would follow that $r_i = r'_i$. We have now reached a contradiction, for $r_1 \in F$ while also $r_1 \notin F$. We can therefore disregard this case.

We treat the implications from left to right for the remaining claims as follows.

- For (ii), we find that U is primitive by definition of sequential unit runs.
- For (iii), suppose that $q \xrightarrow{U}_A q'$ is composite. We then know that $U = V \cdot W$ and there exists a $q'' \in Q$ such that $q \xrightarrow{V}_A q''$ and $q'' \xrightarrow{W}_A q'$ are nontrivial. By (i), we then know that V and W must be non-empty, and hence U is sequential.
- For (iv), suppose that $q \xrightarrow{U}_A q'$ is a parallel unit run. We then know that $U = U_1 \parallel \dots \parallel U_n$ and there exist $r_1, \dots, r_n \in Q$ and $r'_1, \dots, r'_n \in F$, such that for $1 \leq i \leq n$ it holds that $r_i \xrightarrow{U_i}_A r'_i$, and furthermore $q' \in \gamma(q, \{r_1, \dots, r_n\})$. By the premise that A is well-structured, we know that $r_1, \dots, r_n \notin F$ and $n \geq 2$. It then follows that for $1 \leq i \leq n$, the run $r_i \xrightarrow{U_i}_A r'_i$ is non-trivial, and hence U_i is non-empty by the above. From this and the fact that $n \geq 2$, we can conclude that U is parallel.

The implications from right to left for the latter three claims now follow from Lemma 3.6. For instance, if U is primitive, then $q \xrightarrow{U}_A q'$ must be a sequential unit run, for if it were trivial then U would be empty, if it were composite then U would be sequential, and if it were a parallel unit run then U would be parallel. \square

Furthermore, because of the restriction on fork targets in a well-structured PA, pomsets in their languages are not parallel:

Lemma 5.7. *If A is well-structured and q is a fork target in A , then all pomsets in $L_A(q)$ are parallel primes.*

Even for well-structured PAs, however, there are still structural factors confounding language equivalence with which a decision procedure will need to reckon.

Example 5.8 (Distributivity I). In Figure 3b, we have the state q_1 which may fork into q_2 and q_3 , as well as q_4 and q_5 . Now, q_1 accepts the language $\{\mathbf{a} \parallel \mathbf{b}, \mathbf{a} \parallel \mathbf{c}\}$, where the former behaviour stems from forking into q_2 and q_3 , but the latter is obtained by forking into q_4 and q_5 . On the other hand, we have the state q'_1 , which may fork into q'_2 and q'_3 ; here, q'_1 accepts the same language as q_1 , but the two pomsets are due to the *same* fork. This is a consequence of the distributivity of parallel composition over union of pomset languages:

$$L_A(q_1) = (\{\mathbf{a}\} \parallel \{\mathbf{b}\}) \cup (\{\mathbf{a}\} \parallel \{\mathbf{c}\}) = \{\mathbf{a}\} \parallel (\{\mathbf{b}\} \cup \{\mathbf{c}\}) = L_A(q'_1)$$

This example illustrates how behaviour of language-equivalent states may be spread out across different parallel transitions, and that this division may differ locally.

The last example stems from an implicit kind of non-determinism that is supported by PAs, as a result of overlap between the languages of fork targets.

Example 5.9 (Distributivity II). In Figure 3c, q_1 can read $\mathbf{a} \parallel \mathbf{b}$ to arrive in *either* q_3 or q_4 . From that point on, q_3 can read \mathbf{a} to reach an accepting state, while q_4 can read \mathbf{b} to do the same. In contrast, q'_1 can read $\mathbf{a} \parallel \mathbf{b}$ to arrive in *only one state*, q'_4 , whence it can read either \mathbf{a} or \mathbf{b} to arrive in q'_5 and accept. Nevertheless, q_1 and q'_1 accept the same language. Even though the behaviour implemented by the forks from q_1 is the same, the state where they land is not uniquely determined. In a sense, this is a consequence of distributivity of sequential composition over union of pomset languages:

$$L_A(q_1) = (\{\mathbf{a} \parallel \mathbf{b}\} \cdot \{\mathbf{a}\}) \cup (\{\mathbf{a} \parallel \mathbf{b}\} \cdot \{\mathbf{b}\}) = \{\mathbf{a} \parallel \mathbf{b}\} \cdot (\{\mathbf{a}\} \cup \{\mathbf{b}\}) = L_A(q'_1)$$

Examples 5.8 and 5.9 illustrate that there are many structurally different ways to express parallel behaviour in a pomset automaton. It is not so clear how to choose for one particular way of representing parallel behaviour. Instead, we design our algorithm to directly deal with the illustrated equivalences. To understand how this can be done, it is convenient to first shift perspective from trying to decide language equivalence to trying to find out which states do and do not overlap in terms of their language [LS14].

Definition 5.10 (Atoms). Let A be a PA or an NFA, with states Q and $\alpha \subseteq Q$. We write

$$L_A(\alpha) = \left(\bigcap_{q \in \alpha} L_A(q) \right) \setminus \left(\bigcup_{q \notin \alpha} L_A(q) \right)$$

When $L_A(\alpha) \neq \emptyset$, we say that α is an *atom* of A ; we write At_A for the set of atoms of A .

Example 5.11. In the PA in Figure 3b, $\{q_3, q'_3\}$ is an atom, as is $\{q_5, q'_3\}$; the set $\{q_2, q_4, q'_2\}$ is also an atom. In fact, these are all atoms that contain fork targets of Figure 3b.

The PA in Figure 3c has $\{q_3, q_5, q'_2, q'_4\}$ as an atom; similarly, $\{q_2, q_4, q'_3, q'_4\}$ is an atom. These are again all of the atoms that contain fork targets in Figure 3c.

The following lemma shows how atoms can be used to decide language equivalence.

Lemma 5.12. *Let A be a PA or NFA with states q_1 and q_2 . Then $L_A(q_1) = L_A(q_2)$ if and only if for all $\alpha \in \text{At}_A$ it holds that $q_1 \in \alpha$ precisely when $q_2 \in \alpha$.*

Proof. First, suppose that q_1 and q_2 have the same language, and let $\alpha \in \text{At}$ with $q_1 \in \alpha$. Then surely $q_2 \in \alpha$, because otherwise we have that $L_A(\alpha) \subseteq L_A(q_1) \setminus L_A(q_2) = \emptyset$, which contradicts that α is an atom. Similarly, we have that $q_2 \in \alpha$ implies $q_1 \in \alpha$.

For the other implication, let $U \in L_A(q_1)$. Choose $\alpha = \{q \in Q : U \in L_A(q)\}$, and note that α is an atom, since $U \in L_A(\alpha)$ by construction. Because $q_1 \in \alpha$, we know that $q_2 \in \alpha$ by the premise, and thus (by the above), we have that $U \in L_A(q_2)$. This shows that $L_A(q_1) \subseteq L_A(q_2)$; the other inclusion follows by symmetry. \square

Thus, if we can compute the set of atoms of a PA then we can decide language equivalence of states. As it turns out, this is possible if the PA is finite, fork-acyclic and well-structured.

Lemma 5.13. *If A is finite, fork-acyclic, and well-structured, then At_A is computable.*

Proof. We proceed by induction on D_A (c.f. Definition 4.14). In the base, where $D_A = 0$, we have $Q = \emptyset$, since if $q \in Q$, then $D_A(q) \geq 1$; hence, \emptyset is the only atom in this case.²

For the inductive step, let $D_A > 0$ and suppose that the claim holds for pomset automata of strictly smaller depth. We choose $Q' = \{q \in Q : D_A(q) < D_A\}$, and note that Q' is support-closed: if $q' \preceq_A q \in Q'$, then $D_A(q') \leq D_A(q)$, and hence $q' \in Q'$. By Lemma 4.17, we can then restrict A to obtain $A[Q']$, which, by construction, is of depth strictly less than D_A . By applying the induction hypothesis, we can compute the atoms of $A[Q']$.

To compute the atoms of A proper, we shall construct an NFA $A' = \langle Q, \delta', F \rangle$ whose atoms are precisely those of A ; the claim then follows because we can compute the atoms of an NFA using a standard algorithm such as the one by Brzozowski and Tamm [BT14]. The idea behind this NFA is that it contains the δ -transitions of A , and it encodes the γ -transitions by transitions labelled with symbols built from the atoms of $A[Q']$.

The alphabet of our NFA will contain the symbols from Σ , as well as additional symbols that encode (parts of) the languages of pomsets that can label parallel unit transitions, as divided up by the atoms of $A[Q']$; more precisely, we choose

$$\Delta = \Sigma \cup \left\{ \{\alpha_1, \dots, \alpha_n\} \in \mathbb{M}(\text{At}_{A[Q']}) : \begin{array}{l} \exists q \in Q, q_1 \in \alpha_1, \dots, q_n \in \alpha_n. \\ \gamma(q, \{q_1, \dots, q_n\}) \neq \emptyset \end{array} \right\}$$

Here, we assume without loss of generality that the two sets are disjoint, i.e., that none of the multisets are already symbols in Σ . This alphabet is finite because $\text{At}_{A[Q']}$ is finite, and because by definition of pomset automata, there are finitely many multisets ϕ s.t. $\gamma(q, \phi) \neq \emptyset$.

We define the sequential transition function $\delta' : Q \times \Delta \rightarrow 2^Q$, as follows.

$$\delta'(q, \mathbf{a}) = \delta(q, \mathbf{a}) \quad \delta'(q, \{\alpha_1, \dots, \alpha_n\}) = \bigcup \{ \gamma(q, \{q_1, \dots, q_n\}) : q_1 \in \alpha_1, \dots, q_n \in \alpha_n \}$$

Example 5.14. Let A be the PA in Figure 3b. In Example 5.11, we saw that $\alpha_1 = \{q_3, q'_3\}$, $\alpha_2 = \{q_5, q'_3\}$ and $\alpha_3 = \{q_2, q_4, q'_2\}$ are the atoms that contain fork targets. The resulting NFA is drawn in Figure 4a. There, we see that $q_6 \in \delta'(q_2, \mathbf{a})$ because $q_6 \in \delta(q_2, \mathbf{a})$. Furthermore, $q_6 \in \delta'(q_1, \{\alpha_3, \alpha_1\})$ because $q_6 \in \gamma(q_1, \{q_2, q_3\})$ and $q_2 \in \alpha_3$ while $q_3 \in \alpha_1$.

Example 5.15. Let A be the PA in Figure 3c. In Example 5.11, we found atoms $\alpha_1 = \{q_3, q_5, q'_2, q'_4\}$ and $\alpha_2 = \{q_2, q_4, q'_3, q'_4\}$. The resulting NFA is drawn in Figure 4b. There, we see that $q_7 \in \delta'(q_1, \{\alpha_1, \alpha_2\})$ because $q_7 \in \gamma(q_1, \{q_3, q_2\})$ with $q_3 \in \alpha_1$ and $q_2 \in \alpha_2$.

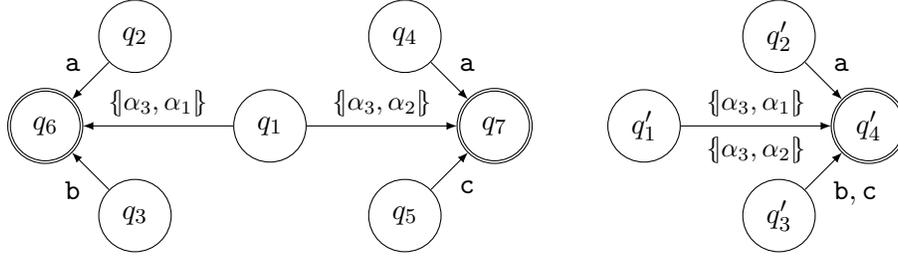
Having defined our target NFA, it remains to show that the atoms of A are the same as those of A' . To this end, we need to relate the languages of A' to the languages of A ; we do this by means of the substitution $\zeta : \Delta \rightarrow 2^{\text{SP}(\Sigma)}$, given by:

$$\zeta(\mathbf{a}) = \{\mathbf{a}\} \quad \zeta(\{\alpha_1, \dots, \alpha_n\}) = L_A(\alpha_1) \parallel \dots \parallel L_A(\alpha_n)$$

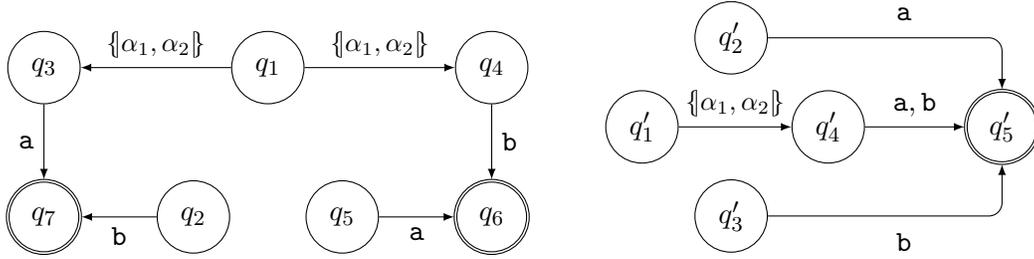
We need the following two technical properties of ζ . The first of these relates the languages of the states of the NFA A' to the languages of the states of the PA A by means of ζ .

Fact 5.16. *For all $q \in Q$, it holds that $L_A(q) = \zeta(L_{A'}(q))$.*

²The empty intersection is assumed to be the set of all sp-pomsets $\text{SP}(\Sigma)$.



(a) NFA created in inductive step of atom computation for Figure 3b.



(b) NFA created in inductive step of atom computation for Figure 3c.

FIGURE 4. Examples of PAs obtained in the inductive step of atom computation.

The second property that we need says that ζ is atomic; this is a consequence of the fact that A is well-structured, and hence if $\{\alpha_1, \dots, \alpha_n\} \in \Delta$, then $n \geq 2$ and $1 \notin L_A(\alpha_i)$ for all $1 \leq i \leq n$, meaning that $\zeta(\{\alpha_1, \dots, \alpha_n\})$ must consist of sequential primes.

Fact 5.17. *The substitution ζ is atomic.*

Let $\alpha \subseteq Q$; we can then use the above observations to calculate that

$$\begin{aligned}
 \zeta(L_{A'}(\alpha)) &= \zeta\left(\left(\bigcap_{q \in \alpha} L_{A'}(q)\right) \setminus \left(\bigcup_{q \notin \alpha} L_{A'}(q)\right)\right) && \text{(def. } L_{A'} \text{ on sets of states)} \\
 &= \left(\bigcap_{q \in \alpha} \zeta(L_{A'}(q))\right) \setminus \left(\bigcup_{q \notin \alpha} \zeta(L_{A'}(q))\right) && \text{(Fact 5.17 and Lemma 3.12)} \\
 &= \left(\bigcap_{q \in \alpha} L_A(q)\right) \setminus \left(\bigcup_{q \notin \alpha} L_A(q)\right) && \text{(Fact 5.16)} \\
 &= L_A(\alpha) && \text{(def. } L_A \text{ on sets of states)}
 \end{aligned}$$

To wrap up, suppose α is an atom of A ; then $\zeta(L_{A'}(\alpha)) = L_A(\alpha) \neq \emptyset$, and hence $L_{A'}(\alpha) \neq \emptyset$ by Fact 5.17 and Lemma 3.12, making α an atom of A' . Conversely, if α is an atom of A' then $L_A(\alpha) = \zeta(L_{A'}(\alpha)) \neq \emptyset$ by Fact 5.17 and Lemma 3.12, meaning α is an atom of A . \square

Lemmas 5.12 and 5.13 now give us the decidability result for well-structured PAs.

Theorem 5.18. *Let A be a finite, fork-acyclic and well-structured PA. Given states q_1 and q_2 , it is decidable whether $L_A(q_1) = L_A(q_2)$.*

The complexity of the above procedure is dominated by the complexity of calculating the atoms of the NFA constructed in each inductive step. Brzozowski and Tamm's algorithm

is doubly exponential in the number of states [BT14], which implies a rough upper bound of $O(D_A \cdot 2^{2^{|Q|}})$ on our algorithm. Despite this, the algorithm in [BT14] is known to have good performance in practice [AMR07].

6. WELL-STRUCTURED AUTOMATA

In this section we show that for every finite and fork-acyclic automaton, we can construct a finite, fork-acyclic, and *well-structured* automaton that implements it. More precisely, we show that this transformation can take a *bounded* and fork-acyclic automaton, and produce a bounded, fork-acyclic and well-structured automaton; the former claim then follows, since every finite automaton is bounded, and every bounded automaton can be made finite while preserving fork-acyclicity, by Lemma 4.17.³ We first decompose the definition of well-structured automata into three simpler properties.

Definition 6.1 (*n*-forking). Let $n \in \mathbb{N}$. A pomset automaton A is *n-forking* if for every state $q \in Q$ and every multiset of states $\phi \in \mathbb{M}(Q)$ such that $\gamma(q, \phi) \neq \emptyset$ we have $|\phi| \geq n$.

Definition 6.2 (Parsimony). A pomset automaton A is said to be *parsimonious* if, whenever $p \in Q$ and $q \in \phi \in \mathbb{M}(Q)$ such that $\gamma(p, \phi) \neq \emptyset$, we have that $1 \notin L_A(q)$.

Definition 6.3 (Flat-branching). A pomset automaton A is *flat-branching* if for all states $q, q' \in Q$ and every multiset $\phi \in \mathbb{M}(Q)$, if $\gamma(q, \phi) \neq \emptyset$ and $q' \in \phi$ then

$$\forall \psi \in \mathbb{M}(Q). \gamma(q', \psi) \cap F = \emptyset.$$

Example 6.4. All automata displayed in Figures 1 to 3 are 2-forking. The automata in Figure 3 are parsimonious, but the ones in Figure 2 are not; for instance, in Figure 2a, we have that $\gamma(q_1, \{q_3, q_4\}) \neq \emptyset$, while $1 \in L_A(q_4)$. The automata in Figures 2, 3b and 3c are flat-branching, but the one in Figure 3a is not; in particular, in that automaton we have that $\gamma(q_1, \{q_3, q_4\}) \neq \emptyset$, while $q_2 \in \gamma(q_3, \{q_5, q_6\}) \cap F$.

One can prove that the above properties are indeed equivalent to well-structuredness.

Lemma 6.5. *A pomset automaton A is well-structured if and only if it is 2-forking, parsimonious, and flat-branching.*

Our task can therefore be reduced to converting a bounded and fork-acyclic PA into an equivalent bounded and fork-acyclic PA that is also 2-forking, parsimonious and flat-branching. Note that these properties can be at cross purposes: for instance, ensuring parsimony may introduce forks of smaller sizes, which could make it so that the automaton is no longer 2-forking. Similarly, eliminating forks into the empty multiset may introduce new accepting states, which can invalidate flat-branching.

Thus, establishing all properties simultaneously requires some care. The remainder of this section describes a series of transformations that establish one property while maintaining the ones already established. More precisely, our construction to convert a bounded PA A into a bounded, 2-forking, parsimonious and flat-branching PA goes as follows:

- (1) we first show how to implement A with a parsimonious automaton A_0 ;
- (2) then we discuss how to implement A with a 1-forking automaton A_1 ;
- (3) we proceed to show how to implement A with a 2-forking automaton A_2 ;
- (4) finally, we show that A can be implemented by a flat-branching PA A_3 .

³It is clear that the construction from this lemma also preserves well-structuredness.

Since each of these transformations preserves the established properties (e.g., A_2 is still bounded, 1-forking and parsimonious), we end up with a bounded PA that implements A and is 2-forking, parsimonious, and flat-branching, and hence well-structured by Lemma 6.5.

Before we get into the weeds, we discuss some technical properties that will help simplify the constructions. First, note that we shall need to rule out a form of ϵ -transitions. Indeed, although our model does not include such transitions explicitly, they may be achieved by branching to states accepting the empty pomset (if the automaton is not parsimonious), or by using a $\gamma(p, \sqcup)$ -transition (if the automaton is not 1-forking). The result is a run $p \xrightarrow{1}_A q$ that is non-trivial. To reason about such transitions, we observe the following:

Lemma 6.6. *For any bounded PA A , the predicate $p \xrightarrow{1}_A q$ is decidable.*

Another useful notion for this section will be that of *weak implementation*. Essentially, a weak implementation of a PA is another PA where the behaviour of each state of the first PA is spread out across a set of states, rather than just one (as is the case for implementation).

Definition 6.7. A PA A' *weakly implements* a PA A if the following hold:

- (i) for every state q in A there is a finite set of states Q_q in A' s.t. $L_A(q) = \bigcup_{x \in Q_q} L_{A'}(x)$.
- (ii) if A is fork-acyclic, then so is A' .

To prove that there exists an automaton implementing A that satisfies some of the properties above, it suffices to find one that weakly implements A .

Lemma 6.8. *If a PA A' weakly implements A , then we can construct a PA A'' implementing A . If A' is bounded (resp. n -forking, flat-branching, parsimonious), then so is A'' .*

6.1. Parsimony. Let $A = \langle Q, F, \delta, \gamma \rangle$ be a bounded and fork-acyclic PA; we want to implement A with a fork-acyclic and parsimonious PA A_0 . There are two key ideas to this translation:

- We introduce a new state \top , which will be the only accepting state of the automaton; in fact, it will be the only state accepting the empty pomset. This state will not have any outgoing transitions, so its language is exactly $\{1\}$. We will modify the transition functions, such that if a transition in A can lead to a state that accepts 1, it can also non-deterministically lead to \top in A_0 .
- To ensure our construction preserves successful runs, we need to add γ transitions to mitigate the previous modification. More precisely, if $p \in Q$ can fork into $\phi \sqcup \psi \in \mathbb{M}(Q)$ to reach $q \in Q$, i.e., $q \in \gamma(p, \phi \sqcup \psi)$, and all states in ψ accept the empty pomset in A , then we make sure that, in A_0 , q can also fork into ϕ to reach p , simulating the acceptance of the empty pomset from states in ψ .

Doing so, we obtain an automaton weakly implementing A — indeed, if $1 \notin L_A(q)$, then $L_A(q) = L_{A_0}(q)$, and otherwise $L_A(q) = L_{A_0}(q) \cup L_{A_0}(\top)$. Since \top cannot be a fork target, and any other state cannot accept the empty pomset, this new automaton is parsimonious. Finiteness and fork-acyclicity are also maintained by this construction.

Definition 6.9 (A_0). The PA A_0 is given by the tuple $\langle Q_0, F_0, \delta_0, \gamma_0 \rangle$ where $Q_0 = Q \cup \{\top\}$ (with $\top \notin Q$), and $F_0 = \{\top\}$. Furthermore, δ_0 is generated by the rules

$$\frac{q' \in \delta(q, \mathbf{a})}{q' \in \delta_0(q, \mathbf{a})} \qquad \frac{q' \in \delta(q, \mathbf{a}) \quad 1 \in L_A(q')}{\top \in \delta_0(q, \mathbf{a})}$$

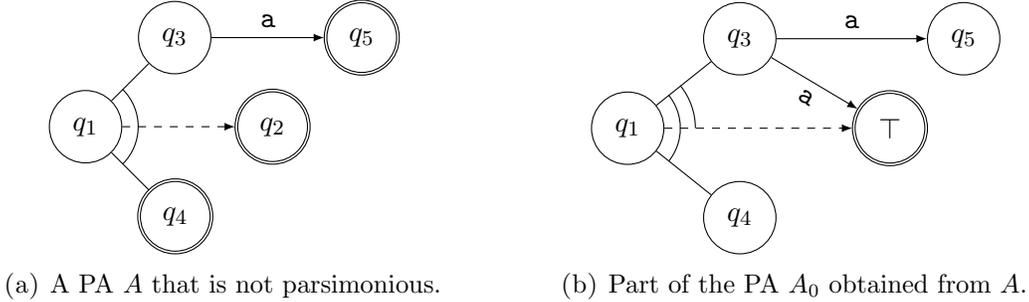


FIGURE 5. Example of construction to ensure parsimony.

Also, γ_0 is generated by the following rules for all $q \in Q$ and $\phi \in \mathbb{M}(Q)$:

$$\frac{q' \in \gamma(q, \phi \sqcup \psi) \quad \forall r \in \psi. 1 \in L_A(r)}{q' \in \gamma_0(q, \phi)} \quad \frac{q' \in \gamma(q, \phi \sqcup \psi) \quad 1 \in L_A(q') \quad \phi \neq \square \quad \forall r \in \psi. 1 \in L_A(r)}{\top \in \gamma_0(q, \phi)}$$

Lastly, $\delta_0(\top, \mathbf{a}) = \emptyset$ for all $\mathbf{a} \in \Sigma$ and $\gamma_0(\top, \phi) = \emptyset$ for all $\phi \in \mathbb{M}(Q_0)$.

Example 6.10. Recall the automaton A from Figure 2a, depicted in Figure 5a. Part of the resulting PA A_0 is drawn in Figure 5b. For instance, since $q_5 \in \delta(q_3, \mathbf{a})$ in A , we have $q_5 \in \delta_0(q_3, \mathbf{a})$ and $\top \in \delta_0(q_3, \mathbf{a})$ in A_0 by the first and second rules generating δ_0 respectively.

Furthermore, since $1 \in L_A(q_4)$ and $q_2 \in \gamma(q_1, \{q_3, q_4\})$ while $1 \in L_A(q_2)$, we have that $\top \in \gamma_0(q_1, \{q_3, q_4\})$ as well as $\top \in \gamma_0(q_1, \{q_3\})$ by the second rule generating γ_0 . Not drawn are the transitions witnessed by $q_2 \in \gamma_0(q_1, \{q_3, q_4\})$ and $q_2 \in \gamma_0(q_1, \{q_3\})$, both of which are consequences of the first rule generating γ_0 , but do not contribute to the language of the automaton. Indeed in the new automaton, no accepting state is present in the support of q_2 , so visiting this state never leads to an accepting run.

Note that even though A was 2-forking, this is no longer the case in A_0 , as a result of the fact that $\top \in \gamma_0(q_1, \{q_3\})$. We will remedy the appearance of unary forks later on.

Our construction is correct, in the following sense.

Lemma 6.11. A_0 is bounded, parsimonious, and weakly implements A .

6.2. Removing nullary forks. Let $A = \langle Q, F, \delta, \gamma \rangle$ be a bounded, parsimonious, and fork-acyclic automaton. We want to implement A with a 1-forking PA A_1 while maintaining boundedness, parsimony and fork-acyclicity. As mentioned, the *nullary* forks that we want to eliminate — i.e., those where $q \in \gamma(p, \square)$ — essentially furnish silent transitions $p \xrightarrow{1}_A q$, analogous to classic NFAs. We proceed eliminate these by means of the classic method, i.e., by saturating the transition relations.

Definition 6.12 (A_1). The PA A_1 is defined to be $\langle Q, F, \delta_1, \gamma_1 \rangle$, where δ_1 and γ_1 are generated by the following inference rules for all $\mathbf{a} \in \Sigma$ and $\phi \in \mathbb{M}(Q)$ with $\phi \neq \square$.

$$\frac{p \xrightarrow{1}_A q \quad r \in \delta(q, \mathbf{a}) \quad r \xrightarrow{1}_A s}{s \in \delta_1(p, \mathbf{a})} \quad \frac{p \xrightarrow{1}_A q \quad r \in \gamma(q, \phi) \quad r \xrightarrow{1}_A s}{s \in \gamma_1(p, \phi)}$$

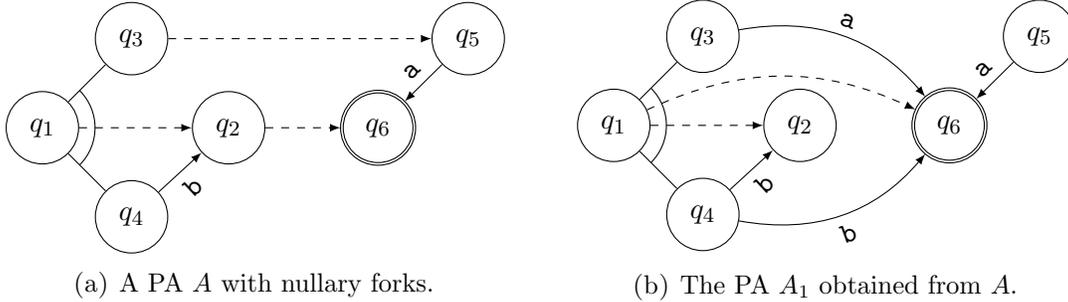


FIGURE 6. Example of nullary fork removal.

Example 6.13. Suppose A is the automaton in Figure 6a. This automaton has two nullary forks: $q_5 \in \gamma(q_3, \emptyset)$ and $q_6 \in \gamma(q_2, \emptyset)$. We have drawn the automaton A_1 obtained from A in Figure 6b. Here, $q_6 \in \delta_1(q_3, a)$, because $q_3 \xrightarrow{1_A} q_5$, $q_6 \in \delta(q_5, a)$ and $q_6 \xrightarrow{1_A} q_6$. Similarly, $q_6 \in \gamma_1(q_1, \{\{q_3, q_4\}\})$ because $q_1 \xrightarrow{1_A} q_1$, $q_2 \in \gamma(q_1, \{\{q_3, q_4\}\})$ and $q_2 \xrightarrow{1_A} q_6$.

We conclude by stating correctness of our translation, in the following sense.

Lemma 6.14. A_1 is bounded, 1-forking, parsimonious, and weakly implements A .

6.3. Removing unary forks. We now show that, given a fork-acyclic, bounded, parsimonious and 1-forking PA $A = \langle Q, F, \delta, \gamma \rangle$, we can implement it using a 2-forking PA A_2 that retains the properties of A . The main idea is to simulate unary forks by keeping a “call stack” in the state. When A follows a unary fork, e.g., $q' \in \gamma(q, \{\{r\}\})$, A_2 will push q' on the stack to “remember” where we should continue after completing the computation in r ; once we reach an accepting state, the transitions of q' will become available. By fork-acyclicity we can bound the height of this stack by the depth of the automaton.

To keep track of this call stack, we need to know which unary forks can occur in sequence from any given state. This is captured by the following.

Definition 6.15. The relation \uparrow , is the least subset of $Q \times Q^*$ satisfying the following rules:

$$\frac{}{q \uparrow q} \qquad \frac{r \uparrow w \quad q' \in \gamma(q, \{\{r\}\})}{q \uparrow wq'}$$

Intuitively, if $q \uparrow q_1 \cdots q_n$, then q can perform a series of unary forks leading to state q_1 . Once the computation starting in q_1 reaches an accepting state, that state will be at the top of the stack; since it is accepting, we can then pop it off the stack to continue trying to resolve q_2 , and so on. The first rule covers the case where no fork takes place, while the second rule allows to extend an existing series of forks with one more.

Example 6.16. Suppose A is the PA in Figure 7a. We first note that $q_6 \uparrow q_6$ by the first rule; hence, since $q_4 \in \gamma(q_3, \{\{q_6\}\})$, we have $q_3 \uparrow q_6q_4$ by applying the second rule. Furthermore, since $q_2 \in \gamma(q_1, \{\{q_3\}\})$, we find that $q_1 \uparrow q_6q_4q_2$, again by the second rule. Hence, q_1 can fork into q_6 , and after completing a computation there and in q_4 , we can carry on in q_2 .

We can now define our transformation, as follows.

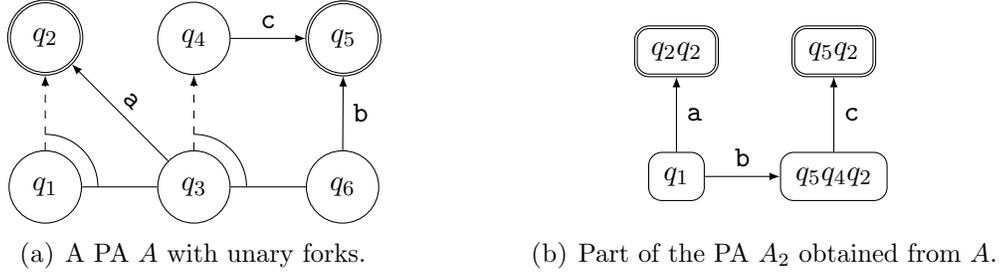


FIGURE 7. An example of unary fork removal.

Definition 6.17. The PA A_2 is defined to be $\langle Q_2, F_2, \delta_2, \gamma_2 \rangle$, where $Q_2 = Q^*$ and $F_2 = F^*$. Also, δ_2 and γ_2 are generated by the following rules for $\mathbf{a} \in \Sigma$ and $\phi \in \mathbb{M}(Q)$ with $|\phi| \geq 2$:

$$\frac{q \uparrow rw \quad q' \in \delta(r, \mathbf{a}) \quad q'wx \in Q^*}{q'wx \in \delta_2(qx, \mathbf{a})} \quad \frac{w' \in \delta_2(w, \mathbf{a}) \quad q \in F \quad qw \in Q^*}{w' \in \delta_2(qw, \mathbf{a})}$$

$$\frac{q \uparrow rw \quad q' \in \gamma(r, \phi) \quad q'wx \in Q^*}{q'wx \in \gamma_2(qx, \phi)} \quad \frac{w' \in \gamma_2(w, \phi) \quad q \in F \quad qw \in Q^*}{w' \in \gamma_2(qw, \phi)}$$

The first rule allows us to look at the state q on top of the stack, and see where it can fork to; if the state r where we end up allows a δ -transition to q' , we push q' onto the stack, along with the unresolved states w gained from the unary fork.⁴ The third rule works analogously, for (non-unary) parallel transitions.

The second rule says that we can also look at states further down the stack, provided that they are preceded only by accepting states. This allows us to pop states off the stack when their computation has reached an accepting state, while continuing in the next state.

Example 6.18. Let A be the PA in Figure 7a. We have drawn the support of q_1 in the automaton A_2 obtained from A in Figure 7b. There, we have $q_2q_2, q_5q_2 \in F_2$ because $q_2, q_5 \in F$. Also, since $q_1 \uparrow q_3q_2$ (see Example 6.16) and $q_2 \in \delta(q_3, \mathbf{a})$, we have $q_2q_2 \in \delta_2(q_1, \mathbf{a})$ by the first rule above. Furthermore, since $q_1 \uparrow q_6q_4q_2$ (see Example 6.16) and $q_5 \in \delta(q_6, \mathbf{b})$, we have $q_5q_4q_2 \in \delta_2(q_1, \mathbf{b})$, again by the first rule above. Lastly, $q_4 \uparrow q_4$, so $q_5q_2 \in \delta_2(q_4q_2, \mathbf{c})$ by the first rule. Since $q_5 \in F$, we find that $q_5q_2 \in \delta_2(q_5q_4q_2, \mathbf{c})$ by the second rule.

The effect of the third and fourth rule is not shown, but is largely analogous.

We again conclude with a statement of correctness of the transformation.

Lemma 6.19. A_2 is bounded, parsimonious, 2-forking, and implements A .

6.4. Flat-branching. In this section, we enforce flat-branching. We start from a PA A that is assumed to be fork-acyclic, bounded, parsimonious, and 2-forking, and we want to construct a flat-branching PA A_3 that weakly implements A , but keeps the properties of A .

The first idea of this construction is fairly obvious: to remove chains of forks while retaining the same language, we will saturate the parallel transitions by unfolding every possible chain of forks. For instance, if $q \in \gamma(p, \{q_1, q_2\})$ and $\gamma(q_1, \{r_1, r_2\}) \cap F \neq \emptyset$, we want to have $q \in \gamma_3(p, \{r_1, r_2, q_2\})$. Fork-acyclicity is instrumental for this construction to

⁴Note that since $q \uparrow q$, we have that $q'x \in \delta_2(qx, \mathbf{a})$ whenever $q' \in \delta(q, \mathbf{a})$.

terminate, as fork cycles could introduce infinitely many $\phi \in \mathbb{M}(Q)$ such that $\gamma(p, \phi) \neq \emptyset$. We make this idea formal as follows:

Definition 6.20. We define \blacktriangleleft as the smallest reflexive relation on $\mathbb{M}(Q)$ satisfying:

$$\frac{\gamma(p, \chi) \cap F \neq \emptyset \quad \phi \blacktriangleleft \psi \sqcup \{p\}}{\phi \blacktriangleleft \psi \sqcup \chi}$$

Intuitively, $\phi \blacktriangleleft \psi$ when a fork into ϕ can be expanded to a fork into ψ , by forking from one or more of its states, provided the new fork reaches an accepting state.

Example 6.21. Recall the left half of the automaton in Figure 3a, depicted in Figure 8a. Since \blacktriangleleft is reflexive, we have $\{q_3\} \blacktriangleleft \{q_3\}$. Then, because $\gamma(q_3, \{q_5, q_6\}) \cap F \neq \emptyset$, we have that $\{q_3\} \blacktriangleleft \{q_5, q_6\}$ applying the rule. Similarly, we have that $\{q_1\} \blacktriangleleft \{q_3, q_4\}$. Combining these two using the rule above then tells us that $\{q_1\} \blacktriangleleft \{q_5, q_6, q_4\}$. Thus, any fork into $\{q_1\}$ to reach some q' can be expanded to a fork into $\{q_5, q_6, q_4\}$ to reach q' .

Next, we want to make sure that the original forks cannot be executed in succession anymore, by forcing all forks to expand maximally before continuing with some non-forking transition. The main idea is to split each state q into q^s and q^p . The state q^s will ensure that $\gamma_3(q^s, \phi) \cap F = \emptyset$ for any multiset ϕ , i.e., no forks are allowed. We leverage this property to get flat-branching, by making sure that for any state $p \in Q_3$ of the new automaton, $\gamma_3(p, \phi) \neq \emptyset$ implies that every state in ϕ is of the q^s variety. On the other hand, from the state q^p , one cannot perform δ -transitions, and furthermore for any multiset ϕ we have $\gamma_3(q^p, \phi) \subseteq \{\top\}$, where \top is the unique accepting state of A_3 (as in Section 6.1).

Definition 6.22 (A_3). The PA A_3 is the quadruple $\langle Q_3, F_3, \delta_3, \gamma_3 \rangle$, where

$$Q_3 = \{q^p : q \in Q\} \cup \{q^s : q \in Q\} \cup \{\top\} \quad F_3 = \{\top\}$$

with \top a fresh state, such that for all $\mathbf{a} \in \Sigma$ and $\phi \in \mathbb{M}(Q_3)$ we have $\delta_3(\top, \mathbf{a}) = \gamma_3(\top, \phi) = \emptyset$. Furthermore, the action of δ_3 and γ_3 on states different from \top is generated by the following rules for all $\mathbf{a} \in \Sigma$ and all $\psi, \phi \in \mathbb{M}(Q)$ with $\psi \blacktriangleleft \phi$:

$$\frac{q \in \delta(p, \mathbf{a})}{q^s, q^p \in \delta_3(p^s, \mathbf{a})} \quad \frac{\delta(p, \mathbf{a}) \cap F \neq \emptyset}{\top \in \delta_3(p^s, \mathbf{a})} \quad \frac{q \in \gamma(p, \psi)}{q^s, q^p \in \gamma_3(p^s, \phi^s)} \quad \frac{\gamma(p, \psi) \cap F \neq \emptyset}{\top \in \gamma_3(p^p, \phi^s)}$$

in which $\phi^s = \{q_1^s, \dots, q_n^s\}$ whenever $\phi = \{q_1, \dots, q_n\}$.

Example 6.23. Let A be the PA in Figure 8a. Part of the support of q_1^p is drawn in Figure 8b. There, we find that, since $q_2 \in \gamma(q_1, \{q_3, q_4\}) \cap F$ and $\{q_3, q_4\} \blacktriangleleft \{q_4, q_5, q_6\}$ (refer to the previous example), also $\top \in \gamma_3(q_1^p, \{q_4^s, q_5^s, q_6^s\})$ by the last rule above. Furthermore, since $q_2 \in \delta(q_5, \mathbf{a}) \cap F$, we have that $\top \in \delta_3(q_5^s, \mathbf{a})$ by the second rule above, but also $q_2^s, q_2^p \in \delta_3(q_5^s, \mathbf{a})$, by the first rule above. Not drawn are the transitions $q_2^s, q_2^p \in \gamma_3(q_1^s, \{q_3^s, q_4^s\})$ which would result from applying the third rule. These transitions do not contribute to the language, since they lead to states from which no accepting run is possible.

We conclude this transformation by stating the desired correctness.

Lemma 6.24. A_3 is bounded, 2-forking, parsimonious, flat-branching (i.e., well-structured) and weakly implements A .

In total, we have proved the following theorem.

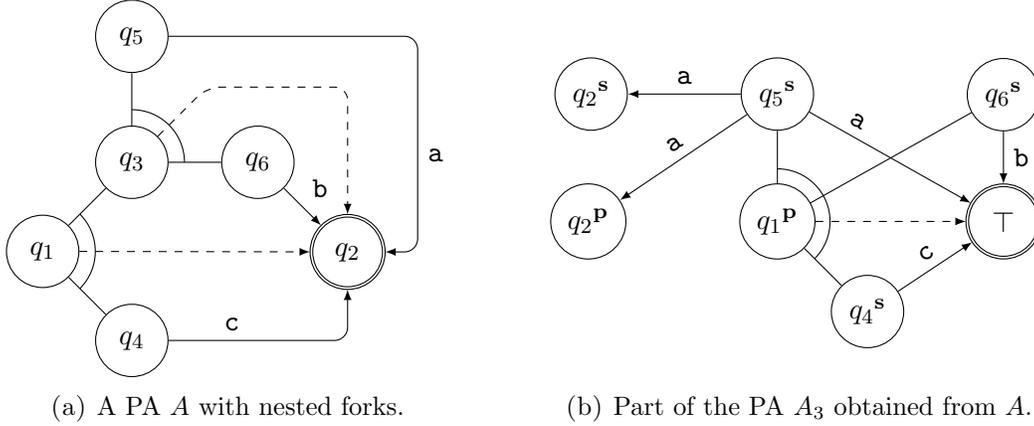


FIGURE 8. Example of construction to ensure flat-branching.

Theorem 6.25. *Let A be a finite and fork-acyclic PA. We can construct a finite and fork-acyclic PA A' that is well-structured and implements A .*

The complexity of the construction above is hard to pin down, because the number of states supporting any single state in the well-nested automaton is obscured by the infinite nature of the intermediate automata. Nevertheless, we anticipate that the removal of unary forks can be expected to provide the worst blowup in automaton size (possibly exponential in D_A), as the number of states grows to accommodate all possible “call stacks”. The removal of nullary forks and the flat-branching construction both contribute a relatively small number of new states; the same is true for parsimonification, although this step may add a significant number of parallel transitions.

7. FROM EXPRESSIONS TO AUTOMATA

Recall that sp-languages can be defined denotationally, in terms of a sr-expression $e \in \mathcal{T}$. In this section, we show how to obtain a fork-acyclic and finite PA that accepts the semantics of e , starting at some state. Concretely, we define a bounded pomset automaton where *every* sr-expression is a state, and the language of this state is intended to be the semantics of said sr-expression. The transition functions are set up such that, starting from $e \in \mathcal{T}$, reading a particular pomset brings us to $e' \in \mathcal{T}$, which describes the sr-language that remains to be read. This generalizes Antimirov’s (partial) derivatives of rational expressions [Ant96]; we therefore refer to the transition functions on expressions as *derivatives*.

Before we define the derivatives, it is convenient to introduce two shorthands. The first of these allows us to include a set of terms based on whether or not another term is in \mathcal{F} , the set of sr-expressions that accept the empty pomset (c.f. Definition 3.15). The second shorthand allows us to concatenate a set of terms with another term on the right.

Definition 7.1 (Derivatives). We define $\delta_\Sigma : \mathcal{T} \times \Sigma \rightarrow 2^\mathcal{T}$ and $\gamma_\Sigma : \mathcal{T} \times \mathbb{M}(\mathcal{T}) \rightarrow 2^\mathcal{T}$ inductively. To this end, the following shorthands are useful, for all $e \in \mathcal{T}$ and $T \subseteq \mathcal{T}$:

$$e \star T = \begin{cases} T & e \in \mathcal{F} \\ \emptyset & \text{otherwise} \end{cases} \quad T \mathbin{\text{;}} e = \{f \cdot e : f \in T\}$$

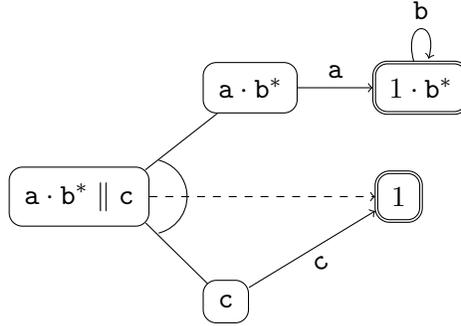


FIGURE 9. Part of the syntactic pomset automaton.

This then allows us to define δ_Σ and γ_Σ inductively, as follows:

$$\begin{array}{ll}
 \delta_\Sigma(0, \mathbf{a}) = \emptyset & \gamma_\Sigma(0, \phi) = \emptyset \\
 \delta_\Sigma(1, \mathbf{a}) = \emptyset & \gamma_\Sigma(1, \phi) = \emptyset \\
 \delta_\Sigma(\mathbf{b}, \mathbf{a}) = \{1 : \mathbf{a} = \mathbf{b}\} & \gamma_\Sigma(\mathbf{b}, \phi) = \emptyset \\
 \delta_\Sigma(e + f, \mathbf{a}) = \delta_\Sigma(e, \mathbf{a}) \cup \delta_\Sigma(f, \mathbf{a}) & \gamma_\Sigma(e + f, \phi) = \gamma_\Sigma(e, \phi) \cup \gamma_\Sigma(f, \phi) \\
 \delta_\Sigma(e \cdot f, \mathbf{a}) = \delta_\Sigma(e, \mathbf{a}) \ ; \ f \cup e \star \delta_\Sigma(f, \mathbf{a}) & \gamma_\Sigma(e \cdot f, \phi) = \gamma_\Sigma(e, \phi) \ ; \ f \cup e \star \gamma_\Sigma(f, \phi) \\
 \delta_\Sigma(e \parallel f, \mathbf{a}) = \emptyset & \gamma_\Sigma(e \parallel f, \phi) = \{1 : \phi = \{e, f\}\} \\
 \delta_\Sigma(e^*, \mathbf{a}) = \delta_\Sigma(e, \mathbf{a}) \ ; \ e^* & \gamma_\Sigma(e^*, \phi) = \gamma_\Sigma(e, \phi) \ ; \ e^*
 \end{array}$$

We can now define our pomset automaton that operates on terms.

Definition 7.2. The *syntactic PA*, denoted by A_Σ , is the PA $\langle \mathcal{T}, \mathcal{F}, \delta_\Sigma, \gamma_\Sigma \rangle$.

We simplify subscripts, writing \rightarrow_Σ instead of \rightarrow_{A_Σ} , and so forth.

Example 7.3. We have drawn part of the syntactic PA, specifically the support of $\mathbf{a} \cdot \mathbf{b}^* \parallel \mathbf{c}$, in Figure 9. There, we see that $1 \cdot \mathbf{b}^*$ is an accepting state, because $1, \mathbf{b}^* \in \mathcal{F}$. The sequential transitions are generated by δ_Σ ; for instance, $1 \cdot \mathbf{b}^* \in \delta_\Sigma(\mathbf{a} \cdot \mathbf{b}^*, \mathbf{a})$, because $1 \in \delta_\Sigma(\mathbf{a}, \mathbf{a})$ and $\delta_\Sigma(\mathbf{a}, \mathbf{a}) \ ; \ \mathbf{b}^* \subseteq \delta_\Sigma(\mathbf{a} \cdot \mathbf{b}^*, \mathbf{a})$; also, $1 \cdot \mathbf{b}^* \in \delta_\Sigma(1 \cdot \mathbf{b}^*, \mathbf{b})$, because $1 \in \delta_\Sigma(\mathbf{b}, \mathbf{b})$, and $\delta_\Sigma(\mathbf{b}, \mathbf{b}) \ ; \ \mathbf{b}^* \subseteq \delta_\Sigma(1 \cdot \mathbf{b}^*, \mathbf{b})$. Lastly, $1 \in \gamma_\Sigma(\mathbf{a} \cdot \mathbf{b}^* \parallel \mathbf{c}, \{\mathbf{a} \cdot \mathbf{b}^*, \mathbf{c}\})$ by definition of γ_Σ .

To show that the language of a state e in the syntactic PA is exactly $\llbracket e \rrbracket$, we first need to show a number of intermediary lemmas that analyze runs that can occur.

7.1. Deconstruction lemmas. The first set of lemmas are called *deconstruction lemmas*: they tell us how, given a run originating in e , we can obtain one or more runs of terms originating in subterms of e . First, we show how to do this for sums.

Lemma 7.4. *Let $e_1, e_2 \in \mathcal{T}$, $f \in \mathcal{F}$ and $U \in \text{SP}(\Sigma)$, such that $e_1 + e_2 \xrightarrow{U}_\Sigma f$. There exists an $f' \in \mathcal{F}$ with $e_1 \xrightarrow{U}_\Sigma f'$ or $e_2 \xrightarrow{U}_\Sigma f'$.*

Proof. Let ℓ be the length of $e_1 + e_2 \xrightarrow{U}_\Sigma f$. There are two cases to consider.

- If $\ell = 0$, then $e_1 + e_2 \xrightarrow{U}_\Sigma f$ is trivial. In that case, $f = e_1 + e_2$, and so $e_i \in \mathcal{F}$ for some $i \in \{1, 2\}$; if we choose $f' = e_i$, we find $e_i \xrightarrow{U}_\Sigma f'$, and the claim is satisfied.

- Otherwise, if $\ell > 0$, then $U = U_0 \cdot U'$ and there exists a $g \in \mathcal{T}$ such that $e_1 + e_2 \xrightarrow{U_0}_\Sigma g$ is a unit run, and $g \xrightarrow{U'}_\Sigma f$ is of length $\ell - 1$. If $e_1 + e_2 \xrightarrow{U_0}_\Sigma g$ is a sequential unit run, then $U_0 = \mathbf{a}$ for some $\mathbf{a} \in \Sigma$, and $g \in \delta_\Sigma(e_1 + e_2, \mathbf{a})$. Without loss of generality, let $g \in \delta_\Sigma(e_1, \mathbf{a})$; in that case, $e_1 \xrightarrow{U_0}_\Sigma g$. If we then choose $f' = f$, we find that $e_1 \xrightarrow{U}_\Sigma f'$. The case where $e_1 + e_2 \xrightarrow{U_0}_\Sigma g$ is a parallel unit run is similar. \square

Next, we show how a run that starts in a sequential composition gives rise to runs that originate in each of the components. In this case, it is useful to also relate the length of the original run to the lengths of the deconstructed runs, so that we can use this result in the analogous deconstruction result for the Kleene star (Lemma 7.7).

Lemma 7.5. *Let $e_1, e_2 \in \mathcal{T}$, $f \in \mathcal{F}$ and $U \in \text{SP}(\Sigma)$, with $e_1 \cdot e_2 \xrightarrow{U}_\Sigma f$ (of length ℓ). Then $U = U_1 \cdot U_2$ and there exist $f_1, f_2 \in \mathcal{F}$ with $e_i \xrightarrow{U_i}_\Sigma f_i$ for $i \in \{1, 2\}$ (of length at most ℓ).*

There is also a deconstruction lemma for runs originating in a parallel composition, where the pomset accepted is a parallel composition of pomsets accepted by the operands.

Lemma 7.6. *Let $e_1, e_2 \in \mathcal{T}$, $f \in \mathcal{F}$ and $U \in \text{SP}(\Sigma)$ such that $e_1 \parallel e_2 \xrightarrow{U}_\Sigma f$. Then $U = U_1 \parallel U_2$, and there exist $f_1, f_2 \in \mathcal{F}$ such that $e_i \xrightarrow{U_i}_\Sigma f_i$ for $i \in \{1, 2\}$.*

For the last deconstruction lemma, we consider terms with the Kleene star as the topmost operator. Here, we show how to obtain a number of runs, each of which originates in the expression directly below the Kleene star.

Lemma 7.7. *Let $e \in \mathcal{T}$, $f \in \mathcal{F}$ and $U \in \text{SP}(\Sigma)$ be such that $e^* \xrightarrow{U}_\Sigma f$. There exist $f_1, f_2, \dots, f_n \in \mathcal{F}$ such that $U = U_1 \cdot U_2 \cdots U_n$ and for $1 \leq i \leq n$ it holds that $e \xrightarrow{U_i}_\Sigma f_i$.*

Proof. The proof proceeds by induction on the length ℓ of $e^* \xrightarrow{U}_\Sigma f$. In the base, where $\ell = 0$, we have that $f = e^*$ and $U = 1$; it suffices to choose $n = 0$.

In the inductive step, let $e^* \xrightarrow{U}_\Sigma f$ be of length $\ell + 1$. We find $g \in \mathcal{T}$ and $U = U_0 \cdot U'$ such that $e^* \xrightarrow{U_0}_\Sigma g$ is a unit run, and $g \xrightarrow{U'}_\Sigma f$ is of length ℓ . If $e^* \xrightarrow{U_0}_\Sigma g$ is a sequential unit run, then $U_0 = \mathbf{a}$ for some $\mathbf{a} \in \Sigma$, and $g \in \delta_\Sigma(e^*, \mathbf{a}) = \delta_\Sigma(e, \mathbf{a}) \circledast e^*$. Hence, $g = g' \cdot e^*$, with $g' \in \delta_\Sigma(e, \mathbf{a})$. By Lemma 7.5, we find $f'', f' \in \mathcal{F}$ such that $U' = V \cdot W$ as well as $g' \xrightarrow{V}_\Sigma f''$ and $e^* \xrightarrow{W}_\Sigma f'$, of length at most ℓ . By induction, we find $f_2, f_3, \dots, f_n \in \mathcal{F}$ such that $W = U_2 \cdot U_3 \cdots U_n$, and for $1 < i \leq n$ we have $e \xrightarrow{U_i}_\Sigma f_i$.

We then choose $f_1 = f''$ and $U_1 = U_0 \cdot V$. For these choices, $U = U_0 \cdot U' = U_0' \cdot V \cdot W = U_1 \cdot U_2 \cdots U_n$. Since $e \xrightarrow{U_0}_\Sigma \delta_\Sigma(e, \mathbf{a}) \xrightarrow{V}_\Sigma f''$, we also have $e \xrightarrow{U_1}_\Sigma f_1$, satisfying the claim.

The case where $e^* \xrightarrow{U_0}_\Sigma g$ is a parallel unit run is similar. \square

7.2. Construction lemmas. Next, we prove the *construction lemmas*, so called because they tell us how to combine runs that originate in some terms into runs that originate in some composition of those terms. Keeping the same order of operators as before, we start by showing how to construct runs originating from terms that are sums.

Lemma 7.8. *Let $e_1, e_2 \in \mathcal{T}$, $f_1, f_2 \in \mathcal{F}$ and $U \in \text{SP}(\Sigma)$ be such that $e_1 \xrightarrow{U}_\Sigma f_1$ or $e_2 \xrightarrow{U}_\Sigma f_2$. There exists an $f \in \mathcal{F}$ such that $e_1 + e_2 \xrightarrow{U}_\Sigma f$.*

Proof. We treat the case where $e_1 \xrightarrow{U}_\Sigma f_1$; the case where $e_2 \xrightarrow{U}_\Sigma f_2$ can be treated similarly. There are two possibilities to consider, based on the length ℓ of $e_1 \xrightarrow{U}_\Sigma f_1$.

- If $\ell = 0$, then $f_1 = e_1$ and $U = 1$. We then choose $f = e_1 + e_2$.
- If $\ell > 0$, then we find $e'_1 \in \mathcal{T}$ and $U = U_0 \cdot U'$ such that $e_1 \xrightarrow{U_0}_{\Sigma} e'_1$ is a unit run, and $e'_1 \xrightarrow{U'}_{\Sigma} f_1$. If $e_1 \xrightarrow{U_0}_{\Sigma} e'_1$ is a sequential unit run, then $U_0 = \mathbf{a}$ for some $\mathbf{a} \in \Sigma$, and $e'_1 \in \delta_{\Sigma}(e_1, \mathbf{a})$. But then $e'_1 \in \delta_{\Sigma}(e_1 + e_2, \mathbf{a})$ as well, and hence $e_1 + e_2 \xrightarrow{U_0}_{\Sigma} e'_1$. Putting this together, we find that $e_1 + e_2 \xrightarrow{U}_{\Sigma} f_1$; choosing $f = f_1$ then satisfies the claim.

The case where $e_1 \xrightarrow{U_0}_{\Sigma} e'_1$ is a parallel unit run is similar. \square

The construction lemma for sequential composition says that we can compose runs to accepting states in the following way.

Lemma 7.9. *Let $e_1, e_2 \in \mathcal{T}$, $f_1, f_2 \in \mathcal{F}$ and $U, V \in \text{SP}(\Sigma)$ such that $e_1 \xrightarrow{U}_{\Sigma} f_1$ and $e_2 \xrightarrow{V}_{\Sigma} f_2$. There exists an $f \in \mathcal{F}$ with $e_1 \cdot e_2 \xrightarrow{U \cdot V}_{\Sigma} f$.*

The next operator is parallel composition; here, the construction lemma is easy.

Lemma 7.10. *Let $e_1, e_2 \in \mathcal{T}$, $f_1, f_2 \in \mathcal{F}$, and $U, V \in \text{SP}(\Sigma)$ such that $e_1 \xrightarrow{U}_{\Sigma} f_1$ and $e_2 \xrightarrow{V}_{\Sigma} f_2$. Then $e_1 \parallel e_2 \xrightarrow{U \parallel V}_{\Sigma} 1$.*

The last construction lemma concerns the Kleene star. Here, a succession of runs originating in the expression below the star allows us to create a run originating in the starred expression, provided each of these runs reaches an accepting state.

Lemma 7.11. *Let $e \in \mathcal{T}$, $f_1, f_2, \dots, f_n \in \mathcal{F}$, and $U_1, U_2, \dots, U_n \in \text{SP}(\Sigma)$ be such that for $1 \leq i \leq n$ it holds that $e \xrightarrow{U_i}_{\Sigma} f_i$. There exists an $f \in \mathcal{F}$ such that $e^* \xrightarrow{U_1 \cdot U_2 \cdots U_n}_{\Sigma} f$.*

Proof. We can assume w.l.o.g. that for $0 \leq i < n$ it holds that $e \xrightarrow{U_i}_{\Sigma} f_i$ is non-trivial. We proceed by induction on n . In the base, where $n = 0$, we can choose $f = e^*$.

For the inductive step, let $n > 0$ and assume that the claim holds for $n - 1$. By induction, we find $f' \in \mathcal{F}$ with $e^* \xrightarrow{U_2 \cdot U_3 \cdots U_n}_{\Sigma} f'$. Since $e \xrightarrow{U_1}_{\Sigma} f_1$ is non-trivial, we find $e' \in \mathcal{T}$ and $U_1 = U_0 \cdot U'_1$ such that $e \xrightarrow{U_0}_{\Sigma} e'$ is a unit run, and $e' \xrightarrow{U'_1}_{\Sigma} f_1$. By Lemma 7.9, we find $f \in \mathcal{F}$ such that $e' \cdot e^* \xrightarrow{U'_1 \cdot U_2 \cdot U_3 \cdots U_n}_{\Sigma} f$. If $e \xrightarrow{U_0}_{\Sigma} e'$ is a sequential unit run, then $U_0 = \mathbf{a}$ for some $\mathbf{a} \in \Sigma$, and $e' \in \delta_{\Sigma}(e, \mathbf{a})$. But then $e' \cdot e^* \in \delta_{\Sigma}(e^*, \mathbf{a})$, and hence $e^* \xrightarrow{U_0}_{\Sigma} e' \cdot e^*$. Thus, $e^* \xrightarrow{U_0}_{\Sigma} e' \cdot e^* \xrightarrow{U'_1 \cdot U_2 \cdot U_3 \cdots U_n}_{\Sigma} f$ and therefore $e^* \xrightarrow{U_1 \cdot U_2 \cdots U_n}_{\Sigma} f$.

The case where $e \xrightarrow{U_0}_{\Sigma} e'$ is a parallel unit run is similar. \square

7.3. Correctness of the syntactic PA. The (de)construction lemmas combine to prove the equations claimed by the lemma below; the deconstruction lemmas show inclusion from left to right, whereas construction lemmas show the inclusion from right to left.

Lemma 7.12. *Let $e, f \in \mathcal{T}$. The following hold:*

$$\begin{array}{lll} L_{\Sigma}(0) = \emptyset & L_{\Sigma}(e + f) = L_{\Sigma}(e) \cup L_{\Sigma}(f) & L_{\Sigma}(e^*) = L_{\Sigma}(e)^* \\ L_{\Sigma}(1) = \{1\} & L_{\Sigma}(e \cdot f) = L_{\Sigma}(e) \cdot L_{\Sigma}(f) & \\ L_{\Sigma}(\mathbf{a}) = \{\mathbf{a}\} & L_{\Sigma}(e \parallel f) = L_{\Sigma}(e) \parallel L_{\Sigma}(f) & \end{array}$$

A straightforward inductive argument now helps us validate the following:

Lemma 7.13. *For all $e \in \mathcal{T}$, we have $L_{\Sigma}(e) = \llbracket e \rrbracket$.*

This covers the correctness of our expressions-to-automata construction as far as the languages are concerned. It remains to show that the automaton is also fork-acyclic, which comes down to showing that if $e \preceq_{\Sigma} f$, then the nesting depth of \parallel in e is bounded by that of f . Since forks only occur to expressions with strictly lower parallel nesting depth, neither e nor f can support $e \parallel f$.

Lemma 7.14. *The syntactic pomset automaton is fork-acyclic.*

Proof sketch. We construct a function $\text{depth}(-) : \mathcal{T} \rightarrow \mathbb{N}$, and we show that if $e \preceq_{\Sigma} f$, then $\text{depth}(e) \leq \text{depth}(f)$. In the process, we also show that if $e, f \in \mathcal{T}$ and $\phi \in \mathbb{M}(\mathcal{T})$ with $\gamma(e, \phi) \neq \emptyset$ and $f \in \phi$, then $\text{depth}(f) < \text{depth}(e)$; together with the earlier property of $\text{depth}(-)$, this tells us that $f \prec_{\Sigma} e$, as desired. \square

Furthermore, each state in the syntactic PA has finite support. This is done by constructing a (finite) overestimation of the expressions that support e in A_{Σ} .

Lemma 7.15. *The syntactic pomset automaton is bounded.*

Proof sketch. We should show that for $e \in \mathcal{T}$, it holds that $\pi_{\Sigma}(e)$ is finite. Since $\pi_{\Sigma}(e)$ is the smallest closed set that contains e , it suffices to find a finite closed set $S(e)$ such that $e \in S(e)$. The construction of this set and the verification of its closure is entirely routine; we refer to the appendix for details. \square

This then allows us to state one half of our Kleene theorem, as follows.

Theorem 7.16 (Expressions to automata). *Let $e \in \mathcal{T}$. We can obtain a fork-acyclic and finite PA A with a state q such that $L_A(q) = \llbracket e \rrbracket$.*

Proof. Choose $A = A_{\Sigma}[e]$. By Lemmas 7.14 and 7.15 as well as Lemma 4.17, A is finite and fork-acyclic. Finally, by Lemmas 7.13 and 4.17, we have that $L_{A[\pi_{\Sigma}(e)]}(e) = L_{\Sigma}(e) = \llbracket e \rrbracket$. \square

It also shows that equivalence of sr-expressions is decidable.

Corollary 7.17. *Given $e, f \in \mathcal{T}$, it is decidable whether $\llbracket e \rrbracket = \llbracket f \rrbracket$.*

8. FROM AUTOMATA TO EXPRESSIONS

Let $A = \langle Q, F, \delta, \gamma \rangle$ be a finite and fork-acyclic pomset automaton. In this section, we set out to obtain the converse of the result in the previous section, that is, to construct for every state $q \in Q$ of the automaton A a series-rational expression e_q such that $\llbracket e_q \rrbracket = L_A(q)$. Since A is finite and fork-acyclic, the strict support relation \prec_A is well-founded. We can therefore proceed building e_q by \prec_A -induction. Our *main induction hypothesis* is as follows:

If $r \in Q$ with $r \prec_A q$, then there exists an sr-expression e_r with $\llbracket e_r \rrbracket = L_A(r)$.

Our task is to define for q a series-rational expression e_q such that $\llbracket e_q \rrbracket = L_A(q)$. As a matter of fact, it is convenient to establish a slightly more general result: we let

$$Q' = \{p \in Q : q \preceq_A p \preceq_A q\}$$

and simultaneously define for every $p \in Q'$ an sr-expression e_p such that $\llbracket e_p \rrbracket = L_A(p)$.

To define these expressions, we use an approach due to McNaughton and Yamada [MY60]. The first step in this approach is to define, for all $p \in Q'$ and $r \in Q$, an sr-expression e_{pr}

that captures the pomsets that can be read when transitioning from p to r by means of a (sequential or parallel) unit run. This can be done as follows:

$$e_{pr} = \sum_{r \in \delta(p, \mathbf{a})} \mathbf{a} + \sum_{r \in \gamma(p, \{\!\{s_1, \dots, s_n\}\!\})} e_{s_1} \parallel \dots \parallel e_{s_n}$$

Intuitively, the first sum accounts for sequential unit runs: if $r \in \delta(p, \mathbf{a})$, then $p \xrightarrow{\mathbf{a}}_A r$ is a unit run. For the second term, we note that if $s_1, \dots, s_n \in Q$ and $r \in \gamma(p, \{\!\{s_1, \dots, s_n\}\!\})$, then $s_1, \dots, s_n \prec_A p$ by fork-acyclicity; therefore, by the main induction hypothesis, e_{s_i} is already defined such that $L_A(s_i) = \llbracket e_{s_i} \rrbracket$. Thus, if for $1 \leq i \leq n$, we have $U_i \in \llbracket e_{s_i} \rrbracket$, then the pomset U_i is accepted by s_i , which means that $p \xrightarrow{U_1 \parallel \dots \parallel U_n}_A r$ is a parallel unit run.

The following asserts that e_{pr} is fit for purpose:

Lemma 8.1. *Let $p \in Q'$ and $r \in Q$. Now $U \in \llbracket e_{pr} \rrbracket$ if and only if $p \xrightarrow{U}_A r$ is a unit run.*

The next step is to extend the expression e_{pr} into an expression that describes the labels of (non-trivial) runs that may pass through intermediate states in Q' . More precisely, for $S \subseteq Q'$, we choose e_{pr}^S to describe the language of pomsets that can be read while transitioning from p to r , while passing through states in S as intermediates.⁵

$$e_{pr}^S = \begin{cases} e_{pr} & S = \emptyset \\ e_{pr}^{S'} + e_{ps}^{S'} \cdot (e_{ss}^{S'})^* \cdot e_{sr}^{S'} & S = S' \cup \{s\}, s \notin S' \end{cases}$$

Intuitively, when $S = \emptyset$, no intermediate states are allowed, so we need a unit run from p to r . When S is non-empty, we single out the state s ; a run from p to r either does not pass through s (the first term), or it passes through s at least once, where it may loop several times, before continuing on to r (the second term).

The following asserts the correctness of e_{pr}^S w.r.t. our intention:

Lemma 8.2. *Let $p \in Q'$, $r \in Q$, and $S \subseteq Q'$. Now it holds that $U \in \llbracket e_{pr}^S \rrbracket$ if and only if there exist $\ell \geq 1$ and $s_1, \dots, s_{\ell-1} \in S$ and $U = U_1 \cdots U_n$, such that*

$$p = s_0 \xrightarrow{U_1}_A s_1 \xrightarrow{U_2}_A s_2 \xrightarrow{U_3}_A \dots \xrightarrow{U_{\ell-1}}_A s_{\ell-1} \xrightarrow{U_\ell}_A s_\ell = r$$

in which each of the above is a unit run.

Proof. We proceed by induction on S . The case where $S = \emptyset$, is covered by Lemma 8.1.

In the inductive step, let $S = S' \cup \{s\}$, with $s \notin S'$. For the implication from left to right, suppose $U \in \llbracket e_{pr}^S \rrbracket$. Then, there are two cases:

- If $U \in \llbracket e_{pr}^{S'} \rrbracket$, then the claim follows immediately by induction.
- If $U \in \llbracket e_{ps}^{S'} \cdot (e_{ss}^{S'})^* \cdot e_{sr}^{S'} \rrbracket$, then, in accordance with the semantics of sequential composition and Kleene star, there exist $U_0 \in \llbracket e_{ps}^{S'} \rrbracket$, $U_1, \dots, U_n \in \llbracket e_{ss}^{S'} \rrbracket$ and $U_{n+1} \in \llbracket e_{sr}^{S'} \rrbracket$ such that $U = U_0 \cdot U_1 \cdots U_n \cdot U_{n+1}$. With $n + 2$ applications of the induction hypothesis we then we get an appropriate sequence of unit runs.

For the converse implication, we can identify all occurrences of s among the $s_1, \dots, s_{\ell-1}$, and partition the sequence into subsequences that do not travel through s as an intermediate state. We can then apply the induction hypothesis to find that $U \in \llbracket e_{ps}^{S'} \cdot (e_{ss}^{S'})^* \cdot e_{sr}^{S'} \rrbracket$. \square

⁵The inductive step is ambiguous, in that we do not specify *which* $s \in S$ to choose — any choice will do.

The final step is to create an expression e_p that characterises $L_A(p)$, i.e., the pomsets U such that $p \xrightarrow{U}_A r$ for some $r \in F$. Note that, if this run has exited Q' to reach s , then $p \prec_A s$; in particular, this means that once a run has left Q' , it cannot pass through (or fork into) states in Q' anymore. Therefore, $p \xrightarrow{U}_A r$ only passes through states in Q' , or $p \xrightarrow{U}_A r$ eventually leaves Q' , never to return again.

This guides us to define e_p as follows:

$$e_p = \sum_{r \in Q' \cap F} e_{pr}^{Q'} + \sum_{r \prec_A p} e_{pr}^{Q'} \cdot e_r + \begin{cases} 1 & p \in F \\ 0 & \text{otherwise} \end{cases}$$

Intuitively, the first term accounts for the possibility that the run from p to some accepting state stays within Q' ; the second term incorporates that we can reach some $r \prec_A p$ (for which we already have the expression e_r , by the main induction hypothesis), whence we reach an accepting state. The last term is there to allow for p to accept immediately.

We can then prove that e_p indeed describes the language of p .

Lemma 8.3. *For all $p \in Q'$ it holds that $\llbracket e_p \rrbracket = L_A(p)$.*

Proof. For the inclusion from left to right, suppose that $U \in \llbracket e_p \rrbracket$; we distinguish three cases:

- If $U \in \llbracket e_{pr}^{Q'} \rrbracket$ for some $r \in Q' \cap F$, then by Lemma 8.2 we have that $p \xrightarrow{U}_A r$ and hence, since $r \in F$, we find that $U \in L_A(p)$.
- If $U \in \llbracket e_{pr}^{Q'} \cdot e_r \rrbracket$ for some $r \prec_A p$, then there exist pomsets V and W such that $U = V \cdot W$ with $V \in \llbracket e_{pr}^{Q'} \rrbracket$ and $W \in \llbracket e_r \rrbracket$. From $V \in \llbracket e_{pr}^{Q'} \rrbracket$ it follows by Lemma 8.2 that $p \xrightarrow{V}_A r$; from $W \in \llbracket e_r \rrbracket$ it follows by the main induction hypothesis that $W \in L_A(r)$, and hence $r \xrightarrow{W}_A r'$ for some $r' \in F$. In total, we have that $p \xrightarrow{V \cdot W}_A r'$, and therefore $U = V \cdot W \in L_A(p)$.
- If $U = 1$ and $p \in F$, then $p \xrightarrow{U}_A p \in F$, and hence $U \in L_A(p)$ immediately.

For the other inclusion, suppose that $U \in L_A(p)$. Then there exists $r \in F$ such that $p \xrightarrow{U}_A r$, and hence, by Lemma 4.6, there exist states $p = s_0, \dots, s_n = r$ and pomsets U_1, \dots, U_n such that $U = U_1 \cdots U_n$ and $s_{i-1} \xrightarrow{U_i}_A s_i$ is a unit run for all $1 \leq i \leq n$. If $n = 0$, then $p \xrightarrow{U}_A r$ is trivial, meaning $p \in F$ and $U = 1$, and hence $U \in \llbracket e_p \rrbracket$ immediately. If $n \geq 1$ and $s_i \in Q'$ for all $0 \leq i < n$, then by Lemma 8.2, we have $U \in \llbracket e_{pr}^{Q'} \rrbracket \subseteq \llbracket e_p \rrbracket$. Otherwise, let s_i be the first intermediate state such that $s_i \notin Q'$. Since $s_i \preceq_A p$, we have that $s_i \prec_A p$, so from $U_i \cdots U_n \in L_A(s_i)$ it follows by the main induction hypothesis that $U_i \cdots U_n \in \llbracket e_{s_i} \rrbracket$. Moreover, since $s_j \in Q'$ for all $0 \leq j < i$, we have by Lemma 8.2 that $U_1 \cdots U_{i-1} \in \llbracket e_{ps_i}^k \rrbracket$. We conclude that $U_1 \cdots U_n \in \llbracket e_{ps_i}^{Q'} \cdot e_{s_i} \rrbracket \subseteq \llbracket e_p \rrbracket$. \square

By induction on strict support, we have now proved the following theorem.

Theorem 8.4 (Automata to expressions). *Let $A = \langle Q, F, \delta, \gamma \rangle$ be a finite and fork-acyclic PA. For every $q \in Q$, we can construct an $e_q \in \mathcal{T}$ such that $\llbracket e_q \rrbracket = L_A(q)$.*

Theorems 7.16 and 8.4 now combine to form our Kleene theorem.

Theorem 8.5 (Kleene theorem). *Let L be a pomset language. The following are equivalent:*

- (1) L is series-rational;
- (2) L is accepted by a state of a finite and fork-acyclic PA.

9. FURTHER WORK

The decision procedure for language equivalence in pomset automata as outlined in Section 5 is not very efficient. For one, the algorithm computes the atoms generated by the languages of *all* states, rather than just the ones being compared. Consequently, to use equivalence checking of pomset automata as a proxy for equivalence checking of sr-expressions in an applied setting, a more efficient procedure is necessary. Bisimulation up-to equivalence [HK71] and extensions thereof would be an obvious class of optimizations to explore.

Another direction for future work is to transfer the techniques from [BPS17] to fork-acyclic pomset automata. The advantage of such an approach is that it would most likely not require the input PA to be well-structured. Furthermore, since the techniques from [BPS17] can also be used to compare safe Petri nets “up-to sequentialisation”, an adaptation of this algorithm could help us decide an analogous property for pomset automata.

In [KBL⁺19], we presented a Kleene theorem that includes the parallel variant of the Kleene star, sometimes known as *parallel star* or *replication*. This version of the theorem covers a strictly larger class of pomset automata, including those that are not fork-acyclic. However, the structural condition that describes this larger class is not overly complicated and we would like to extend the decision procedure of the present paper to work for this class as well. Doing so would distinguish our algorithm from [BPS17], which does not (and cannot) include this operator, as it would necessarily construct Petri nets that are unbounded.

Acknowledgements. We would like to thank the anonymous reviewers for their insightful comments, which helped to substantially improve the initial version of this paper.

REFERENCES

- [AMR07] Marco Almeida, Nelma Moreira, and Rogério Reis. On the performance of automata minimization algorithms. Technical Report DCC-2007-03, Universidade do Porto, June 2007.
- [Ant96] Valentin M. Antimirov. Partial derivatives of regular expressions and finite automaton constructions. *Theor. Comput. Sci.*, 155(2):291–319, 1996.
- [BLMvT16] Jos C. M. Baeten, Bas Luttik, Tim Muller, and Paul van Tilburg. Expressiveness modulo bisimilarity of regular expressions with parallel composition. *Math. Struct. Comput. Sci.*, 26(6):933–968, 2016.
- [BPS17] Paul Brunet, Damien Pous, and Georg Struth. On decidability of concurrent Kleene algebra. In *Proc. Concurrency Theory (CONCUR)*, pages 28:1–28:15, 2017.
- [Brz64] Janusz A. Brzozowski. Derivatives of regular expressions. *J. ACM*, 11(4):481–494, 1964.
- [BT14] Janusz A. Brzozowski and Hellis Tamm. Theory of átomata. *Theor. Comput. Sci.*, 539:13–27, 2014.
- [Gis88] Jay L. Gischer. The equational theory of pomsets. *Theor. Comput. Sci.*, 61:199–224, 1988.
- [GKM72] Ronald L. Graham, Donald E. Knuth, and T. S. Motzkin. Complements and transitive closures. *Discret. Math.*, 2(1):17–29, 1972.
- [Gra81] Jan Grabowski. On partial languages. *Fundam. Inform.*, 4(2):427, 1981.
- [HK71] John E. Hopcroft and Richard M. Karp. A linear algorithm for testing equivalence of finite automata. Technical Report TR71-114, Cornell University, December 1971.
- [HMSW09] Tony Hoare, Bernhard Möller, Georg Struth, and Ian Wehrman. Concurrent Kleene Algebra. In *Proc. Concurrency Theory (CONCUR)*, pages 399–414, 2009.
- [JM16] Peter Jipsen and M. Andrew Moshier. Concurrent Kleene Algebra with tests and branching automata. *J. Log. Algebr. Meth. Program.*, 85(4):637–652, 2016.
- [KBL⁺17] Tobias Kappé, Paul Brunet, Bas Luttik, Alexandra Silva, and Fabio Zanasi. Brzozowski goes concurrent — a Kleene theorem for pomset languages. In *Proc. Concurrency Theory (CONCUR)*, pages 25:1–25:16, 2017.

- [KBL⁺19] Tobias Kappé, Paul Brunet, Bas Luttik, Alexandra Silva, and Fabio Zanasi. On series-parallel pomset languages: Rationality, context-freeness and automata. *J. Log. Algebr. Meth. Program.*, 103:130–153, 2019.
- [KBSZ18] Tobias Kappé, Paul Brunet, Alexandra Silva, and Fabio Zanasi. Concurrent Kleene algebra: Free model and completeness. In *Proc. European Symposium on Programming (ESOP)*, pages 856–882, 2018.
- [Kle56] Stephen C. Kleene. Representation of events in nerve nets and finite automata. *Automata Studies*, pages 3–41, 1956.
- [Koz94] Dexter Kozen. A completeness theorem for Kleene Algebras and the algebra of regular events. *Inf. Comput.*, 110(2):366–390, 1994.
- [Koz96] Dexter Kozen. Kleene algebra with tests and commutativity conditions. In *Proc. Tools and Algorithms for the Construction and Analysis of Systems (TACAS)*, pages 14–33, 1996.
- [LRR03] Kamal Lodaya, D. Ranganayakulu, and K. Rangarajan. Hierarchical structure of 1-safe Petri nets. In *Proc. ASIAN, Programming Languages and Distributed Computation*, pages 173–187, 2003.
- [LS14] Michael R. Laurence and Georg Struth. Completeness theorems for Bi-Kleene Algebras and series-parallel rational pomset languages. In *Proc. Relational and Algebraic Methods in Computer Science (RAMiCS)*, pages 65–82, 2014.
- [LW00] Kamal Lodaya and Pascal Weil. Series-parallel languages and the bounded-width property. *Theor. Comp. Sci.*, 237(1):347–380, 2000.
- [MY60] Robert McNaughton and Hisao Yamada. Regular expressions and state graphs for automata. *IRE Trans. Electronic Computers*, 9(1):39–47, 1960.
- [Pra82] Vaughan R. Pratt. On the composition of processes. In *Proc. Principles of Programming Languages (POPL)*, pages 213–223, 1982.
- [Tho68] Ken Thompson. Regular expression search algorithm. *Commun. ACM*, 11(6):419–422, 1968.
- [Win77] Józef Winkowski. An algebraic characterization of the behaviour of non-sequential systems. *Inf. Process. Lett.*, 6(4):105–109, 1977.
- [Win79] Józef Winkowski. An algebraic approach to concurrence. In *MFCS*, volume 74, pages 523–532, 1979.

APPENDIX A. LEMMAS ABOUT POMSETS AND POMSET AUTOMATA

To prove Lemma 3.9, we first recall the following lemma from [KBSZ18].

Lemma A.1. *Let U, V, W, X be sp-pomsets such that $U \parallel V = W \parallel X$. There exist sp-pomsets Y_0, Y_1, Z_0, Z_1 such that*

$$U = Y_0 \parallel Y_1 \quad V = Z_0 \parallel Z_1 \quad W = Y_0 \parallel Z_0 \quad X = Y_1 \parallel Z_1$$

Lemma 3.9. *Parallel factorisations exist uniquely for sp-pomsets.*

Proof. Let V_1, \dots, V_n and W_1, \dots, W_m be sp pomsets that form parallel factorisations of U . We set out to prove that $\{\!\{V_1, \dots, V_n\}\!\} = \{\!\{W_1, \dots, W_m\}\!\}$. The proof proceeds by induction on $\min(n, m)$. In the base, there are two cases to consider. If $\min(n, m) = 0$, then the claim follows by Lemma 3.6. If $\min(n, m) = 1$, the claim follows by definition of parallel primes.

For the inductive step, we assume that the claim holds for n', m' with $\min(n', m') < \min(n, m)$, and that $\min(n, m) > 1$. By Lemma A.1, we find sp-pomsets Y_0, Y_1, Z_0, Z_1 s.t.

$$V_1 = Y_0 \parallel Y_1 \quad V_2 \parallel \dots \parallel V_n = Z_0 \parallel Z_1 \quad W_1 = Y_0 \parallel Z_0 \quad W_2 \parallel \dots \parallel W_m = Y_1 \parallel Z_1$$

Since V_1 is a parallel prime, we find that Y_0 or Y_1 is empty (but not both). In the former case, we find that $V_1 = Y_1$ and $W_1 = Z_0$. Now, let X_1, \dots, X_k be a parallel factorisation of Z_1 . By induction, we then find that $\{\!\{V_2, \dots, V_n\}\!\} = \{\!\{Z_0, X_1, \dots, X_k\}\!\}$, as well as $\{\!\{W_2, \dots, W_m\}\!\} = \{\!\{Y_1, X_1, \dots, X_k\}\!\}$. Consequently, we have that

$$\begin{aligned} \{\!\{V_1, \dots, V_n\}\!\} &= \{\!\{V_1, Z_0, X_1, \dots, X_k\}\!\} \\ &= \{\!\{Y_1, W_1, X_1, \dots, X_k\}\!\} \\ &= \{\!\{W_1, \dots, W_m\}\!\} \end{aligned}$$

In the latter case, we find that Y_0 is not empty, and thus Z_0 is empty, by the fact that W_1 is a parallel prime — thus, $W_1 = Y_0 = V_1$. Furthermore, $V_2 \parallel \dots \parallel V_n = W_2 \parallel \dots \parallel W_m$. Hence, by induction it follows that $\{\!\{V_2, \dots, V_n\}\!\} = \{\!\{W_2, \dots, W_m\}\!\}$; the claim then follows. \square

Lemma 3.12. *Let $L, L' \subseteq \Sigma^*$, and let ζ be a substitution. If ζ is atomic, then*

$$\zeta(L \cap L') = \zeta(L) \cap \zeta(L') \quad \zeta(L \setminus L') = \zeta(L) \setminus \zeta(L') \quad \zeta(L) = \emptyset \iff L = \emptyset$$

Proof. We first claim that for every $U \in \text{SP}(\Delta)$ there exists at most one $w \in \Sigma^*$ such that $U \in \zeta(w)$. To see this, suppose that $w, w' \in \Sigma^*$ such that $U \in \zeta(w)$ and $U \in \zeta(w')$. This means that we can write $w = \mathbf{a}_1 \dots \mathbf{a}_n$ and $w' = \mathbf{a}'_1 \dots \mathbf{a}'_{n'}$, as well as $U = U_1 \dots U_n$ and $U = U'_1 \dots U'_{n'}$ such that for $1 \leq i \leq n$ we have $U_i \in \zeta(\mathbf{a}_i)$, and for $1 \leq i \leq n'$ we have $U'_i \in \zeta(\mathbf{a}'_i)$. Indeed, by the first restriction on ζ and Lemma 3.8, we find that $n = n'$, and that for $1 \leq i \leq n$ we have that $U_i = U'_i$. This also means that for $1 \leq i \leq n$ we have that $U_i \in \zeta(\mathbf{a}_i) \cap \zeta(\mathbf{a}'_i)$; the second restriction on ζ then tells us that for $1 \leq i \leq n$ we have that $\mathbf{a}_i = \mathbf{a}'_i$, which entails that $w = w'$. Hence, we conclude that $w = w'$.

We now treat the claims in the order given.

- (i) First suppose that $U \in \zeta(L \cap L')$; then there exists a $w \in L \cap L'$ such that $U \in \zeta(w)$. We then have that $U \in \zeta(L)$ and $U \in \zeta(L')$, meaning that $U \in \zeta(L) \cap \zeta(L')$.

For the other inclusion, let $U \in \zeta(L) \cap \zeta(L')$. Then there exist $w \in L$ and $w' \in L'$ such that $U \in \zeta(w)$ and $U \in \zeta(w')$; by the above, we conclude that $w = w'$, and hence $w \in L \cap L'$, which means that $U \in \zeta(L \cap L')$.

- (ii) For the third equality, first suppose that $U \in \zeta(L \cap L')$. In that case there exists a $w \in L \setminus L'$ such that $U \in \zeta(w)$. In particular, this means that $w \in L$ and hence $U \in \zeta(L)$. Furthermore, suppose that $U \in \zeta(L')$; we then find a $w' \in L'$ such that $U \in \zeta(w')$. By the above, this would mean that $w = w'$, and hence $w \in L' — a contradiction. We therefore know that $U \notin \zeta(L')$, and hence $U \in \zeta(L) \setminus \zeta(L')$.$

For the other inclusion, let $U \in \zeta(L) \setminus \zeta(L')$. Then we obtain $w \in L$ such that $U \in \zeta(w)$, and furthermore there cannot be a $w' \in L'$ with $U \in \zeta(w')$, which in particular means that $w \notin L'$. We conclude that $w \in L \setminus L'$, and hence $U \in \zeta(L \setminus L')$.

- (iii) Lastly, if $L = \emptyset$, then $\zeta(L) = \emptyset$ by definition. For the other direction, suppose that $L \neq \emptyset$, i.e., that $w \in L$. By the second condition, we have that $\zeta(\mathbf{a}) \neq \emptyset$ for all $\mathbf{a} \in \Sigma$, and hence $\zeta(w)$ cannot be empty, which means that $\zeta(L)$ cannot be empty either. \square

Lemma 3.16. *Let $e \in \mathcal{T}$. Now $e \in \mathcal{F}$ if and only if $1 \in \llbracket e \rrbracket$.*

Proof. The proof from left to right proceeds by induction on the construction of \mathcal{F} . In the base, we have $e = 1$ or $e = f^*$ for some $f \in \mathcal{T}$; in both cases, we find $1 \in \llbracket e \rrbracket$ by definition.

For the inductive step, there are three cases to consider. First, if $e \in \mathcal{F}$ because $e = f + g$ and $f \in \mathcal{F}$ or $g \in \mathcal{F}$, then $1 \in \llbracket f \rrbracket$ or $1 \in \llbracket g \rrbracket$ by induction, and hence $1 \in \llbracket e \rrbracket$. Second, if $e \in \mathcal{F}$ because $e = f \cdot g$ and $f, g \in \mathcal{F}$, then by induction $1 \in \llbracket f \rrbracket$ and $1 \in \llbracket g \rrbracket$, meaning that $1 = 1 \cdot 1 \in \llbracket e \rrbracket$. Third, if $e \in \mathcal{F}$ because $e = f \parallel g$ and $f, g \in \mathcal{F}$, then by induction $1 \in \llbracket f \rrbracket$ and $1 \in \llbracket g \rrbracket$, meaning that $1 = 1 \parallel 1 \in \llbracket e \rrbracket$.

The proof from right to left proceeds by induction on the structure of e . In the base, the only case to consider is $e = 1$, where we find $e \in \mathcal{F}$ immediately.

For the inductive step, there are four cases to consider. First, $e = f + g$, then $1 \in \llbracket e \rrbracket$ implies that $1 \in \llbracket f \rrbracket$ or $1 \in \llbracket g \rrbracket$. In the former case, $f \in \mathcal{F}$ by induction, and hence $e \in \mathcal{F}$ by definition. Second, if $e = f \cdot g$, then $1 \in \llbracket e \rrbracket$ implies that $1 \in \llbracket f \rrbracket$ and $1 \in \llbracket g \rrbracket$. By induction, we then find that $f, g \in \mathcal{F}$, and thus $e \in \mathcal{F}$ by definition. The case where $e = f \parallel g$ can be argued similarly. Lastly, if $e = f^*$, then $e \in \mathcal{F}$ by definition, and there is nothing to prove. \square

Lemma 4.6. *Let $q \xrightarrow{U}_{\rightarrow A} q'$. There exist $q = q_0, \dots, q_\ell = q' \in Q$ and $U_1, \dots, U_\ell \in \text{SP}(\Sigma)$, such that $U = U_1 \cdots U_\ell$, and for all $1 \leq i \leq \ell$ we have that $q_{i-1} \xrightarrow{U_i}_{\rightarrow} q_i$ is a unit run.*

Proof. We proceed by induction on \rightarrow_A . In the base, there are two cases to consider.

- If $q \xrightarrow{U}_{\rightarrow A} q'$ because $q = q'$ and $U = 1$, then we choose $\ell = 0$ to satisfy the claim.
- If $q \xrightarrow{U}_{\rightarrow A} q'$ because $U = \mathbf{a}$ for some $\mathbf{a} \in \Sigma$ and $q' \in \delta(q, \mathbf{a})$, then choose $\ell = 1$ and $U_1 = \mathbf{a}$.

For the inductive step, there are also two cases to consider.

- On the one hand, if $q \xrightarrow{U}_{\rightarrow A} q'$ because $U = V \cdot W$ and there exists a $q'' \in Q$ such that $q \xrightarrow{V}_{\rightarrow A} q''$ and $q'' \xrightarrow{W}_{\rightarrow A} q'$, then we can obtain the necessary states and pomsets by induction from those runs.
- On the other hand, if $q \xrightarrow{U}_{\rightarrow A} q'$ because $U = U_1 \parallel \cdots \parallel U_n$ and there exist $r_1, \dots, r_n \in Q$ as well as $r'_1, \dots, r'_n \in F$ such that for $1 \leq i \leq n$ we have $r_i \xrightarrow{U'_i}_{\rightarrow} r'_i$, and furthermore $q' \in \gamma(q, \{\!\{r_1, \dots, r_n\}\!\})$, then we can choose $\ell = 1$ and $U_1 = U$ to satisfy the claim. \square

APPENDIX B. LEMMAS TOWARDS THE DECISION PROCEDURE

Lemma 5.7. *If A is well-structured and q is a fork target in A , then all pomsets in $L_A(q)$ are parallel primes.*

Proof. Let $p, q \in Q$ and $\phi \in \mathbb{M}(Q)$ such that $\gamma(p, \phi) \neq \emptyset$ and $q \in \phi$. Suppose, towards a contradiction, that $U \in L_A(q)$ is parallel. We then know that there exists a $q' \in F$ with $q \xrightarrow{U}_A q'$, which is a parallel unit run by Lemma 5.6. In particular, $q' \in \gamma(q, \phi')$ for some $\phi' \in \mathbb{M}(Q)$. However, the latter contradicts the fact that $q' \in F$, by the premise that A is well-structured. It thus follows that U must be either primitive or sequential. \square

Fact 5.16. *For all $q \in Q$, it holds that $L_A(q) = \zeta(L_{A'}(q))$.*

Proof. It suffices to prove that, for $q, q' \in Q$ as well as $U \in \text{SP}(\Sigma)$, it holds that $q \xrightarrow{U}_A q'$ is a unit run if and only if there exists an $\mathbf{a} \in \Delta$ with $U \in \zeta(\mathbf{a})$ and $q' \in \delta'(q, \mathbf{a})$; straightforward inductive arguments on run length (for the forward inclusion, using Lemma 4.6) and word length (for the backwards inclusion) then complete the proof.

First suppose that $q \xrightarrow{U}_A q'$ is a unit run. We have two cases to consider.

- If $q \xrightarrow{U}_A q'$ is a sequential unit run, then $U = \mathbf{b}$ for some $\mathbf{b} \in \Sigma$, and $q' \in \delta(q, \mathbf{b})$. We can then choose $\mathbf{a} = \mathbf{b}$ to find that $q' \in \delta'(q, \mathbf{a})$ by definition of δ' .
- If $q \xrightarrow{U}_A q'$ is a parallel unit run, then $U = U_1 \parallel \dots \parallel U_n$ and we obtain $q_1, \dots, q_n \in Q$ such that $q' \in \gamma(q, \{q_1, \dots, q_n\})$ as well as $U_i \in L_A(q_i)$. For $1 \leq i \leq n$, we then choose $\alpha_i = \{r \in Q' : U_i \in L_A(r)\} \in \text{At}_{A[Q']}$, and set $\mathbf{a} = \{\alpha_1, \dots, \alpha_n\} \in \Delta$. We now find that $U \in \zeta(\mathbf{a})$ as well as $q' \in \delta'(q, \{\alpha_1, \dots, \alpha_n\})$.

Conversely, suppose $\mathbf{a} \in \Delta$ s.t. $U \in \zeta(\mathbf{a})$ and $q' \in \delta'(q, \mathbf{a})$. We have two cases to consider.

- If $\mathbf{a} \in \Sigma$, then $q' \in \delta(q, \mathbf{a})$ and $U = \mathbf{a}$. Hence $q \xrightarrow{U}_A q'$ is a sequential unit run.
- If $\mathbf{a} = \{\alpha_1, \dots, \alpha_n\}$, then $q' \in \gamma(q, \{q_1, \dots, q_n\})$ with $q_i \in \alpha_i$. Since $U \in \zeta(\mathbf{a})$, also $U = U_1 \parallel \dots \parallel U_n$ with $U_i \in L_A(q_i)$. Thus, $q \xrightarrow{U}_A q'$ is a parallel unit run. \square

Fact 5.17. *The substitution ζ is atomic.*

Proof. If $\mathbf{a} \in \Delta$ and $U \in \zeta(\mathbf{a})$, then $U = \mathbf{b}$ for some $\mathbf{b} \in \Sigma$, or $U \in L_A(\alpha_1) \parallel \dots \parallel L_A(\alpha_n)$ for atoms $\alpha_1, \dots, \alpha_n \in \text{At}_{A[Q']}$. In the former case, we know that U is primitive, and hence a sequential prime. In the latter case, we know that none of these atoms contain the empty set and $n \geq 2$ (since A is well-structured); it follows that U is a sequential prime.

Furthermore, suppose $\mathbf{a}, \mathbf{b} \in \Delta$ such that $U \in \zeta(\mathbf{a})$ and $U \in \zeta(\mathbf{b})$. By the above, we know that U is either primitive and $\mathbf{a}, \mathbf{b} \in \Sigma$, or U is parallel and $\mathbf{a}, \mathbf{b} \in \Delta \setminus \Sigma$. In the former case, it follows that $\mathbf{a} = U = \mathbf{b}$ by definition of ζ . In the latter case, we know that $U = U_1 \parallel \dots \parallel U_n$ and $U = U'_1 \parallel \dots \parallel U'_{n'}$, as well as $\mathbf{a} = \{\alpha_1, \dots, \alpha_n\}$ and $\mathbf{b} = \{\alpha'_1, \dots, \alpha'_{n'}\}$, such that for $1 \leq i \leq n$ we have that $U_i \in L_A(\alpha_i)$ and for $1 \leq i \leq n'$ we have that $U'_i \in L_A(\alpha'_i)$. Because A is well-structured, each of the U_i and U'_i is a parallel prime; by Lemma 3.9, we then find that $n = n'$, and without loss of generality we have that $U_i = U'_i$. Since $U_i \in L_A(\alpha_i) \cap L_A(\alpha'_i)$, it follows that $\alpha_i = \alpha'_i$, and hence $\mathbf{a} = \mathbf{b}$. \square

APPENDIX C. LEMMAS ABOUT WELL-STRUCTURED AUTOMATA

The following characterisation of runs labelled by the empty pomset will be useful.

Fact C.1. *Let \rightsquigarrow_A be the smallest relation on Q satisfying the rules*

$$\frac{}{p \rightsquigarrow_A p} \quad \frac{p \rightsquigarrow_A r \quad r \rightsquigarrow_A q}{p \rightsquigarrow_A q} \quad \frac{q \in \gamma(p, \{q_1, \dots, q_n\}) \quad \forall 1 \leq i \leq n. q_i \rightsquigarrow_A q'_i \in F}{p \rightsquigarrow_A q}$$

Now $p \rightsquigarrow_A q$ holds if and only if $p \xrightarrow{1}_A q$.

Proof. The first inclusion is proved by induction on the construction of \rightsquigarrow_A . In the base, $p \rightsquigarrow_A q$ because $p = q$, thus $p \xrightarrow{1}_A q$ immediately. In the inductive step, there are two cases.

- If $p \rightsquigarrow_A q$ because there exists an $r \in Q$ such that $p \rightsquigarrow_A r$ and $r \rightsquigarrow_A q$, then by induction we have that $p \xrightarrow{1}_A r$ and $r \xrightarrow{1}_A q$. It then follows that $p \xrightarrow{1}_A q$.
- Suppose $p \rightsquigarrow_A q$ because there exist $q_1, \dots, q_n \in Q$ such that $q \in \gamma(p, \{q_1, \dots, q_n\})$ and $q'_1, \dots, q'_n \in F$ such that for $1 \leq i \leq n$ we have that $q_i \xrightarrow{1}_A q'_i$. By induction, we find for $1 \leq i \leq n$ that $q_i \xrightarrow{1}_A q'_i$. We can then conclude that $p \xrightarrow{1}_A q$.

For the other inclusion, we proceed by induction on the construction of $\xrightarrow{1}_A$. In the base, we have $p \xrightarrow{1}_A q$ because $p = q$, and so we are done. The case where $p \xrightarrow{1}_A q$ is a sequential unit run can be disregarded. For the inductive step, there are two cases to consider.

- If $p \xrightarrow{1}_A q$ because $1 = U \cdot V$ and there exists an $r \in Q$ such that $p \xrightarrow{U}_A r$ and $r \xrightarrow{V}_A q$, then note that $U = V = 1$. Hence $p \rightsquigarrow_A r$ and $r \rightsquigarrow_A q$ by induction, meaning $p \rightsquigarrow_A q$.
- Suppose $p \xrightarrow{1}_A q$ because there exist $q_1, \dots, q_n \in Q$ such that $q \in \gamma(p, \{q_1, \dots, q_n\})$ and there exist $q'_1, \dots, q'_n \in F$ such that $1 = U_1 \parallel \dots \parallel U_n$ and for $1 \leq i \leq n$ we have $q_i \xrightarrow{U_i}_A q'_i$. By induction, we find for $1 \leq i \leq n$ that $q_i \rightsquigarrow_A q'_i$. We can then conclude that $p \rightsquigarrow_A q$. \square

To prove Lemma 6.5, the following is useful.

Fact C.2. *Let A be a 1-forking PA such that if q is a fork target, then $q \notin F$. For any states $p, q \in Q$, we have $p \xrightarrow{1}_A q$ iff $p = q$. Furthermore, $1 \in L_A(p)$ iff $p \in F$.*

Proof. For the first part, the implication from right to left holds by definition of \rightarrow_A . For the other implication, note that $p \xrightarrow{1}_A q$ is equivalent to $p \rightsquigarrow_A q$ by Fact C.1. We proceed by induction on the derivation $p \rightsquigarrow_A q$. In the base, we have that $p \rightsquigarrow_A q$ because $p = q$, and we are done. In the inductive step, there are two cases to consider.

- If $p \rightsquigarrow_A q$ because $p \rightsquigarrow_A r$ and $r \rightsquigarrow_A q$ for some $r \in Q$, then by induction $p = r = q$.
- If $p \rightsquigarrow_A q$ because there exist $q_1, \dots, q_n \in Q$ such that $r \in \gamma(p, \{q_1, \dots, q_n\})$, and $q'_1, \dots, q'_n \in F$ such that for $1 \leq i \leq n$ we have that $q_i \rightsquigarrow_A q'_i$, then by induction we have for $1 \leq i \leq n$ that $q_i = q'_i$. Since A is 1-forking, it follows that $q_1 = q'_1 \in F$. But this contradicts our premise, because q_1 is a fork target; we can therefore exclude this case.

As for the second part, the implication from right to left holds by definition of $L_A(p)$. For the other implication, note that if $1 \in L_A(p)$, then $p \xrightarrow{1}_A q$ for some $q \in F$. Since $p = q$ by the first part, we conclude that $p \in F$. \square

Lemma 6.5. *A pomset automaton A is well-structured if and only if it is 2-forking, parsimonious, and flat-branching.*

Proof. By definition, an automaton is well-structured iff it is 2-forking, flat-branching, and satisfies the property that for all $p \in Q$ and $q \in \phi \in \mathbb{M}(Q)$ with $\gamma(p, \phi) \neq \emptyset$, it holds that $q \notin F$. If the automaton is parsimonious, then this property is satisfied. By Fact C.2, if A is well-structured, then $q \notin F \Leftrightarrow 1 \notin L_A(q)$, so it is in particular parsimonious. \square

Lemma 6.6. *For any bounded PA A , the predicate $p \xrightarrow{1}_A q$ is decidable.*

Proof. By Fact C.1, it suffices to show that \rightsquigarrow_A is decidable. Since A finite, we can build \rightsquigarrow_A by saturation, starting from the identity and adding pairs until we reach a fixpoint; this fixpoint is reached after finitely many steps as a result of boundedness. \square

Lemma 6.8. *If a PA A' weakly implements A , then we can construct a PA A'' implementing A . If A' is bounded (resp. n -forking, flat-branching, parsimonious), then so is A'' .*

Proof. Let $A = \langle Q, F, \delta, \gamma \rangle$, and $A' = \langle Q', F', \delta', \gamma' \rangle$. We can assume without loss of generality that Q and Q' are disjoint. We define $A'' = \langle Q'', F'', \delta'', \gamma'' \rangle$ as follows:

$$\begin{aligned} Q'' &= Q' \uplus Q & F'' &= F' \uplus \{q \in Q : 1 \in L_A(q)\} \\ \delta''(q'', \mathbf{a}) &= \begin{cases} \delta'(q'', \mathbf{a}) & q'' \in Q' \\ \bigcup_{x \in Q_{q''}} \delta'(x, \mathbf{a}) & q'' \in Q \end{cases} & \gamma''(q'', \phi) &= \begin{cases} \gamma'(q'', \phi) & q'' \in Q' \\ \bigcup_{x \in Q_{q''}} \gamma'(x, \phi) & q'' \in Q \end{cases} \end{aligned}$$

Here, $Q_{q''}$ is the set of states of Q implementing $q'' \in Q$, by weak implementation. Note that if $q'' \in Q$, then there are finitely many $\phi \in \mathbb{M}(Q'')$ such that $\gamma''(q'', \phi) \neq \emptyset$, because $Q_{q''}$ is finite, and for each $x \in Q_{q''}$ there are finitely many $\phi \in \mathbb{M}(Q')$ such that $\gamma'(x, \phi) \neq \emptyset$.

We now prove that A'' implements A . First, we show for $q \in Q$ that $L_A(q) = L_{A''}(q)$.

- For the inclusion from left to right, suppose that $U \in L_A(q)$. If $U = 1$, then $q \in F''$ by definition of F'' , and thus $U = 1 \in L_{A''}(q'')$ immediately.

On the other hand, suppose $U \neq 1$. There exists an $x \in Q_q$ such that $U \in L_{A'}(x)$, because A' weakly implements A . Hence, there must exist a $q' \in F'$ such that $x \xrightarrow{U}_{A'} q'$; a straightforward inductive argument then shows that $x \xrightarrow{U}_{A''} q'$. Now, since $U \neq 1$, the latter run must be non-trivial. We thus find $x' \in Q''$ and $U = V \cdot W$ such that $x \xrightarrow{V}_{A''} x'$ is a unit run, and $x' \xrightarrow{W}_{A''} q'$. By construction of A'' , it follows that $q \xrightarrow{V}_{A''} x'$ is a unit run too, and hence $q \xrightarrow{U}_{A''} q'$. Since $q' \in F''$, it follows that $U \in L_{A''}(q)$.

- For the inclusion from right to left, suppose that $U \in L_{A''}(q)$. There must then exist a $q' \in F''$ such that $q \xrightarrow{U}_{A''} q'$. Now, if this run is trivial, then $q = q'$ and $U = 1$. This means that $q \in F''$; since $q \in Q$, it follows that $q \notin F'$, and hence $U = 1 \in L_A(q)$.

Otherwise, if $q \xrightarrow{U}_{A''} q'$ is non-trivial, then there exists a $q'' \in Q''$ and $U = V \cdot W$ such that $q \xrightarrow{V}_{A''} q''$ is a unit run, and $q'' \xrightarrow{W}_{A''} q'$. By construction of A'' , we find an $x \in Q_q$ such that $x \xrightarrow{V}_{A''} q''$ is a unit run, and hence $x \xrightarrow{U}_{A''} q'$. A simple inductive argument then shows that $x \xrightarrow{U}_{A'} q'$ and $q' \in F'$, and thus $U \in L_{A'}(x)$. Because A' weakly implements A , we conclude that $U \in L_A(q)$.

We should also show that if A is fork-acyclic, then so is A'' . To this end, note that fork-acyclicity of A already implies fork-acyclicity of A' ; hence, we need only show that our construction above preserves fork-acyclicity. The following property is useful.

Fact C.3. *For any $q \in Q'$ and $q' \in Q''$, if $q' \preceq_{A''} q$ then $q' \in Q'$ and $q' \preceq_{A'} q$.*

Proof. It suffices to argue that if $q' \preceq_{A''} q$ arises from one of the rules that generate $\preceq_{A''}$, then $q' \in Q'$ and $q' \preceq_{A'} q$. This is straightforward, for if $q \in Q'$ then $\delta''(q, \mathbf{a})$ coincides with $\delta'(q, \mathbf{a})$ for all $\mathbf{a} \in \Sigma$, and $\gamma''(q, \phi)$ coincides with $\gamma'(q, \phi)$ for all $\phi \in \mathbb{M}(Q'')$. \square

Let $r, q \in Q''$ and $\phi \in \mathbb{M}(Q'')$ with $r \in \phi$ and $\gamma''(q, \phi) \neq \emptyset$. By construction of A'' , we have that $r \in Q'$. If $q \preceq_{A''} r$, then by Fact C.3 we have that $q \in Q'$ and $q \preceq_{A'} r$. In that case, also $\gamma'(q, \phi) \neq \emptyset$, and hence $r \preceq_{A'} q$. This, however, contradicts the premise that A' is fork-acyclic; we thus have that $r \prec_{A''} q$. Hence, A'' must be fork-acyclic as well.

We should also show that our construction preserves boundedness, n -forking, flat-branching and parsimony. For boundedness, the following is useful.

Fact C.4. *For any $q \in Q$ and $q' \in Q''$, if $q' \preceq_{A''} q$, then $q = q'$, or $q' \in Q'$ such that $q' \preceq_{A'} x$ for some $x \in Q_q$.*

Proof. To see this, we proceed by induction on the construction of $\preceq_{A''}$. In the base, either $q = q'$ by reflexivity, or one of three cases applies.

- If $q' \preceq_{A''} q$ because $q' \in \delta''(q, \mathbf{a}) \subseteq Q'$ for some $\mathbf{a} \in \Sigma$, then $q' \in \delta'(x, \mathbf{a})$ for some $x \in Q_q$, and therefore $q' \preceq_{A'} x$.
- If $q' \preceq_{A''} q$ because $q \in \gamma''(q, \phi)$ for some $\phi \in \mathbb{M}(Q'')$, then $q' \in \gamma'(x, \phi) \subseteq Q'$ for some $x \in Q_q$, and therefore $q' \preceq_{A'} x$.
- If $q' \preceq_{A''} q$ because there exists a $\phi \in \mathbb{M}(Q'')$ such that $q' \in \phi$ and $\gamma''(q, \phi) \neq \emptyset$, then $\phi \in \mathbb{M}(Q')$ and there exists an $x \in Q_q$ such that $\gamma'(x, \phi) \neq \emptyset$, and hence $q' \preceq_{A'} x$.

For the inductive step, $q' \preceq_{A''} q$ because $q'' \in Q''$ and $q' \preceq_{A''} q''$ and $q'' \preceq_{A''} q$. By induction, either $q = q''$ (in which case the claim follows by applying induction to $q' \preceq_{A''} q'' = q$), or $q'' \preceq_{A'} x$ for some $x \in Q_q$, in which case $q' \preceq_{A'} q''$ by Fact C.3. \square

By Fact C.4 we have that the support of $q \in Q$ in A'' is given by $\{q\} \cup \bigcup_{x \in Q_q} \pi_{A'}(x)$; if A' is bounded, this set must be finite. Furthermore, by C.3, we know that the support of $q \in Q'$ in A'' is given by the support of q in A' , which is finite if A' is bounded. Thus, if A' is bounded, then so is A'' .

We now prove that our construction preserves n -forking, parsimony and flat-branching.

- If A' is n -forking, let $\phi \in \mathbb{M}(Q'')$ and $q \in Q''$ such that $\gamma''(q, \phi) \neq \emptyset$. By construction there exists $q' \in Q'$ such that $\gamma'(q', \phi) \neq \emptyset$. Since A' is n -forking we may conclude $|\phi| \geq n$.
- Assume, towards a contradiction, that A'' is not parsimonious. Then there exist $q \in Q''$ and $r \in \phi \in \mathbb{M}(Q'')$ such that $\gamma''(q, \phi) \neq \emptyset$ but $1 \in L_{A''}(r)$. By construction there exists $q' \in Q'$ such that $\gamma'(q', \phi) \neq \emptyset$. We also know that it must be the case that $r \in Q'$, hence $1 \in L_{A'}(r)$. This contradicts the premise that A' is parsimonious.
- Assume, towards a contradiction, that A is not flat-branching. There must then exist $q \in Q''$ and $\phi, \psi \in \mathbb{M}(Q'')$ with $r \in \phi$ such that $\gamma''(q, \phi) \neq \emptyset$ and $\gamma''(r, \psi) \cap F \neq \emptyset$. By construction of A'' , we find that $\phi \in \mathbb{M}(Q')$, and there exists a $q' \in Q'$ such that $\gamma'(q', \phi) \neq \emptyset$. Also by construction of A'' and the fact that $r \in \phi$ (and hence $r \in Q'$), we find that $\psi \in \mathbb{M}(Q')$ and $\gamma'(r, \psi) \cap F' \neq \emptyset$. This contradicts the premise that A' is flat-branching. \square

C.1. Ensuring parsimony.

Lemma 6.11. A_0 is bounded, parsimonious, and weakly implements A .

Proof. We relate runs in A to runs in A_0 and vice versa, as follows:

Fact C.5. For any run $p \xrightarrow{U}_A q$ we have $p \xrightarrow{U}_{A_0} q$, and if in addition we know that $1 \in L_A(q)$ and $U \neq 1$, then $p \xrightarrow{U}_{A_0} \top$.

Proof. In the base, there are two cases.

- If $p \xrightarrow{U}_A q$ because $p = q$ and $U = 1$, then $p \xrightarrow{U}_{A_0} q$; the second claim holds vacuously.
- If $p \xrightarrow{U}_A q$ because $U = \mathbf{a}$ for some $\mathbf{a} \in \Sigma$ and $q \in \delta(p, \mathbf{a})$, then $q \in \delta_0(p, \mathbf{a})$ by construction, so indeed $p \xrightarrow{U}_{A_0} q$. If $1 \in L_A(q)$, then $\top \in \delta_0(p, \mathbf{a})$ as well, so $p \xrightarrow{\mathbf{a}}_{A_0} \top$.

For the inductive step, we have two more cases.

- If $p \xrightarrow{U}_A q$ because $U = V \cdot W$ and there exists $r \in Q$ such that $p \xrightarrow{V}_A r$ and $r \xrightarrow{W}_A q$, then by induction we know that $p \xrightarrow{V}_{A_0} r$ and $r \xrightarrow{W}_{A_0} q$, so $p \xrightarrow{U}_{A_0} q$.

If furthermore $1 \in L_A(q)$ and $V \cdot W \neq 1$, then we distinguish two subcases.

- If $W = 1$, then $V = U$, and $1 \in L_A(r)$; therefore by induction we get $p \xrightarrow{U}_{A_0} \top$.
- If $W \neq 1$, then by induction we get $r \xrightarrow{W}_{A_0} \top$, hence $p \xrightarrow{V \cdot W}_{A_0} q$.
- Suppose that $p \xrightarrow{U}_A q$ because there exist $q_1, \dots, q_n \in Q$ with $q \in \gamma(p, \{q_1, \dots, q_n\})$, and there exist $q'_1, \dots, q'_n \in F$ such that $U = U_1 \parallel \dots \parallel U_n$ and for $1 \leq i \leq n$ we have $q_i \xrightarrow{U_i}_A q'_i$. We then partition $\{q_1, \dots, q_n\}$ into ϕ and ψ such that:

$$\{q_1, \dots, q_n\} = \phi \sqcup \psi \quad \forall i. q_i \in \phi \Rightarrow U_i \neq 1 \quad \forall i. q_i \in \psi \Rightarrow U_i = 1.$$

By construction of γ_0 , we have $q \in \gamma_0(p, \phi)$. By induction, since $\forall i. q_i \in \phi \Rightarrow U_i \neq 1$, we get $\forall i. q_i \in \phi \Rightarrow q_i \xrightarrow{U_i}_{A_0} \top$. Consequently, we have $p \xrightarrow{U_1 \parallel \dots \parallel U_n}_{A_0} q$.

If additionally $1 \in L_A(q)$ and $U_1 \parallel \dots \parallel U_n \neq 1$, then there must exist a $1 \leq i \leq n$ with $U_i \neq 1$, so we know that $\phi \neq \emptyset$; hence, $\top \in \gamma_0(p, \phi)$. As result, we get $p \xrightarrow{U_1 \parallel \dots \parallel U_n}_{A_0} \top$. \square

Fact C.6. For any $p \in Q$, if $p \xrightarrow{U}_{A_0} q$, then if $q \in Q$ we have $p \xrightarrow{U}_A q$, and if $q = \top$ we have $U \in L_A(p)$ but $U \neq 1$.

Proof. In the base, there are two cases.

- If $p \xrightarrow{U}_{A_0} q$ because $U = 1$ and $p = q$, then $p \xrightarrow{U}_A q$ immediately.
- If $p \xrightarrow{U}_{A_0} q$ because $U = \mathbf{a}$ for some $\mathbf{a} \in \Sigma$ and $q \in \delta_0(p, \mathbf{a})$, then there are two subcases:
 - If $q \in Q$, then $q \in \delta(p, \mathbf{a})$, and hence $p \xrightarrow{U}_A q$.
 - If $q = \top$, then there exists a $q' \in \delta(p, \mathbf{a})$ such that $1 \in L_A(q')$. In particular, this means that $p \xrightarrow{U}_A q'$ and $q' \xrightarrow{1}_A q''$ for some $q'' \in F$. This means that $p \xrightarrow{U}_A q''$, and hence $U \in L_A(p)$. Since $\mathbf{a} \neq 1$, we rightfully get $U \neq 1$.

For the inductive step, there are again two cases.

- Suppose $p \xrightarrow{U}_{A_0} q$ because $U = V \cdot W$ and there exists an $r \in Q_0$ with $p \xrightarrow{V}_{A_0} r$ and $r \xrightarrow{W}_{A_0} q$. We distinguish two cases based on r :
 - If $r = \top$, then by construction of A_0 we have $W = 1$ and $q = \top$, so by induction it follows that $U = V \in L_A(p)$ while $U = V \neq 1$;
 - If $r \in Q$, then by induction hypothesis we have $p \xrightarrow{U}_A r$. We distinguish two cases:
 - * If $q \in Q$, then by induction $r \xrightarrow{V}_A q$, so $p \xrightarrow{U \cdot V}_A q$.
 - * If $q = \top$, then by induction $W \in L_A(r)$ while $W \neq 1$, so $U \in L_A(p)$ while $U \neq 1$.
- Suppose $p \xrightarrow{U}_{A_0} q$ because there exist $q_1, \dots, q_n \in Q_0$ such that $q \in \gamma(p, \{q_1, \dots, q_n\})$ such that $U = U_1 \parallel \dots \parallel U_n$ and for $1 \leq i \leq n$ we have $q_i \xrightarrow{U_i}_{A_0} \top$. By construction of γ_0 , we also know that for $1 \leq i \leq n$ we have $q_i \in Q$. By induction, we then know that for $1 \leq i \leq n$ we have $U_i \in L_A(q_i)$ but $U_i \neq 1$. We distinguish the two cases for q :
 - If $q \in Q$, then by construction of γ_0 there exists a multiset ψ such that:

$$q \in \gamma(p, \{q_1, \dots, q_n\} \sqcup \psi) \quad \forall r \in \psi. 1 \in L_A(r).$$

If $\psi = \{q_{n+1}, \dots, q_m\}$, we may then complete the run by setting $\forall n < i \leq m. U_i = 1$ and performing the following γ step:

$$\frac{\forall 1 \leq i \leq m. U_i \in L_A(q_i) \quad q \in \gamma(p, \{q_1, \dots, q_m\})}{p \xrightarrow{U_1 \parallel \dots \parallel U_m}_A q}$$

This concludes this case, since we have:

$$U_1 \parallel \dots \parallel U_m = U_1 \parallel \dots \parallel U_n \parallel 1 \parallel \dots \parallel 1 = U_1 \parallel \dots \parallel U_n.$$

– Similarly, if $q = \top$, then $\phi \neq \sqcup$ and there exist $\psi \in \mathbb{M}(Q)$ and $q' \in Q$ such that:

$$q' \in \gamma(p, \{q_1, \dots, q_n\} \sqcup \psi) \quad \forall r \in \psi. 1 \in L_A(r) \quad 1 \in L_A(q').$$

We may proceed as in the previous case to show that $p \xrightarrow{U_1 \parallel \dots \parallel U_n}_A q'$. Since $1 \in L_A(q')$, this implies $U_1 \parallel \dots \parallel U_n \in L_A(p)$. To conclude, note that for all $1 \leq i \leq n$ we have $U_i \neq 1$; since $\phi \neq \sqcup$, also $n > 0$. This entails that $U = U_1 \parallel \dots \parallel U_n \neq 1$. \square

Summing up, we get that $L_A(q) \setminus \{1\} = L_{A_0}(q)$. So, if $1 \notin L_A(q)$, then we have $L_A(q) = L_{A_0}(q)$, and otherwise $L_A(q) = L_{A_0}(q) \cup L_{A_0}(\top)$. Therefore, A_0 weakly implements A , since fork-acyclicity is clearly preserved by this construction.

To show that A_0 is parsimonious, notice that by construction, if $\gamma_0(p, \phi) \neq \emptyset$ and $q \in \phi$, then $q \in Q$. As we showed above, this implies $1 \notin L_{A_0}(q)$.

To show that A_0 is bounded, note that if $q, q' \in Q_0$ with $q' \prec_{A_0} q$, then either $q' = \top$, $q = \top$ or $q, q' \in Q$ with $q' \preceq_A q$. Hence, if $q \in Q_0$, then either $q = \top$ and $\pi_{A_0}(q) = \{\top\}$, or $q \in Q$ and $\pi_{A_0}(q) \subseteq \pi_A(q)$; in either case, $\pi_{A_0}(q)$ is finite, since A is bounded; hence, A_0 is bounded, too. \square

C.2. Removing nullary forks.

Lemma 6.14. A_1 is bounded, 1-forking, parsimonious, and weakly implements A .

Proof. To see that A_1 is bounded, note that $q' \preceq_{A_1} q$ if and only if $q' \preceq_A q$; hence, $\pi_{A_1}(q) \subseteq \pi_A(q)$ for all $q \in Q$; since A is bounded, this implies that A_1 is bounded, too. This also show that A_1 is fork-acyclic if A_1 is.

The condition $\phi \neq \sqcup$ in the definition of γ_1 ensures that A_1 is 1-forking. Before we show that A_1 is parsimonious, note that if we have $r \in \phi$ and $\gamma_1(q, \phi) \neq \emptyset$, then there is a state q' such that $q \xrightarrow{1}_A q'$ and $\gamma(q', \phi) \neq \emptyset$. By parsimony of A , we have $1 \notin L_A(r)$, hence $r \notin F$. By Fact C.2, we conclude that $1 \notin L_{A_1}(p)$, hence A_1 is parsimonious.

We now check that A_1 weakly implements A , by relating their runs.

Fact C.7. If $q \xrightarrow{1}_A p \xrightarrow{U}_A p' \xrightarrow{1}_A q'$ with $U \neq 1$, then $q \xrightarrow{U}_{A_1} q'$.

Proof of Fact C.7. We proceed by induction on the construction of $p' \xrightarrow{U}_A q'$. In the base, we can exclude the case where $p' = q'$ and $U = 1$, because it contradicts the premise. This leaves the case where $p \xrightarrow{U}_A p'$ because $U = \mathbf{a}$ for some $\mathbf{a} \in \Sigma$ and $p' \in \delta(p, \mathbf{a})$. By construction of δ_1 , we then find that $q' \in \delta_1(q, \mathbf{a})$, hence $q \xrightarrow{U}_{A_1} q'$. For the inductive step, there are two cases:

- Suppose $p \xrightarrow{U}_A p'$ because $U = V \cdot W$ and there exists $p'' \in Q$ such that $p \xrightarrow{V}_A p''$ and $p'' \xrightarrow{W}_A p'$. We distinguish four subcases:
 - If $V = 1 = W$, then $U = V \cdot W = 1$, contradicting the premise; we disregard this case.
 - If $V \neq 1 = W$, then $p'' \xrightarrow{1}_A q'$ and $U = V$; by induction, we find that $q \xrightarrow{U}_{A_1} q'$.
 - If $V = 1 \neq W$, then $q \xrightarrow{1}_A p''$ and $U = W$; by induction, we find that $q \xrightarrow{U}_{A_1} q'$.
 - If $V \neq 1 \neq W$, then induction we get $q \xrightarrow{V}_{A_1} p''$ and $p'' \xrightarrow{W}_{A_1} q'$, hence $q \xrightarrow{U}_{A_1} q'$.
- Suppose $p \xrightarrow{U}_A p$ because there exist $r_1, \dots, r_n \in Q$ such that $p' \in \gamma(p, \{r_1, \dots, r_n\})$, and $U = U_1 \parallel \dots \parallel U_n$ such that for $1 \leq i \leq n$ there exists $r'_i \in F$ with $r_i \xrightarrow{U_1}_A r'_i$. By parsimony of A , we know that each of the U_i is different from 1. Hence, $U_1 \parallel \dots \parallel U_n$ is itself non-empty, and we have $n > 0$. Therefore $q' \in \gamma_1(q, \{r_1, \dots, r_n\})$. By induction we have for every $1 \leq i \leq n$ that $r_i \xrightarrow{U_i}_{A_1} r'_i \in F$, so we may conclude $q \xrightarrow{U}_{A_1} q'$. \square

Fact C.8. *If $q \xrightarrow{U}_{A_1} q'$, then $q \xrightarrow{U}_A q'$.*

Proof of Fact C.8. We proceed by induction on $q \xrightarrow{U}_{A_1} q'$. In the base, there are two cases.

- If $q \xrightarrow{U}_{A_1} q'$ because $q = q'$ and $U = 1$, then $q \xrightarrow{U}_A q'$ immediately.
- If $q \xrightarrow{U}_{A_1} q'$ because $U = \mathbf{a}$ for some $\mathbf{a} \in \Sigma$ and $q' \in \delta_1(q, \mathbf{a})$, then there exist $p, p' \in Q$ such that $q \xrightarrow{1}_A p$ and $p' \xrightarrow{1}_A q'$ with $p' \in \delta(p, \mathbf{a})$, by construction of δ_1 . We can string these together to find that $q \xrightarrow{1}_A p \xrightarrow{\mathbf{a}}_A p' \xrightarrow{1}_A q'$, hence $q \xrightarrow{U}_A q'$.

In the inductive step, there are again two cases:

- Suppose $q \xrightarrow{U}_{A_1} q'$ because $U = V \cdot W$ and there exists a $q'' \in Q$ such that $q \xrightarrow{V}_{A_1} q''$ and $q'' \xrightarrow{W}_{A_1} q'$. By induction, $q \xrightarrow{V}_A q''$ and $q'' \xrightarrow{W}_A q'$, and hence $q \xrightarrow{U}_A q'$.
- Suppose $q \xrightarrow{U}_{A_1} q'$ because there exist $r_1, \dots, r_n \in Q$ with $q' \in \gamma_1(q, \{\!\!| r_1, \dots, r_n \!\!\})$, and $U = U_1 \parallel \dots \parallel U_n$ and for $1 \leq i \leq n$ there exists $r'_i \in F$ with $r_i \xrightarrow{U_i}_{A_1} r'_i$. By induction, it must be the case that for $1 \leq i \leq n$ we have $r_i \xrightarrow{U_i}_A r'_i$. By construction of γ_1 , we know that there are $p, p' \in Q$ such that $q \xrightarrow{1}_A p$, $p' \xrightarrow{1}_A q'$, and $p' \in \gamma(p, \{\!\!| q_1, \dots, q_n \!\!\})$. We then know that $p \xrightarrow{U_1 \parallel \dots \parallel U_n}_A p'$, and hence $q \xrightarrow{U_1 \parallel \dots \parallel U_n}_A q'$. \square

We can now wrap up by showing that, for $q \in Q$, we have $L_A(q) = \bigcup_{q \xrightarrow{1}_{A_1} p} L_{A_1}(p)$.

- Let $U \in L_A(q)$, meaning there is $q' \in F$ such that $q \xrightarrow{U}_A q'$. On the one hand, if $U = 1$, then $1 \in L_{A_1}(q')$; since $q \xrightarrow{1}_A q'$, U is contained in the right-hand side. On the other hand, if $U \neq 1$, then $q \xrightarrow{U}_{A_1} q'$ by Fact C.7, hence $U \in L_{A_1}(p)$. Since $q \xrightarrow{1}_A q$, we are done.
- Let $p \in Q$ such that $q \xrightarrow{1}_A p$, and let $U \in L_{A_1}(p)$. This means there is $p' \in F$ such that $p \xrightarrow{U}_{A_1} p'$, and hence $p \xrightarrow{U}_A p'$ by Fact C.8. Therefore $q \xrightarrow{U}_A p' \in F$, so $U \in L_A(q)$.

Since boundedness and fork-acyclicity are preserved, A_1 weakly implements A . \square

C.3. Removing unary forks.

Lemma 6.19. *A_2 is bounded, parsimonious, 2-forking, and implements A .*

Proof. We start by showing that our construction preserves fork-acyclicity, and that A_2 is bounded. For preservation of fork-acyclicity, we first verify the following.

Fact C.9. *The following hold for all $w = q_1 \cdots q_n \in Q_2$ and $w' = q'_1 \cdots q'_m \in Q_2$:*

- (i) *If $q \in Q$ such that $q \uparrow w$, then for all q_i we have $q_i \preceq_A q$.*
- (ii) *If $w' \in \delta_2(w, \mathbf{a})$, then for every q'_i there exists a q_j with $q'_i \preceq_A q_j$.*
- (iii) *If $w' \in \gamma_2(w, \phi)$, then for every q'_i there exists a q_j with $q'_i \preceq_A q_j$.*
- (iv) *If $\gamma_2(w, \phi) \neq \emptyset$ with $r \in \phi$, then there exists a q_i with $r \prec_A q_i$.*
- (v) *If $w' \preceq_{A_2} w$, then for every q'_i there exists a q_j with $q'_i \preceq_A q_j$.*

Proof of Fact C.9. We treat the claims in the order given.

- (i) We proceed by induction on the construction of \uparrow . In the base, where $q = q_1$, the claim holds vacuously. For the inductive step, we have $r \in Q$ such that $r \uparrow q_1 \cdots q_{n-1}$ and $q_n \in \gamma(q, \{\!\!| r \!\!\})$. In that case, $q_n \preceq_A q$ and $r \preceq_A q$ immediately. Also, we find by induction that for all $1 \leq i < n$ it holds that $q_i \preceq_A r$, and hence $q_i \preceq_A q$.

- (ii) We proceed by induction on the construction of δ_2 . In the base, we have $w' \in \delta_2(w, \mathbf{a})$ because $w' = q' \cdot x \cdot y$ and $w = q \cdot y$, and there exists an $r \in Q$ such that $q \uparrow r \cdot x$ and $q' \in \delta(r, \mathbf{a})$. By (i), we know that $q'_1 = q' \preceq_A r \preceq_A q = q_1$. Furthermore, if q'_i appears in x , then also by (i) we know that $q'_i \preceq_A q = q_1$. Lastly, if q'_i appears in y , then note that it also appears in w , and hence we can conclude by $q'_i \preceq_A q_j$ for some j . In the inductive step, $w = q \cdot x$ such that $q \in F$ and $w' \in \delta_2(x, \mathbf{a})$. The claim then follows immediately by induction.
- (iii) This proof proceeds analogously to the one above.
- (iv) We proceed by induction on the construction of γ_2 . In the base, we have that $w = q_1 \cdot x$ and $w' = q'_1 \cdot y \cdot x$ such that there exists an $p \in Q$ with $q_1 \uparrow p \cdot y$ and $q'_1 \in \gamma(p, \phi)$. Since A is fork-acyclic, it follows that $r \prec_A p$; because $p \preceq_A q_1$ by (i), the claim follows.
For the inductive step, we have that $\gamma_2(w, \phi) \neq \emptyset$ because $w = q \cdot x$ such that $q \in F$ and $\gamma_2(x, \phi) \neq \emptyset$. By induction, we then find a q_i such that $r \prec_A q_i$.
- (v) This can be shown by induction on \preceq_{A_2} , noting that the base cases are covered by (ii)–(iv). For the inductive step, it suffices to note that the claimed property is transitive in nature. \square

Now, if $w \in \phi \in \mathbb{M}(Q_2)$ and $x \in Q_2$ with $\gamma_2(x, \phi) \neq \emptyset$, we should show that $w \prec_{A_2} w$. First, note that $w = r \in Q$ for some $r \in Q$ by construction of γ_2 . By Fact C.9(iv), we know that $w = q_1 \cdots q_n$, and there exists a $1 \leq i \leq n$ with $r \prec_A q_i$. Suppose, towards a contradiction, that $w \preceq_{A_2} r$; then by Fact C.9(v) we also know that $q_i \preceq_A r$, contradicting that $r \prec_A q_i$.

To argue that A_2 is bounded, we first record the following.

Fact C.10. *Let $w = q_1 \cdots q_n \in Q_2$ and $x = q'_1 \cdots q_{n'} \in Q_2$. If $w \preceq_{A_2} x$, then for every $1 \leq i \leq n$ there exists a $1 \leq j \leq n'$ such that $n + D_A(q_i) - i \leq n' + D_A(q'_j) - j$.*

Proof of Fact C.10. It suffices to verify the claim for the pairs generating \preceq_{A_2} .

- If $w \in \delta_2(x, \mathbf{a})$ for some $\mathbf{a} \in \Sigma$, then we proceed by induction on the construction of δ_2 . In the base, there exist $r \in Q$ and $1 \leq k \leq n$ such that $q'_1 \uparrow r \cdot q_2 \cdots q_k$ and $q_1 \in \delta(r, \mathbf{a})$, while $n' + k - 1 = n$, and for $k < i \leq n'$ we have $q_i = q'_{i-k+1}$. We now consider two cases.
 - When $1 \leq i \leq k$, we choose $j = 1$; since $n \leq n'$ and $i \geq 1$ and $D_A(q_i) \leq D_A(q'_1)$ by Fact C.9(i), we find that $n + D_A(q_i) - i \leq n' + D_A(q'_1) - 1$.
 - Otherwise, when $k < i \leq n$, we choose $j = i - k + 1$ to find that $n + D_A(q_i) - i = n + D_A(q'_j) - i = n' + k - 1 - D_A(q'_j) - i = n' + D_A(q'_j) - j$.
In the inductive step, $q'_1 \in F$ and $w \in \delta_2(q'_2 \cdots q'_{n'}, \mathbf{a})$. The claim then follows by induction.
- If $w \in \gamma_2(x, \phi)$ for some $\phi \in \mathbb{M}(Q_2)$, then the proof is similar to the previous case.
- If there exists a $\phi \in \mathbb{M}(Q_2)$ such that $w \in \phi$ and $\gamma_2(x, \phi) \neq \emptyset$, then note that $\phi \in \mathbb{M}(Q)$ by construction of γ_2 , and thus that $w = r$ for some $r \in Q$. The proof proceeds by induction on γ_2 , where it suffices to show that $D_A(r) < D_A(q'_j)$ for some $1 \leq j \leq n'$. This is a direct consequence of Fact C.9(iv). \square

To see that A_2 is bounded, let $q_1 \cdots q_n \in Q_2$ and choose $m = \max_{1 \leq i \leq n} D_A(q_i)$. If $q'_1, \dots, q'_{n'} \in Q$ such that $q'_1 \cdots q'_{n'} \preceq_A q_1 \cdots q_n$, then by Fact C.10, we find $1 \leq j \leq n'$ such that $n' \leq n' + D_A(q'_1) - 1 \leq n + D_A(q_j) - j \leq n + m$. By Fact C.9(v), $q'_1, \dots, q'_{n'} \in \pi_A(q_1) \cup \cdots \cup \pi_A(q_n)$; the latter set is finite. Hence, the states supporting $q_1 \cdots q_n$ in A_2 are words of length at most $n + m$ over a finite alphabet; thus $\pi_{A_2}(q_1 \cdots q_n)$ must be finite.

To show that A_2 can accept the same languages as A , the following facts are useful.

Fact C.11. *If $q \in Q$ and $q \uparrow q_1 \cdots q_n$ such that for $1 \leq i \leq n$ there exist $q'_i \in Q$ and $U_i \in \text{SP}$ with $q_i \xrightarrow{U_i}_A q'_i$, and for $1 \leq i < n$ it holds that $q'_i \in F$, then $q \xrightarrow{U_1 \cdots U_n}_A q'_n$.*

Proof of Fact C.11. The proof proceeds by induction on the construction of \uparrow . In the base, we have that $q = q_1$ and $n = 1$; it then follows immediately that $q = q_1 \xrightarrow{U_1}_A q'_1 = q'_n$.

For the inductive step, we have $q \uparrow q_1 \cdots q_n$ because there exists $r \in Q$ with $r \uparrow q_1 \cdots q_{n-1}$ and $q_n \in \gamma(q, \{\uparrow\})$. By induction, we then know that $r \xrightarrow{U_1 \cdots U_{n-1}}_A q_{n-1}$; since $q_{n-1} \in F$ and $q_n \in \gamma(q, \{\uparrow\})$ it follows that $q \xrightarrow{U_1 \cdots U_{n-1}}_A q_n \xrightarrow{U_n}_A q'_n$, and hence $q \xrightarrow{U_1 \cdots U_n}_A q'_n$. \square

Fact C.12. *If $w \in Q_2$ and $w' \in F_2$ and $U \in \text{SP}$ such that $w \xrightarrow{U}_{A_2} w'$, then $w = q_1 \cdots q_n$ and $U = U_1 \cdots U_n$ such that for $1 \leq i \leq n$ there exists a $q'_i \in F$ with $q_i \xrightarrow{U_i}_A q'_i$.*

Proof of Fact C.12. We proceed by induction on the length ℓ of $w \xrightarrow{U}_{A_2} w'$. In the base, where $\ell = 0$, we have $U = 1$ and $w = w' \in F_2$. We choose for $1 \leq i \leq n$ that $U_i = 1$ and $q'_i = q_i \in F$.

For the inductive step, we have that $U = V \cdot W$ and $w'' \in Q_2$ such that $w \xrightarrow{V}_{A_2} w''$ is a unit run, and $w'' \xrightarrow{W}_{A_2} w'$ of length $\ell - 1$. By induction, $w'' = r_1 \cdots r_m$ and $W = W_1 \cdots W_m$ such that for $1 \leq i \leq m$ there exists an $r'_i \in F$ with $r_i \xrightarrow{W_i}_A r'_i$. Suppose that $w \xrightarrow{V}_{A_2} w''$ is a sequential unit run. Then $V = \mathbf{a}$ for some $\mathbf{a} \in \Sigma$, and $w'' \in \delta_2(w, \mathbf{a})$. We proceed by induction on δ_2 .

In the base, we have that $w = q_1 \cdots q_n$ and $w'' = q' \cdot v \cdot q_2 \cdots q_n$ such that there exists an $r \in Q$ with $q_1 \uparrow r \cdot v$ and $q' \in \delta(r, \mathbf{a})$. Note that $q' \cdot v = r_1 \cdots r_k$ for some $k \leq m$, and that $r_{i+k-1} = q_i$ for $2 \leq i \leq n$. We choose $U_1 = \mathbf{a} \cdot W_1 \cdots W_k$. Since $r \xrightarrow{\mathbf{a} \cdot W_1}_A r'_1 \in F$ and for $2 \leq i \leq k$ we have $r_i \xrightarrow{W_i}_A r'_i \in F$, it follows that $q_1 \xrightarrow{U_1}_A r'_k$ by Fact C.11; we set $q'_1 = r'_k$. For $i \geq 2$, we choose $q'_i = r'_{i+k-1}$ and $U_i = W_{i+k-1}$, to find that $r_{i+k-1} \xrightarrow{W_{i+k-1}}_A r'_{i+k-1}$, and hence $q_i \xrightarrow{U_i}_A q'_i$. Finally, we note that $U_1 \cdots U_n = \mathbf{a} \cdot W_1 \cdots W_k \cdot W_{k+1} \cdots W_m = V \cdot W = U$.

In the inductive step, $w = q_1 \cdots q_n$ and $w'' \in \delta_2(q_2 \cdots q_n, \mathbf{a})$, with $q_1 \in F$. By induction, $U = U_2 \cdots U_n$ where, for $2 \leq i \leq n$, $q'_i \in F$ with $q_i \xrightarrow{U_i}_A q'_i$. Here, $U_1 = 1$ and $q'_1 = q_1$ suffices.

The case where $w \xrightarrow{V}_{A_2} w''$ is a parallel unit run can be treated similarly. \square

Fact C.12 tells us that for $q \in Q$ we have that $L_{A_2}(q) \subseteq L_A(q)$. After all, if $q \xrightarrow{U}_{A_2} w$ for some $w \in F_2$, then we find $q' \in F$ such that $q \xrightarrow{U}_A q'$, and hence $U \in L_A(q)$.

For the converse inclusion, the following facts tell us how we can compose runs in A_2 .

Fact C.13. *If $w \xrightarrow{U}_{A_2} w'$ and $x \in Q_2$, then $w \cdot x \xrightarrow{U}_{A_2} w' \cdot x$.*

Proof of Fact C.13. We proceed by induction on the length ℓ of $w \xrightarrow{U}_{A_2} w'$. In the base, where $\ell = 0$, we have that $w = w'$ and $U = 1$. We then know that $w \cdot x = w' \cdot x$; hence $w \cdot x \xrightarrow{U}_{A_2} w' \cdot x$.

For the inductive step, let $\ell > 1$. We then find $w'' \in Q_2$ and $U = V \cdot W$ such that $w \xrightarrow{V}_{A_2} w''$ is a unit run, and $w'' \xrightarrow{W}_{A_2} w'$ is of length $\ell - 1$. Hence, $w'' \cdot x \xrightarrow{W}_{A_2} w' \cdot x$ by induction. If $w \xrightarrow{V}_{A_2} w''$ is a sequential unit run, then $V = \mathbf{a}$ for some $\mathbf{a} \in \Sigma$, and $w'' \in \delta_2(w, \mathbf{a})$. By construction of δ_2 , we have $w'' \cdot x \in \delta_2(w \cdot x, \mathbf{a})$, which means that $w \cdot x \xrightarrow{V}_{A_2} w'' \cdot x$. In total, we have $w \cdot x \xrightarrow{U}_{A_2} w' \cdot x$.

The case where $w \xrightarrow{V}_{A_2} w''$ is a parallel unit run can be treated similarly. \square

Fact C.14. *Let $q, r \in Q$ with $q' \in \gamma(q, \{r\})$, and let $r \xrightarrow{U}_{A_2} w$ be nontrivial. Then $q \xrightarrow{U}_{A_2} w \cdot q'$.*

Proof of Fact C.14. Since $r \xrightarrow{U}_{A_2} w$ is nontrivial, $U = V \cdot W$ and $x \in Q_2$ such that $r \xrightarrow{V}_{A_2} x$ is a unit run, and $x \xrightarrow{W}_{A_2} w$. If $r \xrightarrow{V}_{A_2} x$ is a sequential unit run, then $V = \mathbf{a}$ for some $\mathbf{a} \in \Sigma$, and $x \in \delta_2(r, \mathbf{a})$. By construction of δ_2 , we obtain $q'', r \in Q$ and $y \in Q_2$ such that $x = q'' \cdot y$ and $r \uparrow r' \cdot y$ as well as $q'' \in \delta(r', \mathbf{a})$. By definition of \uparrow , also $q \uparrow r' \cdot y \cdot q'$, and thus $x \cdot q' = q'' \cdot y \cdot q' \in \delta_2(q, \mathbf{a})$. We then find that $q \xrightarrow{V}_{A_2} x \cdot q'$. Since $x \cdot q' \xrightarrow{W}_{A_2} w \cdot q'$ by Fact C.13, we conclude that $q \xrightarrow{U}_{A_2} w \cdot q'$.

The case where $r \xrightarrow{V}_{A_2}$ is a parallel unit run can be argued similarly. \square

Fact C.15. *Let $w \xrightarrow{U}_{A_2} w'$ be nontrivial, and $x \in F_2$. Then $x \cdot w \xrightarrow{U}_{A_2} w'$.*

Proof of Fact C.15. Since $w \xrightarrow{U}_{A_2} w'$ is non-trivial, we find that $U = V \cdot W$ and $w'' \in Q_2$ such that $w \xrightarrow{V}_{A_2} w''$ is a unit run, and $w'' \xrightarrow{W}_{A_2} w'$. If $w \xrightarrow{V}_{A_2} w''$ is a sequential unit run, then $V = \mathbf{a}$ for some $\mathbf{a} \in \Sigma$, and $w'' \in \delta_2(w, \mathbf{a})$. A simple inductive argument on the length of x then tells us that $w'' \in \delta_2(x \cdot w, \mathbf{a})$ as well. From this, it follows that $x \cdot w \xrightarrow{V}_{A_2} w''$, and thus $x \cdot w \xrightarrow{U}_{A_2} w'$.

The case where $w \xrightarrow{V}_{A_2}$ is a parallel unit run can be argued similarly. \square

Finally, we can use the above to show that A_2 can simulate the unary forks of A .

Fact C.16. *If $q \xrightarrow{U}_A q'$, then there exists $x \in F_2$ such that $q \xrightarrow{U}_{A_2} x \cdot q'$.*

Proof of Fact C.16. We proceed by induction on $q \xrightarrow{U}_A q'$. If $q \xrightarrow{U}_A q'$ is trivial, then the claim is satisfied by choosing $x = 1$. Otherwise, suppose $q \xrightarrow{U}_A q'$ is a sequential unit run, i.e., $U = \mathbf{a}$ for some $\mathbf{a} \in \Sigma$, and $q' \in \delta(q, \mathbf{a})$. Since $q \uparrow q$, we then have that $q' \in \delta_2(q, \mathbf{a})$, and hence $q \xrightarrow{U}_{A_2} q'$.

For the inductive step, there are again two cases to consider.

- Suppose $q \xrightarrow{U}_A q'$ because $U = V \cdot W$ and there exists $q'' \in Q$ such that $q \xrightarrow{V}_A q''$ and $q'' \xrightarrow{W}_A q'$. By induction, we obtain $x'', x' \in F_2$ such that $q \xrightarrow{V}_{A_2} x'' \cdot q''$ and $q'' \xrightarrow{W}_{A_2} x' \cdot q'$. Without loss of generality, $q'' \xrightarrow{W}_A q'$ is non-trivial, and hence neither is $q'' \xrightarrow{W}_{A_2} x' \cdot q'$. By Fact C.15, we find that $x'' \cdot q'' \xrightarrow{W}_{A_2} x' \cdot q'$. In total, we find that $q \xrightarrow{U}_{A_2} x \cdot q'$.
- Suppose $q \xrightarrow{U}_A q'$ because $U = U_1 \parallel \dots \parallel U_n$, and there exist $r_1, \dots, r_n \in Q$ and $r'_1, \dots, r'_n \in F$ such that for $1 \leq i \leq n$ we have $r_i \xrightarrow{U_i}_A r'_i$, and $q' \in \gamma(q, \{r_1, \dots, r_n\})$. There are two subcases.
 - If $n = 1$, then by induction we find $x_1 \in F_2$ such that $r_1 \xrightarrow{U_1}_{A_2} x'_1 \cdot r'_1$. By Fact C.14, we then find $q \xrightarrow{U}_{A_2} x'_1 \cdot r'_1 \cdot q'$. Choosing $x' = x'_1 \cdot r'_1$ satisfies the claim.
 - If $n \geq 2$, then $q' \in \gamma_2(q, \{r_1, \dots, r_n\})$. By induction, we find for $1 \leq i \leq n$ an $x_i \in F_2$ with $r_i \xrightarrow{U_i}_{A_2} x'_i \cdot r'_i \in F_2$. Thus $q \xrightarrow{U}_{A_2} q'$; choosing $x' = 1$ satisfies the claim. \square

The above allows us to prove that, for $q \in Q$, we have $L_A(q) \subseteq L_{A_2}(q)$. To this end, suppose $U \in L_A(q)$; then there exists a $q' \in F$ with $q \xrightarrow{U}_A q'$. By Fact C.16 we find $x' \in F_2$ with $q \xrightarrow{U}_{A_2} x \cdot q' \in F_2$, and hence $U \in L_{A_2}(q)$. Since $L_{A_2}(q) \subseteq L_A(q)$, it follows that $L_A(q) = L_{A_2}(q)$.

Note that it is 2-forking by construction. For parsimony, observe that if $w \in \phi \in \mathbb{M}(Q_2)$ and $x \in Q_2$ such that $\gamma_2(x, \phi)$, then $\phi \in \mathbb{M}(Q)$ by definition of γ_2 , and hence $w = q$ for

some $q \in Q$. A simple inductive argument then tells us that there exists an $r \in Q$ such that $\gamma(r, \phi) \neq \emptyset$. Since A is parsimonious, we know that $1 \notin L_A(q)$; since $L_A(q) = L_{A_2}(q)$, it follows that $1 \notin L_{A_2}(q)$. \square

C.4. Ensuring flat-branching.

Lemma 6.24. A_3 is bounded, 2-forking, parsimonious, flat-branching (i.e., well-structured) and weakly implements A .

Proof. The proof of this statement consists of several steps. In the sequel, we will write $Q^{\mathbf{P}} = \{q^{\mathbf{P}} : q \in Q\}$ and $Q^{\mathbf{S}} = \{q^{\mathbf{S}} : q \in Q\}$. We start by making the following observations:

Fact C.17. *The following hold for all $\psi \in \mathbb{M}(Q)$:*

- (i) *If $\phi \blacktriangleleft \psi$, then $|\phi| \leq |\psi|$.*
- (ii) *If $p \in Q_3$ and $\phi \in \mathbb{M}(Q_3)$ such that $\gamma_3(p, \phi) \neq \emptyset$, then $\phi \in \mathbb{M}(Q^{\mathbf{S}})$.*
- (iii) *If $p \in Q$ and $\phi \in \mathbb{M}(Q_3)$, then $\gamma_3(p^{\mathbf{S}}, \phi) \subseteq Q^{\mathbf{S}} \cup Q^{\mathbf{P}}$.*

Proof. We treat the claims in the order given.

- (i) This claim is proved by induction on \blacktriangleleft . In the base, $\phi \blacktriangleleft \psi$ because $\phi = \psi$, and so the claim holds immediately. For the inductive step, we have that $\phi \blacktriangleleft \psi$ because $\psi = \psi_1 \sqcup \psi_2$, such that $\phi \blacktriangleleft \psi_1$, and $\gamma(p, \psi_2) \cap F \neq \emptyset$ with $\phi \blacktriangleleft \psi_1 \sqcup \{p\}$. By induction, we have that $|\phi| \leq |\psi_1| + 1$. Since A is 1-forking, we have that $|\psi_2| \geq 1$; hence, we conclude that $|\phi| \leq |\psi_1| + |\psi_2| = |\psi|$.
- (ii) If $p \in Q_3$ and $\phi \in \mathbb{M}(Q_3)$ such that $\gamma_3(p, \phi) \neq \emptyset$, then by definition of γ_3 we have that $\phi = \psi^{\mathbf{S}}$ for some $\psi \in \mathbb{M}(Q)$. Hence, $\phi \in \mathbb{M}(Q^{\mathbf{S}})$.
- (iii) Suppose that $p \in Q$ and $\phi \in \mathbb{M}(Q_3)$, and let $q \in \gamma_3(p^{\mathbf{S}}, \phi)$. By definition of γ_3 , we have that $q \in \{r^{\mathbf{S}}, r^{\mathbf{P}}\}$ such that $r \in \gamma(p, \psi)$ for some $\psi \in \mathbb{M}(Q)$. Hence, $q \in Q^{\mathbf{S}} \cup Q^{\mathbf{P}}$. \square

We are now set to prove that A_3 indeed satisfies the right properties.

Fact C.18. A_3 is 2-forking, parsimonious, and flat-branching.

Proof. For 2-forking, suppose $p \in Q_3$ and $\phi \in \mathbb{M}(Q_3)$ such that $\gamma_3(p, \phi) \neq \emptyset$. Then by definition of γ_3 we find $\chi, \psi \in \mathbb{M}(Q)$ and $r \in Q$, such that $\phi = \chi^{\mathbf{S}}$, $\psi \blacktriangleleft \chi$, and $\gamma(r, \psi) \neq \emptyset$. Since A is 2-forking, we can conclude by Fact C.17(i) that $2 \leq |\psi| \leq |\chi| = |\phi|$.

For parsimony, suppose $\gamma_3(p, \phi) \neq \emptyset$ and $q \in \phi$. Then by Fact C.17(ii) we know that $q \in Q^{\mathbf{S}}$, so $q \notin F_3$. Since A_3 is 1-forking, $1 \notin L_{A_3}(q)$ by Fact C.2, hence A_3 is parsimonious.

For flat-branching, suppose $p \in Q_3$ is a fork target. Then by Fact C.17(ii), we know that $p \in Q^{\mathbf{S}}$. By Fact C.17(iii), we can then conclude that $\gamma_3(p, \psi) \cap F_3 \subseteq (Q^{\mathbf{P}} \cup Q^{\mathbf{S}}) \cap F_3 = \emptyset$. \square

We can now relate the runs of A_3 to those in A as follows.

Fact C.19. *If $p \xrightarrow{U}_{A_3} q$, then the following hold:*

- (i) *If $p = p'^{\mathbf{S}}$ and $q \in \{q'^{\mathbf{S}}, q'^{\mathbf{P}}\}$, then $p' \xrightarrow{U}_A q'$.*
- (ii) *If $p \in \{p'^{\mathbf{S}}, p'^{\mathbf{P}}\}$ and $q = \top$, then there exists a $q' \in F$ with $p' \xrightarrow{U}_A q'$.*

Proof. We proceed by induction on $p \xrightarrow{U}_{A_3} q$. In the base, the case where $p \xrightarrow{U}_{A_3} q$ is trivial holds vacuously. Otherwise, if $p \xrightarrow{U}_{A_3} q$ because $U = \mathbf{a}$ for some $\mathbf{a} \in \Sigma$ and $q \in \delta_3(p, \mathbf{a})$, then we know that $p \notin Q^{\mathbf{P}}$ by definition of δ_3 . Therefore, assume $p = p'^{\mathbf{S}}$ for some $p' \in Q$. There are two cases.

- If $q \in \{q'^{\text{S}}, q'^{\text{P}}\}$, then it must be the case that $q' \in \delta(p', \mathbf{a})$, so $p' \xrightarrow{\mathbf{a}}_A q'$.
- If $q = \top$, then $\delta(p', \mathbf{a}) \cap F \neq \emptyset$. Choose $q' \in \delta(p', \mathbf{a}) \cap F$ to find that so $p' \xrightarrow{\mathbf{a}}_A q'$.

For the inductive step, there are again two cases.

- Suppose $p \xrightarrow{U}_{A_3} q$ because $U = V \cdot W$ and there exists an $r \in Q_3$ such that $p \xrightarrow{V}_{A_3} r$ and $r \xrightarrow{W}_{A_3} q$. Furthermore, we may assume w.l.o.g. that neither of these runs is trivial. Since \top does not permit nontrivial runs, we have $r \in Q^{\text{P}} \cup Q^{\text{S}}$. Furthermore, $p \notin Q^{\text{P}}$, because if $p \in Q^{\text{P}}$ then $r = \top$; we set $p = p'^{\text{S}}$. We do a case analysis on r .
 - If $r = r'^{\text{P}}$, then necessarily $q = \top$. By the induction hypothesis we get that $p' \xrightarrow{V}_A r'$ and we find $q' \in F$ such that $r' \xrightarrow{W}_A q'$, hence $p' \xrightarrow{U}_A q'$.
 - If $r = r'^{\text{S}}$, then by induction we get that $p' \xrightarrow{V \cdot W}_A r'$. We now look at q :
 - * If $q \in \{q'^{\text{S}}, q'^{\text{P}}\}$, then by induction we have $r' \xrightarrow{W}_A q'$, so $p' \xrightarrow{V \cdot W}_A q'$.
 - * If $q = \top$, then induction gives us $q' \in F$ such that $r' \xrightarrow{W}_A q'$, and thus $p' \xrightarrow{V \cdot W}_A q'$.
- Suppose $p \xrightarrow{U}_{A_3} q$ because $q_1, \dots, q_n \in Q_3$ with $q \in \gamma_3(p, \{\!\!\{q_1, \dots, q_n\}\!\!\})$, and $U = U_1 \parallel \dots \parallel U_n$ such that for $1 \leq i \leq n$ we have $q_i \xrightarrow{U_i}_{A_3} \top$. Since A_3 is 1-forking and parsimonious, each U_i is non-empty by Fact C.2. By definition of γ_3 , for each $1 \leq i \leq n$ there exists a $q'_i \in Q$ such that $q_i = q_i'^{\text{S}}$. By induction, we obtain for each $1 \leq i \leq n$ a $q''_i \in F$ such that $q'_i \xrightarrow{U_i}_A q''_i$.

On the one hand, suppose $p = p'^{\text{S}}$ for $p' \in Q$. In that case, $q \in \{q'^{\text{S}}, q'^{\text{P}}\}$ for some $q' \in \gamma(p', \phi)$, with $\phi \triangleleft \{\!\!\{q_1, \dots, q_n\}\!\!\}$. We show $p' \xrightarrow{U}_A q'$ by induction on \triangleleft . In the base, $\phi = \{\!\!\{q_1, \dots, q_n\}\!\!\}$, and so the claim follows. In the inductive step, $\{\!\!\{q_1, \dots, q_n\}\!\!\} = \{\!\!\{q_1, \dots, q_k\}\!\!\} \sqcup \{\!\!\{q_{k+1}, \dots, q_n\}\!\!\}$ and there exists an $r \in Q$ such that $\gamma(r, \{\!\!\{q_{k+1}, \dots, q_n\}\!\!\}) \cap F \neq \emptyset$ and $\phi \triangleleft \{\!\!\{q_1, \dots, q_k\}\!\!\} \sqcup \{\!\!\{r\}\!\!\}$. In that case, $r \xrightarrow{U_{k+1} \parallel \dots \parallel U_n}_A r'$ for some $r' \in F$; hence, by induction, $p' \xrightarrow{U_1 \parallel \dots \parallel U_n}_A q'$.

On the other hand, if $p = p'^{\text{P}}$ for some $p' \in Q$, then $q' = \top$ by definition of γ_3 . Furthermore, there exists $q' \in \gamma(p', \phi) \cap F$ for some $\phi \in \mathbb{M}(Q)$ with $\phi \triangleleft \{\!\!\{q_1, \dots, q_n\}\!\!\}$. A similar inductive argument to the previous case then shows that $p' \xrightarrow{U}_A q'$. \square

Fact C.20. *If $p \xrightarrow{U}_A q$ is nontrivial, then $p^{\text{S}} \xrightarrow{U}_{A_3} q^{\text{S}}$ and $p^{\text{S}} \xrightarrow{U}_{A_3} q^{\text{P}}$.*

Furthermore, if $q \in F$, then either $p^{\text{S}} \xrightarrow{U}_{A_3} \top$ or $p^{\text{P}} \xrightarrow{U}_{A_3} \top$.

Proof. We proceed by induction on $p \xrightarrow{U}_A q$. In the base, $p \xrightarrow{U}_A q$ because $U = \mathbf{a}$ for some $\mathbf{a} \in \Sigma$, and $q \in \delta(p, \mathbf{a})$. Thus $q^{\text{S}}, q^{\text{P}} \in \delta_3(p^{\text{S}}, \mathbf{a})$, and hence $p^{\text{S}} \xrightarrow{U}_{A_3} q^{\text{S}}$ and $p^{\text{S}} \xrightarrow{U}_{A_3} q^{\text{P}}$. Furthermore, $q \in F$, then $\top \in \delta_3(p^{\text{S}}, \mathbf{a})$, and hence $p^{\text{S}} \xrightarrow{U}_{A_3} \top$. In the inductive step, there are two cases.

- If $p \xrightarrow{U}_A q$ because $U = V \cdot W$ and there exists an $r \in Q$ with $p \xrightarrow{V}_A r$ and $r \xrightarrow{W}_A q$, then we can assume without loss of generality that neither of these runs is trivial. By induction, we then find that $p^{\text{S}} \xrightarrow{V}_{A_3} r^{\text{S}}$ and $p^{\text{S}} \xrightarrow{W}_{A_3} r^{\text{S}}$, as well as $r^{\text{S}} \xrightarrow{W}_{A_3} q^{\text{S}}$ and $r^{\text{S}} \xrightarrow{W}_{A_3} q^{\text{P}}$. Putting this together, we have that $p^{\text{S}} \xrightarrow{U}_{A_3} q^{\text{S}}$ and $p^{\text{S}} \xrightarrow{U}_{A_3} q^{\text{P}}$.

Furthermore, if $q = \top$, then it suffices to prove that $r^{\text{S}} \xrightarrow{W}_{A_3} \top$ or $r^{\text{P}} \xrightarrow{W}_{A_3} \top$, which we obtain from $r \xrightarrow{W}_A q$ by induction.

- Suppose $p \xrightarrow{U}_A q$ because there exist $q_1, \dots, q_n \in Q$ such that $q \in \gamma(p, \{\!\!\{q_1, \dots, q_n\}\!\!\})$, and $U = U_1 \parallel \dots \parallel U_n$ such that for $1 \leq i \leq n$ there exists a $q'_i \in F$ with $q_i \xrightarrow{U_i}_A q'_i$. Since A is parsimonious, we can assume without loss of generality that none of these runs is trivial.

We then claim that, for $1 \leq i \leq n$, there exists a $\phi_i = \{\{q_{i,1}, \dots, q_{i,n_i}\} \in \mathbb{M}(Q)$ such that $\{\{q_i\} \blacktriangleleft \phi_i$, and $U_i = U_{i,1} \parallel \dots \parallel U_{i,n_i}$, such that for $1 \leq i \leq n_i$ we have that $q_{i,j}^s \xrightarrow{U_{i,j}}_{A_3} \top$. Applying the induction hypothesis to each $q_i \xrightarrow{U_i}_A q'_i \in F$, there are two cases to consider:

- If $q_i^s \xrightarrow{U_i}_{A_3} \top$, then we choose $n_i = 1$ and $q_{i,1} = q_i$ and $U_{i,1} = U_i$.
- If $q_i^p \xrightarrow{U_i}_{A_3} \top$, then by construction of A_3 this must be a parallel unit run. Consequently, there exist $q_{i,1}^s, \dots, q_{i,n_i}^s \in Q^s$ with $\top \in \gamma_3(q_i^p, \{\{q_{i,1}, \dots, q_{i,n_i}\})$, and $U_i = U_{i,1} \parallel \dots \parallel U_{i,n_i}$ such that for $1 \leq j \leq n_i$ we have that $q_{i,j}^s \xrightarrow{U_{i,j}}_{A_3} \top$. By definition of γ_3 , we then obtain $\psi \in \mathbb{M}(Q)$ such that $\gamma(q_i, \psi) \cap F \neq \emptyset$ and $\psi \blacktriangleleft \{\{q_{i,1}, \dots, q_{i,n_i}\}$. A straightforward inductive argument on the definition of \blacktriangleleft shows that it is transitive; hence, since $\{\{q_i\} \blacktriangleleft \psi$, we have that $\{\{q_i\} \blacktriangleleft \{\{q_{i,1}, \dots, q_{i,n_i}\}$.

Using the above, it follows that $\{\{q_1, \dots, q_n\} \blacktriangleleft \{\{q_{1,1}, \dots, q_{n,n_n}\}$. Hence,

$$q^s, q^p \in \gamma_3(p^s, \{\{q_{1,1}^s, \dots, q_{n,n_n}^s\})$$

Since $U = U_{1,1} \parallel \dots \parallel U_{n,n}$, it follows that $p^s \xrightarrow{U}_{A_3} q^s$ and $p^s \xrightarrow{U}_{A_3} q^p$.

Furthermore, if $q \in F$, then $\top \in \gamma_3(p^p, \{\{q_{1,1}^s, \dots, q_{n,n_n}^s\})$, and hence $p^p \xrightarrow{U}_{A_3} \top$. \square

We are now ready to show that our construction preserves languages. More specifically, Facts C.19 and C.20 together imply that for $p \in Q$, we have

$$L_{A_1}(q) = L_{A_3}(q^p) \cup L_{A_3}(q^s) \cup \begin{cases} L_{A_3}(\top) & q \in F \\ \emptyset & \text{otherwise} \end{cases}$$

To see that our construction preserves fork-acyclicity and that A_3 is bounded, one can show that if $p, q \in Q$ are such that $p^s \preceq_{A_3} q^s$, $p^s \preceq_{A_3} q^p$, $p^p \preceq_{A_3} q^s$ or $p^p \preceq_{A_3} q^p$, then $p \preceq_{A_3} q$. A fork cycle in A_3 thus gives rise to a fork cycle in A , which means that if A is fork-acyclic, then so is A_3 . Furthermore, the support of a state q^s or q^p in A_3 is contained in $\{p^s, p^p : p \in \pi_A(q)\} \cup \{\top\}$; since the latter is finite as A is bounded, it follows that A_3 must also be bounded. \square

APPENDIX D. LEMMAS ABOUT THE SYNTACTIC POMSET AUTOMATON

D.1. Deconstruction lemmas.

Lemma 7.5. *Let $e_1, e_2 \in \mathcal{T}$, $f \in \mathcal{F}$ and $U \in \text{SP}(\Sigma)$, with $e_1 \cdot e_2 \xrightarrow{U}_{\Sigma} f$ (of length ℓ). Then $U = U_1 \cdot U_2$ and there exist $f_1, f_2 \in \mathcal{F}$ with $e_i \xrightarrow{U_i}_{\Sigma} f_i$ for $i \in \{1, 2\}$ (of length at most ℓ).*

Proof. We proceed by induction on the length ℓ of $e_1 \cdot e_2 \xrightarrow{U}_{\Sigma} f$. In the base, where $\ell = 0$, we have $f = e_1 \cdot e_2$ (hence $e_1, e_2 \in \mathcal{F}$) and $U = 1$. We can then choose $f_1 = e_1$ and $f_2 = e_2$ as well as $U_1 = U_2 = 1$, to find that $e_1 \xrightarrow{U_1}_{\Sigma} f_1$ and $e_2 \xrightarrow{U_2}_{\Sigma} f_2$, of length zero.

For the inductive step, let $e_1 \cdot e_2 \xrightarrow{U}_{\Sigma} f$ be of length $\ell + 1$. We find that $U = U_0 \cdot U'$, and a $g \in \mathcal{T}$ such that $e_1 \cdot e_2 \xrightarrow{U_0}_{\Sigma} g$ is a unit run, and $g \xrightarrow{U'}_{\Sigma} f$ is of length ℓ . If $e_1 \cdot e_2 \xrightarrow{U_0}_{\Sigma} g$ is a sequential unit run, then $U_0 = \mathbf{a}$ for some $\mathbf{a} \in \Sigma$, and $g \in \delta_{\Sigma}(e_1 \cdot e_2, \mathbf{a}) = \delta_{\Sigma}(e_1, \mathbf{a}) \ddagger e_2 \cup e_1 \star \delta_{\Sigma}(e_2, \mathbf{a})$. This gives us two cases to consider.

- If $g \in \delta_\Sigma(e_1, \mathbf{a}) \ddagger e_2$, then $g = g_1 \cdot e_2$ such that $g_1 \in \delta_\Sigma(e_1, \mathbf{a})$. By induction we find $f_1, f_2 \in \mathcal{F}$ and $U' = U'_1 \cdot U'_2$ such that $g_1 \xrightarrow{U'_1}_\Sigma f_1$, and $e_2 \xrightarrow{U'_2}_\Sigma f_2$, of length at most ℓ . We choose $U_1 = U_0 \cdot U'_1$ and $U_2 = U'_2$ to find $U = U_0 \cdot U' = U_0 \cdot U'_1 \cdot U'_2 = U_1 \cdot U_2$, as well as $e_1 \xrightarrow{U_1}_\Sigma f_1$ of length at most $\ell + 1$, and $e_2 \xrightarrow{U_2}_\Sigma f_2$ of length at most ℓ .
- If $g \in e_1 \star \delta_\Sigma(e_2, \mathbf{a})$, then first note that $e_1 \in \mathcal{F}$, and $g \in \delta_\Sigma(e_2, \mathbf{a})$. We choose $U_1 = 1$ and $U_2 = U$ as well as $f_1 = e_1$ and $f_2 = f'$ to find that $U = U_1 \cdot U_2$ as well as $e_1 \xrightarrow{U_1}_\Sigma f_1$ of length zero. Lastly, $e_2 \xrightarrow{U_0}_\Sigma g \xrightarrow{U'}_\Sigma f' = f_2$, meaning $e_2 \xrightarrow{U}_\Sigma f_2$ of length at most $\ell + 1$.

The case where $e_1 \cdot e_2 \xrightarrow{U_0}_\Sigma g$ is a parallel unit run can be treated similarly. \square

Lemma 7.6. *Let $e_1, e_2 \in \mathcal{T}$, $f \in \mathcal{F}$ and $U \in \text{SP}(\Sigma)$ such that $e_1 \parallel e_2 \xrightarrow{U}_\Sigma f$. Then $U = U_1 \parallel U_2$, and there exist $f_1, f_2 \in \mathcal{F}$ such that $e_i \xrightarrow{U_i}_\Sigma f_i$ for $i \in \{1, 2\}$.*

Proof. If $e_1 \parallel e_2 \xrightarrow{U}_\Sigma f$ is trivial, then $U = 1$ and $e_1 \parallel e_2 = f \in \mathcal{F}$. Hence, $e_1, e_2 \in \mathcal{F}$; we can choose $f_1 = e_1$, $f_2 = e_2$ and $U_1 = U_2 = 1$ to satisfy the claim.

Otherwise, there exist $U_0, U' \in \text{SP}(\Sigma)$ and $g \in \mathcal{T}$ such that $U = V \cdot W$ and $e_1 \parallel e_2 \xrightarrow{V}_\Sigma g$ is a unit run, and $g \xrightarrow{W}_\Sigma f$. We can discount the possibility that $e_1 \parallel e_2 \xrightarrow{V}_\Sigma g$ is a sequential unit run, because $\delta_\Sigma(e_1 \parallel e_2, \mathbf{a}) = \emptyset$ for all $\mathbf{a} \in \Sigma$. Hence, $e_1 \parallel e_2 \xrightarrow{V}_\Sigma g$ is a parallel unit run, meaning that $V = V_1 \parallel \dots \parallel V_n$ and there exists a $\phi = \{h_1, \dots, h_n\} \in \mathbb{M}(\mathcal{T})$ such that $g \in \gamma_\Sigma(e_1 \parallel e_2, \phi)$, and for $1 \leq i \leq n$ there exists an $h'_i \in \mathcal{F}$ with $h_i \xrightarrow{V_i}_\Sigma h'_i$. By definition of γ_Σ , it then follows that $n = 2$ and $g = 1$ as well as (without loss of generality) $e_1 = h_1$ and $e_2 = h_2$. Since $g = 1$, it must be that $g \xrightarrow{W}_\Sigma f$ is trivial, and hence $W = 1$, meaning that $U = V$. We choose $f_1 = h'_1$, $f_2 = h'_2$, $U_1 = V_1$ and $U_2 = V_2$ to satisfy the claim. \square

D.2. Construction lemmas.

Lemma 7.9. *Let $e_1, e_2 \in \mathcal{T}$, $f_1, f_2 \in \mathcal{F}$ and $U, V \in \text{SP}(\Sigma)$ such that $e_1 \xrightarrow{U}_\Sigma f_1$ and $e_2 \xrightarrow{V}_\Sigma f_2$. There exists an $f \in \mathcal{F}$ with $e_1 \cdot e_2 \xrightarrow{U \cdot V}_\Sigma f$.*

Proof. The proof consists of two phases; first, we verify the following.

Fact D.1. *We have that $e_1 \cdot e_2 \xrightarrow{U}_\Sigma f_1 \cdot e_2$.*

Proof. The proof proceeds by induction on the length ℓ of $e_1 \xrightarrow{U}_\Sigma f_1$. In the base, where $\ell = 0$ and $f_1 = e_1$ as well as $U = 1$, we the claim holds immediately.

In the inductive step, let $e_1 \xrightarrow{U}_\Sigma f_1$ be of length $\ell + 1$. We find $e'_1 \in \mathcal{T}$ and $U = U_0 \cdot U'$ such that $e_1 \xrightarrow{U_0}_\Sigma e'_1$ is a unit run, and $e'_1 \xrightarrow{U'}_\Sigma f_1$ is of length ℓ . By induction, $e'_1 \cdot e_2 \xrightarrow{U'}_\Sigma f_1 \cdot e_2$. If $e_1 \xrightarrow{U_0}_\Sigma e'_1$ is a sequential unit run, then $U_0 = \mathbf{a}$ for some $\mathbf{a} \in \Sigma$, and $e'_1 \in \delta_\Sigma(e_1, \mathbf{a})$, meaning $e'_1 \cdot e_2 \in \delta_\Sigma(e_1 \cdot e_2, \mathbf{a})$, hence $e_1 \cdot e_2 \xrightarrow{U_0}_\Sigma e'_1 \cdot e_2$. We conclude that $e_1 \cdot e_2 \xrightarrow{U}_\Sigma f_1 \cdot e_2$.

The case where $e_1 \xrightarrow{U_0}_\Sigma e'_1$ is a parallel unit run is similar. \square

Next, we note the following.

Fact D.2. *There exists an $f \in \mathcal{F}$ such that $f_1 \cdot e_2 \xrightarrow{V}_\Sigma f$.*

Proof. There are two cases to consider, based on the length of $e_2 \xrightarrow{V}_\Sigma f_2$.

- If $\ell = 0$, we know that $f_2 = e_2$ and $V = 1$. We can then choose $f = f_1 \cdot e_2$.

- In the inductive step, let $e_2 \xrightarrow{V} f_1$ be of length $\ell + 1$. We find $e'_2 \in \mathcal{T}$ and $V = V_0 \cdot V'$ such that $e_2 \xrightarrow{V_0} e'_2$ is a unit run, and $e'_2 \xrightarrow{V'} f_2$ is of length ℓ . If $e_2 \xrightarrow{V_0} e'_2$ is a sequential unit run, then $V_0 = \mathbf{a}$ for some $\mathbf{a} \in \Sigma$, and $e'_2 \in \delta_\Sigma(e_2, \mathbf{a})$, and thus $e'_2 \in \delta_\Sigma(f_1 \cdot e_2, \mathbf{a})$. Hence, $f_1 \cdot e_2 \xrightarrow{V_0} e'_2$, meaning $f_1 \cdot e_2 \xrightarrow{V} f_2$; choosing $f = f_2$ satisfies the claim.

The case where $e_2 \xrightarrow{V_0} e'_2$ is a parallel unit run is similar. \square

Putting these together, we find $f \in \mathcal{F}$ such that $e_1 \cdot e_2 \xrightarrow{U \cdot V} f$. \square

Lemma 7.10. *Let $e_1, e_2 \in \mathcal{T}$, $f_1, f_2 \in \mathcal{F}$, and $U, V \in \text{SP}(\Sigma)$ such that $e_1 \xrightarrow{U} f_1$ and $e_2 \xrightarrow{V} f_2$. Then $e_1 \parallel e_2 \xrightarrow{U \parallel V} 1$.*

Proof. Since $1 \in \gamma_\Sigma(e_1 \parallel e_2, \{e_1, e_2\})$, the claim follows immediately. \square

D.3. Correctness of the syntactic PA.

Lemma 7.12. *Let $e, f \in \mathcal{T}$. The following hold:*

$$\begin{aligned} L_\Sigma(0) &= \emptyset & L_\Sigma(e + f) &= L_\Sigma(e) \cup L_\Sigma(f) & L_\Sigma(e^*) &= L_\Sigma(e)^* \\ L_\Sigma(1) &= \{1\} & L_\Sigma(e \cdot f) &= L_\Sigma(e) \cdot L_\Sigma(f) \\ L_\Sigma(\mathbf{a}) &= \{\mathbf{a}\} & L_\Sigma(e \parallel f) &= L_\Sigma(e) \parallel L_\Sigma(f) \end{aligned}$$

Proof. We treat the claims case-by-case.

- To show $L_\Sigma(0) = \emptyset$, suppose that $U \in L_\Sigma(0)$. In that case, $0 \xrightarrow{U} e$ for some $e \in \mathcal{F}$. Since $0 \notin \mathcal{F}$, this means that $0 \xrightarrow{U} e$ cannot be trivial. In that case, there exists an $e' \in \mathcal{T}$ such that $0 \xrightarrow{U} e'$ is a unit run. However, this contradicts that $\delta_\Sigma(0, \mathbf{a}) = \emptyset$ for all $\mathbf{a} \in \Sigma$, and $\gamma_\Sigma(e, \phi) = \emptyset$ for all $\phi \in \mathbb{M}(\mathcal{T})$. Therefore, our assumption that $U \in L_\Sigma(0)$ must be false. We conclude that $L_\Sigma(0) = \emptyset$.
- To show $L_\Sigma(1) = \{1\}$, suppose that $U \in L_\Sigma(1)$, i.e., $1 \xrightarrow{U} e$ for some $e \in \mathcal{F}$. By an argument similar to the previous case, we can argue that $1 \xrightarrow{U} e$ is trivial, and hence $U = 1$. The other inclusion follows from the fact that $1 \in \mathcal{F}$ and $1 \xrightarrow{1} 1$.
- To show $L_\Sigma(\mathbf{a}) = \{\mathbf{a}\}$, suppose that $U \in L_\Sigma(\mathbf{a})$, i.e., $\mathbf{a} \xrightarrow{U} e$ for some $e \in \mathcal{F}$. Since $\mathbf{a} \notin \mathcal{F}$, we know $\mathbf{a} \xrightarrow{U} e$ must be non-trivial. This means that we can write $U = U_0 \cdot U'$, and there exists an $f \in \mathcal{T}$ such that $\mathbf{a} \xrightarrow{U_0} f$ is a unit run, and $f \xrightarrow{U'} e$. A quick glance at δ_Σ and γ_Σ then tells us that $f = 1$. By the previous case, we know that $f \xrightarrow{U'} e$ must be trivial; hence $U' = 1$ and $f = e$. Indeed, $\mathbf{a} \xrightarrow{U_0} f$ must be a sequential unit run, for $\gamma_\Sigma(\mathbf{a}, \phi) = \emptyset$ for all $\phi \in \mathbb{M}(\mathcal{T})$. This tells us that $U = \mathbf{b}$ and $f \in \delta_\Sigma(\mathbf{a}, \mathbf{b})$ for some $\mathbf{b} \in \Sigma$; by definition of δ_Σ , it follows that $\mathbf{b} = \mathbf{a}$.

For the other inclusion, let $U = \mathbf{a}$; then $\mathbf{a} \xrightarrow{\mathbf{a}} 1$ immediately, and hence $\mathbf{a} \in L_\Sigma(\mathbf{a})$.

- To show $L_\Sigma(e + f) = L_\Sigma(e) \cup L_\Sigma(f)$, suppose $U \in L_\Sigma(e + f)$, i.e., $e + f \xrightarrow{U} g$ for $g \in \mathcal{F}$. By Lemma 7.4, we find $g' \in \mathcal{F}$ with $e \xrightarrow{U} g'$ or $f \xrightarrow{U} g'$, and hence $U \in L_\Sigma(e) \cup L_\Sigma(f)$.

For the other inclusion, suppose that $U \in L_\Sigma(e)$. We then have that $e \xrightarrow{U} g$ for some $g \in \mathcal{F}$. By Lemma 7.8, there exists a $g' \in \mathcal{F}$ such that $e + f \xrightarrow{U} g'$, and hence $U \in L_\Sigma(e + f)$. The case where $U \in L_\Sigma(f)$ can be treated similarly.

- To show $L_\Sigma(e \cdot f) = L_\Sigma(e) \cdot L_\Sigma(f)$, suppose that $U \in L_\Sigma(e \cdot f)$, i.e., $e \cdot f \xrightarrow{U} g$ for some $g \in \mathcal{F}$. By Lemma 7.5, we find $g_0, g_1 \in \mathcal{F}$ such that $U = U_0 \cdot U_1$ as well as $e \xrightarrow{U_0} g_0$ and $f \xrightarrow{U_1} g_1$. This means that $U_0 \in L_\Sigma(e)$ and $U_1 \in L_\Sigma(f)$, and thus $U \in L_\Sigma(e) \cdot L_\Sigma(f)$.

If $U \in L_\Sigma(e) \cdot L_\Sigma(f)$, then $U = U_0 \cdot U_1$ such that $U_0 \in L_\Sigma(e)$ and $U_1 \in L_\Sigma(f)$. This means that there exist $g_0, g_1 \in \mathcal{F}$ such that $e \xrightarrow{U_0}_\Sigma g_0$ and $f \xrightarrow{U_1}_\Sigma g_1$. By Lemma 7.9, there exists a $g \in \mathcal{F}$ such that $e \cdot f \xrightarrow{U}_\Sigma g$, and hence $U \in L_\Sigma(e \cdot f)$.

- To show $L_\Sigma(e \parallel f) = L_\Sigma(e) \parallel L_\Sigma(f)$, suppose $U \in L_\Sigma(e \parallel f)$, i.e., $e \parallel f \xrightarrow{U}_\Sigma g$ for some $g \in \mathcal{F}$. By Lemma 7.6, we find $g_1, g_2 \in \mathcal{F}$ and $U_1, U_2 \in \text{SP}(\Sigma)$ such that $U = U_1 \parallel U_2$ as well as $e_1 \xrightarrow{U_1}_\Sigma g_1$ and $e_2 \xrightarrow{U_2}_\Sigma g_2$. It then follows that $U = U_1 \parallel U_2 \in L_\Sigma(e) \parallel L_\Sigma(f)$.

If $U \in L_\Sigma(e) \parallel L_\Sigma(f)$, then $U = U_1 \parallel U_2$ such that $U_1 \in L_\Sigma(e)$ and $U_2 \in L_\Sigma(f)$. This means that there exist $g_1, g_2 \in \mathcal{F}$ such that $e \xrightarrow{U_1}_\Sigma g_1$ and $f \xrightarrow{U_2}_\Sigma g_2$. By Lemma 7.10, we find that $e \parallel f \xrightarrow{U}_\Sigma 1 \in \mathcal{F}$, and hence $U \in L_\Sigma(e \parallel f)$.

- To show $L_\Sigma(e^*) = L_\Sigma(e)^*$, suppose $U \in L_\Sigma(e^*)$, i.e., $e^* \xrightarrow{U}_\Sigma f$ for $f \in \mathcal{F}$. By Lemma 7.7, we find that $U = U_1 \cdots U_n$ and $f_1, \dots, f_n \in \mathcal{F}$ such that for $1 \leq i \leq n$ we have $e \xrightarrow{U_i}_\Sigma f_i$. Hence, we know for $1 \leq i \leq n$ that $U_i \in L_\Sigma(e)$, and therefore $U = U_1 \cdots U_n \in L_\Sigma(e)^*$.

For the other direction, let $U \in L_\Sigma(e)^*$. Then we can write $U = U_1 \cdots U_n$ such that for $1 \leq i \leq n$ it holds that $U_i \in L_\Sigma(e)$. We find for $1 \leq i \leq n$ an $f_i \in \mathcal{F}$ such that $e \xrightarrow{U_i}_\Sigma f_i$. By Lemma 7.11, we find an $f \in \mathcal{F}$ such that $e^* \xrightarrow{U}_\Sigma f$, and hence $U \in L_\Sigma(e^*)$. \square

Lemma 7.13. *For all $e \in \mathcal{T}$, we have $L_\Sigma(e) = \llbracket e \rrbracket$.*

Proof. The proof proceeds by induction on the structure of e . In the base, we need to consider the cases where $e \in \{0, 1\} \cup \Sigma$, all of which go through by Lemma 7.12.

For the inductive step, there are four cases. We argue the case where $e = e_1 + e_2$; the other cases are similar. Using Lemma 7.12 and the induction hypothesis, we can derive that

$$L_\Sigma(e_1 + e_2) = L_\Sigma(e_1) \cup L_\Sigma(e_2) = \llbracket e_1 \rrbracket \cup \llbracket e_2 \rrbracket = \llbracket e_1 + e_2 \rrbracket \quad \square$$

Lemma 7.14. *The syntactic pomset automaton is fork-acyclic.*

Proof. For fork-acyclicity, we define the following. For $e \in \mathcal{T}$, we define $\text{depth}(e)$ as follows:

$$\begin{aligned} \text{depth}(0) &= 0 & \text{depth}(e + f) &= \max(\text{depth}(e), \text{depth}(f)) & \text{depth}(e^*) &= \text{depth}(e) \\ \text{depth}(1) &= 0 & \text{depth}(e \parallel f) &= \max(\text{depth}(e), \text{depth}(f)) + 1 \\ \text{depth}(a) &= 0 & \text{depth}(e \cdot f) &= \max(\text{depth}(e), \text{depth}(f)) \end{aligned}$$

We now claim that if $e \preceq_\Sigma f$, then $\text{depth}(e) \leq \text{depth}(f)$. To see this, it suffices to prove the claim for the rules that generate \preceq_Σ ; this gives us three cases to consider.

- If $e \preceq_\Sigma f$ because there exists an $a \in \Sigma$ with $e \in \delta_\Sigma(f, a)$, we proceed by induction on f . In the base, $f \in \Sigma$ and $e = 1$; but then $\text{depth}(e) = 0 \leq \text{depth}(f)$ immediately.

For the inductive step, there are four cases to consider.

- If $f = f_1 + f_2$, then assume without loss of generality that $e \in \delta_\Sigma(f_1, a)$. By induction, $\text{depth}(e) \preceq_\Sigma \text{depth}(f_1)$; since $\text{depth}(f_1) \leq \text{depth}(f)$, the claim follows.
- If $f = f_1 \cdot f_2$, then there are two subcases to consider.
 - If $e \in \delta_\Sigma(f_1, a) \circledast f_2$, then $e = f'_1 \cdot f_2$ with $f'_1 \in \delta_\Sigma(f_1, a)$. By induction, $\text{depth}(f'_1) \leq \text{depth}(f_1)$. We then know that

$$\text{depth}(e) = \max(\text{depth}(f'_1), \text{depth}(f_2)) \leq \max(\text{depth}(f_1), \text{depth}(f_2)) = \text{depth}(f)$$

- If $e \in f_1 \star \delta_\Sigma(f_2, a)$, then $e \in \delta_\Sigma(f_2, a)$. By induction, $\text{depth}(e) \leq \text{depth}(f_2)$. Since $\text{depth}(f_2) \leq \text{depth}(f)$, the claim follows.

- We can disregard the case where $f = f_1 \parallel f_2$, for $\delta_\Sigma(f, a) = \emptyset$.

- If $f = f_1^*$, then $e = f'_1 \cdot f_1^*$ with $f'_1 \in \delta_\Sigma(f_1, \mathbf{a})$. By induction, $\text{depth}(f'_1) \leq \text{depth}(f_1)$. We then know that $\text{depth}(e) = \max(\text{depth}(f'_1), \text{depth}(f_1)) \leq \text{depth}(f_1) = \text{depth}(f)$.
- (ii) If $e \preceq_\Sigma f$ because there exists a $\phi \in \mathbb{M}(\mathcal{T})$ with $e \in \gamma_\Sigma(f, \phi)$, we proceed by induction on f . In the base, where $f \in \{0, 1\} \cup \Sigma$, the claim holds vacuously, because $\gamma_\Sigma(f, \phi) = \emptyset$.

For the inductive step, all cases except the one for parallel composition are similar to the argument above. Now, if $f = f_1 \parallel f_2$, then $e = 1$, and hence $\text{depth}(e) = 0 \leq \text{depth}(f)$.

- (iii) If $e \preceq_\Sigma f$ because there exists a $\phi \in \mathbb{M}(\mathcal{T})$ with $e \in \phi$ and $\gamma_\Sigma(f, \phi) \neq \emptyset$, we proceed by induction on f , showing $\text{depth}(e) < \text{depth}(f)$. In the base, the claim holds vacuously.

For the inductive step, there are four cases to consider.

- If $f = f_1 + f_2$, then assume without loss of generality that $\gamma_\Sigma(f_1, \phi) \neq \emptyset$. By induction, we have $\text{depth}(e) < \text{depth}(f_1)$. Since $\text{depth}(f_1) \leq \text{depth}(f)$, we are done.
- If $f = f_1 \cdot f_2$, then there are two subcases to consider.
 - If $\gamma_\Sigma(f_1, \phi) \neq \emptyset$ and $f_2 \neq \emptyset$, then $\gamma_\Sigma(f_1, \phi) \neq \emptyset$. By induction, we have that $\text{depth}(e) < \text{depth}(f_1)$. Since $\text{depth}(f_1) \leq \text{depth}(f)$, we are done.
 - If $f_1 \star \gamma_\Sigma(f_2, \phi) \neq \emptyset$, then $\gamma_\Sigma(f_2, \phi) \neq \emptyset$. By induction, we have that $\text{depth}(e) < \text{depth}(f_1)$. Since $\text{depth}(f_1) \leq \text{depth}(f)$, we are done.
- If $f = f_1 \parallel f_2$, then without loss of generality $\phi = \{f_1, f_2\}$ and $e = f_1$. By definition of $\text{depth}()$, we then find that $\text{depth}(e) = \text{depth}(f_1) < \text{depth}(f)$.
- If $f = f_1^*$, then $\gamma_\Sigma(f_1, \phi) \neq \emptyset$. By induction, $\text{depth}(e) < \text{depth}(f_1)$. Since $\text{depth}(f_1) = \text{depth}(f)$, we are done.

Now, if $e, f \in \mathcal{T}$ such that there exists a $\phi \in \mathcal{T}$ with $\gamma_\Sigma(e, \phi) \neq \emptyset$ and $f \in \phi$, then first of all $f \preceq_\Sigma e$ by definition. Suppose towards a contradiction that $e \preceq_\Sigma f$; then $\text{depth}(e) \leq \text{depth}(f)$ by the above. However, we also know that $\text{depth}(f) < \text{depth}(e)$ by (iii) above — a contradiction. We can therefore conclude that $e \not\preceq_\Sigma f$, and hence $f \prec_\Sigma e$. \square

Lemma 7.15. *The syntactic pomset automaton is bounded.*

Proof. We should show that for $e \in \mathcal{T}$, it holds that $\pi_\Sigma(e)$ is finite. Since $\pi_\Sigma(e)$ is the smallest closed set that contains e , it suffices to find a finite closed set $S(e)$ such that $e \in S(e)$; since $\pi_\Sigma(e) \subseteq S(e)$ by definition, the claim then follows. In particular, to show that $S(e)$ is support-closed, it suffices to verify that for $f \in S(e)$ the following hold:

- For all $\mathbf{a} \in \Sigma$, we have $\delta_\Sigma(e, \mathbf{a}) \subseteq S(e)$.
- For all $\phi \in \mathbb{M}(\mathcal{T})$, we have $\gamma_\Sigma(e, \phi) \subseteq S(e)$.
- If $e \in \phi \in \mathbb{M}(\mathcal{T})$ and $\gamma_\Sigma(e, \phi) \neq \emptyset$, then $e \in S(e)$.

We proceed by induction on e . In the base, there are two cases to consider.

- If $e \in \{0, 1\}$, then we choose $S(e) = \{e\}$. Note that $\delta_\Sigma(e, \mathbf{a}) = \emptyset$ for all $\mathbf{a} \in \Sigma$, and $\gamma_\Sigma(e, \phi) = \emptyset$ for all $\phi \in \mathbb{M}(\mathcal{T})$ — hence, the three conditions above hold vacuously.
- If $e = \mathbf{a}$ for some $\mathbf{a} \in \Sigma$, then we choose $S(e) = \{1, \mathbf{a}\}$. To see that $S(e)$ is support-closed, we inspect the derivatives of \mathbf{a} ; the derivatives for 1 satisfy the right conditions by the previous case. First, for all $\mathbf{b} \in \Sigma$, we have that $\delta_\Sigma(\mathbf{a}, \mathbf{b}) \subseteq \{1\}$, and hence $\delta_\Sigma(\mathbf{b}, \mathbf{a}) \subseteq S(e)$. Second, for all $\phi \in \mathbb{M}(\mathcal{T})$ we have that $\gamma_\Sigma(\mathbf{a}, \phi) = \emptyset$, and hence $\gamma_\Sigma(\mathbf{a}, \phi) \subseteq S(e)$. The case where $f \in \phi \in \mathbb{M}(\mathcal{T})$ with $\gamma_\Sigma(\mathbf{a}, \phi) \neq \emptyset$ cannot occur.

For the inductive step, there are four cases to consider.

- If $e = e_1 + e_2$, then we choose

$$S(e) = \pi_\Sigma(e_1) \cup \pi_\Sigma(e_2) \cup \{e\}$$

First, we note that $S(e)$ is finite by induction. To see that $S(e)$ is closed, it suffices to consider the derivatives of $\{e\}$, since $\pi_\Sigma(e_1)$ and $\pi_\Sigma(e_2)$ are closed by definition.

- For all $\mathbf{a} \in \Sigma$, we have that $\delta_\Sigma(e_1 + e_2, \mathbf{a}) = \delta_\Sigma(e_1, \mathbf{a}) \cup \delta_\Sigma(e_2, \mathbf{a})$. Now, since $\delta_\Sigma(e_1, \mathbf{a}) \subseteq \pi_\Sigma(e_1)$ and $\delta_\Sigma(e_2, \mathbf{a}) \subseteq \pi_\Sigma(e_2)$, we find that $\delta_\Sigma(e_1 + e_2, \mathbf{a}) \subseteq S(e)$ as well.
- For all $\phi \in \mathbb{M}(\mathcal{T})$, we have that $\gamma_\Sigma(e_1 + e_2, \phi) = \gamma_\Sigma(e_1, \phi) \cup \gamma_\Sigma(e_2, \phi)$. By an argument similar to the above, we find that $\gamma_\Sigma(e_1 + e_2, \phi) \subseteq S(e)$.
- If $f \in \phi \in \mathbb{M}(\mathcal{T})$ such that $\gamma_\Sigma(e_1 + e_2, \phi) \neq \emptyset$, then either $\gamma_\Sigma(e_1, \phi) \neq \emptyset$ or $\gamma_\Sigma(e_2, \phi) \neq \emptyset$. Hence, either $\gamma_\Sigma(e_1, \phi) \neq \emptyset$ or $\gamma_\Sigma(e_2, \phi) \neq \emptyset$, and hence $f \in \pi_\Sigma(e_1) \cup \pi_\Sigma(e_2) \subseteq S(e)$.
- If $e = e_1 \cdot e_2$, then we choose

$$S(e) = \pi_\Sigma(e_1) \ ; e_2 \cup e_1 \star \pi_\Sigma(e_2) \cup \pi_\Sigma(e_1)$$

First, note that $S(e)$ is finite by induction. Also, $e \in S(e)$, because $e_1 \in \pi_\Sigma(e_1)$. To see that $S(e)$ is support-closed, it suffices to consider the elements of the first set above, since $e_1 \star \pi_\Sigma(e_2)$ and $\pi_\Sigma(e_1)$ are already support-closed. Let $e' = e'_1 \cdot e_2$ for some $e'_1 \in \pi_\Sigma(e_1)$.

- For all $\mathbf{a} \in \Sigma$, we have that

$$\delta_\Sigma(e'_1 \cdot e_2, \mathbf{a}) = \delta_\Sigma(e'_1, \mathbf{a}) \ ; e_1 \cup e'_1 \star \delta_\Sigma(e_2, \mathbf{a})$$

Since $\delta_\Sigma(e'_1, \mathbf{a}) \cdot e_1 \subseteq \pi_\Sigma(e_1)$ and $\delta_\Sigma(e_2, \mathbf{a}) \subseteq \pi_\Sigma(e_2)$, the claim then holds.

- For all $\phi \in \mathbb{M}(\mathcal{T})$, we can show that $\gamma_\Sigma(e', \phi)$ again occurs in $S(e)$, by a similar argument.
- If $f \in \phi \in \mathbb{M}(\mathcal{T})$ such that $\gamma_\Sigma(e', \phi) \neq \emptyset$, then $\gamma_\Sigma(e'_1, \phi) \ ; e_2 \neq \emptyset$, or $e'_1 \star \gamma_\Sigma(e_2, \phi) \neq \emptyset$. In the former case, $f \in \pi_\Sigma(e'_1) \subseteq \pi_\Sigma(e_1) \subseteq S(e)$, while in the latter case $f \in \pi_\Sigma(e_2) \subseteq S(e)$.
- If $e = e_1 \parallel e_2$, then we choose

$$S(e) = \{e_1 \parallel e_2, 1\} \cup \pi_\Sigma(e_1) \cup \pi_\Sigma(e_2)$$

We again have that $S(e)$ is finite by induction. To see that $S(e)$ is support-closed, it suffices to consider the derivatives of $e_1 \parallel e_2$.

- For all $\mathbf{a} \in \Sigma$, we have that $\delta_\Sigma(e_1 \parallel e_2, \mathbf{a}) = \emptyset \subseteq S(e)$.
- For $\phi \in \mathbb{M}(\mathcal{T})$, we have that $\gamma_\Sigma(e_1 \parallel e_2, \phi) \subseteq \{1\} \subseteq S(e)$ by definition.
- For $f \in \phi \in \mathbb{M}(\mathcal{T})$ such that $\gamma_\Sigma(e_1 \parallel e_2, \phi) \neq \emptyset$, we have that $\phi = \{e_1, e_2\}$. In that case, $f \in \pi_\Sigma(e_1)$ or $f \in \pi_\Sigma(e_2)$.
- If $e = e_1^*$, then we choose

$$S(e) = \pi_\Sigma(e_1) \ ; e_1^* \cup \pi_\Sigma(e_1) \cup \{e_1^*\}$$

First, we note that $S(e)$ is again finite by induction. To see that $S(e)$ is support-closed, it suffices to consider $\pi_\Sigma(e_1) \ ; e_1^* \cup \{e_1^*\}$. To this end, let $e' = e'_1 \cdot e_1^*$ with $e'_1 \in \pi_\Sigma(e_1)$.

- For all $\mathbf{a} \in \Sigma$, we have that

$$\begin{aligned} \delta_\Sigma(e'_1 \cdot e_1^*, \mathbf{a}) &= \delta_\Sigma(e'_1, \mathbf{a}) \ ; e_1^* \cup e_1 \star \delta_\Sigma(e_1^*, \mathbf{a}) \\ &\subseteq \delta_\Sigma(e'_1, \mathbf{a}) \ ; e_1^* \cup \delta_\Sigma(e_1^*, \mathbf{a}) \\ &= \delta_\Sigma(e'_1, \mathbf{a}) \ ; e_1^* \cup \delta_\Sigma(e_1, \mathbf{a}) \ ; e_1^* \\ &\subseteq \pi_\Sigma(e_1) \ ; e_1^* \subseteq S(e) \end{aligned}$$

Furthermore, $\delta_\Sigma(e_1^*, \mathbf{a}) \subseteq S(e)$ by a similar argument.

- If $\phi \in \mathbb{M}(\mathcal{T})$, then $\gamma_\Sigma(e'_1 \cdot e_1^*, \phi) \subseteq S(e)$ and $\gamma_\Sigma(e_1^*, \phi) \subseteq S(e)$ by a similar argument.
- If $f \in \phi \in \mathbb{M}(\mathcal{T})$ and $\gamma_\Sigma(e'_1 \cdot e_1^*, \phi) \neq \emptyset$, then $\gamma_\Sigma(e'_1, \phi) \neq \emptyset$, hence $f \in \pi_\Sigma(e'_1) \subseteq \pi_\Sigma(e_1) \subseteq S(e)$. When $\gamma_\Sigma(e_1^*, \phi) \neq \emptyset$, we have that $f \in S(e)$ by a similar argument. \square

APPENDIX E. LEMMAS ABOUT EXPRESSIONS TO AUTOMATA

Lemma 8.1. *Let $p \in Q'$ and $r \in Q$. Now $U \in \llbracket e_{pr} \rrbracket$ if and only if $p \xrightarrow{U}_A r$ is a unit run.*

Proof. For the implication from left to right, suppose that $U \in \llbracket e_{pr} \rrbracket$; we have two cases:

- If $U = \mathbf{a}$ such that $r \in \delta(p, \mathbf{a})$, then $p \xrightarrow{U}_A r$ is a sequential unit run immediately.
- If $s_1, \dots, s_n \in Q$ with $U \in \llbracket e_{s_1} \parallel \dots \parallel e_{s_n} \rrbracket$ and $r \in \gamma(p, \{s_1, \dots, s_n\})$, then $U = U_1 \parallel \dots \parallel U_n$ such that for all $1 \leq i \leq n$ we have $U_i \in \llbracket e_{s_i} \rrbracket$. By the main induction hypothesis we have for each $1 \leq i \leq n$ an $s'_i \in F$ with $s_i \xrightarrow{U_i}_A s'_i$. Hence $p \xrightarrow{U}_A r$ is a parallel unit run.

For the converse implication, there are two cases to consider.

- If $p \xrightarrow{U}_A r$ is a sequential unit run, then $U = \mathbf{a}$ for some $\mathbf{a} \in \Sigma$ and $r \in \delta(p, \mathbf{a})$. In that case $U \in \llbracket e_{pr} \rrbracket$ by construction.
- If $p \xrightarrow{U}_A r$ is a parallel unit run, then $U = U_1 \parallel \dots \parallel U_n$ and there exist $s_1, \dots, s_n \in Q$ and $s'_1, \dots, s'_n \in F$ such that for $1 \leq i \leq n$ it holds that $s_i \xrightarrow{U_i}_A s'_i$, and furthermore $r \in \gamma(p, \{s_1, \dots, s_n\})$. By the main induction hypothesis, we have that $U_i \in \llbracket e_{s_i} \rrbracket$ for $1 \leq i \leq n$, and thus $U \in \llbracket U_1 \parallel \dots \parallel U_n \rrbracket$. We can then conclude that $U \in \llbracket e_{pr} \rrbracket$ again. \square