UNIVERSITÄT TRIER

Series-Parallel Graphs with Loops

Graphs Encoded by Regular Expressions

Stefan Gulan
FB IV, Informatik
gulan@uni-trier.de

TRIERER FORSCHUNGSBERICHTE
INFORMATIK / MATHEMATIK

NO. 10-2

April 2010

Abstract

We introduce a class of digraphs that generalizes the well-known class of arc-series-parallel-digraph. The new class is shown to be effectively recognizable, and a characterization by forbidden subgraphs is given. We argue that the forbidden subgraphs represent the structural features of finite automata that cannot be encoded by regular expressions, i.e., those causing an exponential blowup upon converting automata to expressions.

1 Motivation

A fundamental results in the theory of formal languages is the equivalent descriptive power of regular expressions and finite automata, as originally shown by Kleene [Kle56]. This brings up the problem of converting between the two representations, as regular expressions, being linear entities, are easily readable by humans, whereas finite automata are preferable on the machine level. Such conversions should go beyond mere proof-of-concept-constructions, i.e., it is desirable to optimize actual computations wrt. to time and/or memory. We take a quick look at some conversions for either direction.

Everyday use, notably pattern-matching, renders the conversion from expressions to automata the prevalent one. An established and intuitive paradigm in this translation is to interpret the parse of an expression as the structural information of a graph underlying an equivalent automaton [OF61, Tho68, SSS88, GF08b]; see Watson [Wat94] for a survey of such and other approaches. To the author's knowledge, the only work as yet asking for structural properties of automata constructed this way is by McNaughton [McN], although the discussion remains on an intuitive level. To quote:

> [...] although every regular expression can be transformed into a graph that has the same structure, the converse is not true. I will not define here precisely what I mean by the structure of a regular expression or graph, and hope that my point is made on an intuitive level. (p.35)

In contrast, the conversion of automata to expressions seems to be relevant mostly in the academical domain. It is generally performed via state-elimination, an algorithm originally proposed by Brzozowski & McCluskey [BJ63]. The size (and 'readability') of expressions generated by this method is highly dependent the chosen elimination-ordering. That aside, an exponential blowup cannot be avoided in the general case whatsoever, as was shown by Ehrenfeucht & Zeiger [EZ74]. Even so, no efficient general procedure is known that provides elimination-orderings resulting in outputs that are minimal wrt. to input-size. Still, several heuristics have been proposed in order to get reasonably good elimination-sequences [DM04, HW07, GF08a, AH09]; these approaches all rely on graph-theoretic properties of the input-automaton. In the (informal) spirit of above citation, it is worth noting that the automata which are drawn upon to compare the efficiency of such heuristics always show structural properties which admit no obvious corresponding regular expression.

The present work takes on above questions on a rigorous and formally sound graph-theoretic basis. As was already addressed, there is a structural and a quantitative aspect to this: we will concentrate on developing the structural theory of graphs that are in some sense equivalent to regular expressions; questions about efficiency will be treated only superficially.

2 Preliminaries

The graphs we consider are directed and may contain loops and multiple arcs. These are in general called *directed pseudographs*, we will refer to them as just *graphs*. Formally, a graph is a 4-tuple $G = (V, A, t, h)$ where V and A are finite disjoint sets, called the *vertices* resp. *arcs* of G, while t and h are maps from A to V. If G is not given explicitly, let $G = (V_G, A_G, t_G, h_G)$. The image of $a \in A_G$ under t_G resp. h_G is called the *tail* resp. *head* of a in G. If $t(a) = x$ and $h(a) = y$, we say that a *leaves* x and *enters* y, or that a is an xy-arc. For brevity, we write an xy-arc a as $a = xy$ and/or $a = xy \in A$. Tail and head are referred to as the *endpoints* of an arc. Distinct arcs of a graph with coinciding head and tail are called *parallel*. An xx-arc is an x-*loop* or just loop. An arc that is not a loop is *proper*. The set of arcs entering, resp. leaving x in G are denoted $\mathrm{I}(x)$, resp. $\mathrm{O}(x)$; the *in-degree* of x in G is $\mathrm{d}^-_G(x) = |\,\mathrm{I}(x)|$ and its *out-degree* in G is $\mathrm{d}^+_G(x) = |\,\mathrm{O}(x)|$. A *constriction* of G is a proper arc $a = xy$ where $\mathrm{d}^+_G(x) = 1 = \mathrm{d}^-_G(y)$. A vertex $x \in V_G$ is *simple*, if $\mathrm{d}^-_G(x) \leq 1$ and $\mathrm{d}^+_G(x) \leq 1$. Throughout this work, F, G and H denote graphs, while x, y and z denote vertices, and subscripts are omitted, if the graph they should indicate is understood.

F is a *subgraph* of G, denoted $F \subseteq G$, if $V_F \subseteq V_G$, $A_F \subseteq A_G$ and t_F and h_F are the appropriate restrictions of t_G resp. h_G; we say that G *contains* F. If F and G are subgraphs of H and $a = xy \in A_H$ with $x \in V_F$ and $y \in V_G$, then a is called an (F, G)-arc, as well as an (x, G)- or an (F, y)-arc of H.

A *path* of length n is a graph on $n+1$ vertices and n arcs s.t. every arc is a constriction; let $\mathbf{P_n}$ denote the path of length n and note that $\mathbf{P_0}$, called the *empty path*, is well-defined. Every path P contains exactly one vertex x with $\mathrm{d}^-_P(x) = 0$ and one vertex y with $\mathrm{d}^+_P(y) = 0$, P is then called a path from x to y, or an xy-path; x and y are called the *endpoints* of P, while the remaining vertices are its *internal* vertices. Two paths are *internally disjoint*, if their sets of internal vertices are disjoint. A *cycle* is a graph on n vertices and n arcs s.t. every arc is a constriction. A graph containing an xy-path for all $x, y \in V_G$ is called *strong*.

New graphs will be derived from given ones by means of several basic operations; we generally call both the operation and the graph it derives by the same name.
The *subdivision*, of an arc a in G is the replacement of a with $\mathbf{P_2}$ with the same orientation as a. Formally, the subdivision of a in G, for $a = xy \in A_G$, is the graph H satisfying $V_H = V_G \uplus z$, $A_H = A_G \setminus \{a\} \uplus \{a_1, a_2\}$, $t_H = t_G \setminus \{(a, x)\} \uplus \{(a_1, x), (a_2, z)\}$, and $h_H = h_G \setminus \{(a, y)\} \uplus \{(a_1, z), (a_2, y)\}$. More generally, a subdivision of G, referred to as a $\mathrm{D}G$, is any graph H s.t. there are graphs $G_1, \dots, G_n$ where $G = G_1$, G_{i+1} is the subdivision of an arc in G_i and $G_n = H$.
The *split* of a vertex x in G is the replacement of x with two vertices that separate the arcs entering x from those leaving x. Formally the split of x in G is the graph H satisfying $V_H = V_G \setminus \{x\} \uplus \{x_i, x_o\}$, $A_H = A_G \cup (x_i, x_o)$, $t_H = t_G \setminus \{(a, x) \mid a \in A_G\} \cup \{(a, x_o) \mid (a, x) \in t_G\}$, and h_H equivalently.
The *merge* of two vertices $x, y \in V_G$ is their identification by replacing them with one new vertex z and redirecting all arcs entering or leaving x or y to enter or leave z.

A graph G is *two-terminal*, if there are $s, t \in V_G$ s.t. every $x \in V_G$ lies on some st-path of G. The vertices s and t are respectively called the *source* and *sink* of G; we write $G = (G, s, t)$ to express that G is two-terminal with source s and sink t. A two-

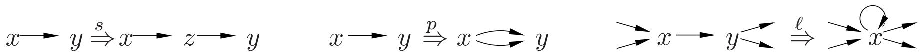

(a) series expansion **(b)** parallel expansion **(c)** loop expansion

Figure 1: Expanding an xy-arc, resp. the containing graph.

terminal graph (G, s, t) is a *hammock*, if $\mathrm{d}_G^-(s) = \mathrm{d}_G^+(t) = 0$. Let x and y be vertices of (G, s, t): x *dominates* y in G, if x lies on every sy-path in G, and x *co-dominates* y if x lies on every yt-path. Furthermore, x *guards* a vertex y, if x dominates and co-dominates y; also x guards an arc $a = yz$, if x guards both y and z. Note that in particular, every vertex guards itself and an x-loop is guarded by x. More generally, x guards a subgraph F of (G, s, t) if x guards every arc and/or vertex of F.

3 Series-Parallel-Loop - Graphs

Throughout this paper, we consider three graph-transformations, which are intended to represent the operators occurring in regular expression: concatenation, sum and iteration. This should serve as a motivation and reminder, as we will not consider language-theoretic aspects before Sec. 5.

Definition 1. The relations $\overset{s}{\Rightarrow}$, $\overset{p}{\Rightarrow}$ and $\overset{\ell}{\Rightarrow}$ are defined on graphs as follows: Let G be a graph and $a = xy \in A_G$, then

i) $G \overset{s}{\Rightarrow} H$ if H is the subdivision of a in G

ii) $G \overset{p}{\Rightarrow} H$ if $H = (V_G, A_G \uplus a_1, t_G \uplus \{(a_1, x)\}, h_G \uplus \{(a_1, y)\})$.

iii) $G \overset{\ell}{\Rightarrow} H$ if a is a constriction and H is obtained by merging x and y in G.

We say that H is derived from G by means of *series*-, *parallel*- or *loop*-expansion, if $G \overset{s}{\Rightarrow} H$, $G \overset{p}{\Rightarrow} H$ or $G \overset{\ell}{\Rightarrow} H$, respectively. The changes from G to H upon expansion are strictly local, as sketched in Fig. 1. We write $G \Rightarrow H$ if the particular type of relation is irrelevant, and $G \Rightarrow^\star H$ if H is derived from G by a (possibly empty) sequence of expansions.

Definition 2. The class $\mathcal{SPL}$ is generated by $\Rightarrow$ from $\mathbf{P_1}$ as follows

- $\mathbf{P_1} \in \mathcal{SPL}$; we call $\mathbf{P_1}$ the *axiom* of $\mathcal{SPL}$
- Let $G \in \mathcal{SPL}$, then $H \in \mathcal{SPL}$ if $G \overset{s}{\Rightarrow} H$ or $G \overset{p}{\Rightarrow} H$, or if $G \neq \mathbf{P_1}$ and $G \overset{\ell}{\Rightarrow} H$

The step-wise construction of an spl-graph is shown in Fig. 2. It is easily seen that every spl-graph is two-terminal. Excluding the axiom from being ℓ-expanded is done for technical reasons: this restriction guarantees that every spl-graph is a hammock. Note that the acyclic spl-graphs coincide with the arc-series-parallel graphs introduced by Valdes et al. [VTL81]; we will resort to their results on several occasions and elaborate on properties of $\mathcal{SPL}$ only that arise from its non-acyclic members.
To decide whether (G, s, t) is an spl-graph, the natural choice is a set of operations dual to expansions; we express them again as relationally. Some care must be taken with the removal of loops — this causes the new relations to be restricted to hammocks.

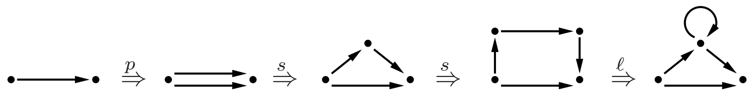

Figure 2: Constructing an spl-graph from $\mathbf{P_1}$ by a sequence of expansions.

Definition 3. The relations $\overset{s}{\Leftarrow}$, $\overset{p}{\Leftarrow}$ and $\overset{\ell}{\Leftarrow}$ are defined on hammocks follows: Let $G = (G, s, t)$ be a hammock, then

i) $G \overset{s}{\Leftarrow} H$, if some x is simple in G with incident arcs $a_1 = yx$ and $a_2 = xz$ and H satisfies $V_H = V_G \setminus \{x\}$, $A_H = A_G \setminus \{a_1, a_2\} \uplus \{a\}$, $t_H = t_G \setminus \{(a_1, y), (a_2, x)\} \uplus \{(a, y)\}$, and $h_H = h_G \setminus \{(a_1, x)(a_2, z)\} \uplus \{(a, z)\}$.

ii) $G \overset{p}{\Leftarrow} H$, if G contains distinct xy-arcs a_1, a_2, and $H = (V_G, A_G \setminus \{a_2\}, t_G \setminus \{(a_2, x)\}, h_G \setminus \{(a_2, y)\})$.

iii) $G \overset{\ell}{\Leftarrow} H$, if a is an x-loop in G s.t. x does not guard any arc besides a, and H is the split of x in $G \setminus \{a\}$.

If $G \overset{c}{\Leftarrow} H$ for $c \in \{s, p, \ell\}$ we say that G c-*reduces* to H and call both the replacement operation(s) in G yielding H and H itself a c-*reduction* of G. As with expansions, we write just $G \Leftarrow H$, if the particular type of reduction is not important, and $G \Leftarrow^\star H$ if H can be derived from G by a sequence of reductions.

Clearly, due to $\Leftarrow$ being defined on hammocks only, reduction is not the proper dual of expansion; however, even when restricted to hammocks we have

$$\overset{s}{\Leftarrow} = (\overset{s}{\Rightarrow})^{-1} \qquad \text{and} \qquad \overset{p}{\Leftarrow} = (\overset{p}{\Rightarrow})^{-1} \qquad \text{but} \qquad \overset{\ell}{\Leftarrow} \subsetneq (\overset{\ell}{\Rightarrow})^{-1},$$

i.e. ℓ-reducibility from G to H implies ℓ-expandability from H to G but not vice versa. The latter is due to the fact that if ℓ-expansion introduces an loop $a = xx$, x might guard some arc besides a, so a dual reduction is not guaranteed. Conveniently, this does not happen within $\mathcal{SPL}$, which is a consequence of the following

Proposition 3.1. *Let (G, s, t) be an spl-graph containing a cycle C. Then there is exactly one vertex $v \in V_C$ that guards C.*

Proof. The axiom $\mathbf{P_1}$ of $\mathcal{SPL}$ satisfies the claim, so assume $G \in \mathcal{SPL}$ does, and let $G \Rightarrow H$. Clearly, H inherits the claimed property from G if $G \overset{s}{\Rightarrow} H$ or $G \overset{p}{\Rightarrow} H$. In case $G \overset{\ell}{\Rightarrow} H$, let $a = xy$ be the constriction of G allowing for expansion and let $l = zz$ be the loop that emerges from it. The cycle given by l is guarded by z alone. For each cycle C of H beside that, G contains a cycle C'; by assumption C' is guarded by exactly one $q \in V_{C'}$. If $a \notin A_{C'}$, then C and C' are the same cycles and C is also guarded by q in H alone. Otherwise, if $q \in \{x, y\}$ then z guards C in H and if q is distinct from x and y, it still guards C. □

The intuition of Prop. 3.1 is that every cycle in an spl-graph contains a vertex that serves both as 'entry' and 'exit' of this cycle. This is crucial in proving that spl-reducibility is equivalent to -membership.

Theorem 3.2. $G \in \mathcal{SPL}$ iff $G \Leftarrow^\star \mathbf{P_1}$

Proof. Obviously $G \in \mathcal{SPL}$ if G can be reduced to $\mathbf{P_1}$, since the reversed sequence of reductions is an expansion-sequence. Sufficiency is shown by induction. The claim holds for $\mathbf{P_1}$, so assume that $G \in \mathcal{SPL}$ can be reduced and let $G \Rightarrow H$. We attend to ℓ-expansion only, the other cases are trivial. So let $G \overset{\ell}{\Rightarrow} H$ with $a = uv$ being the relevant constriction of G and $l = xx$ be the loop introduced in H. Assume that x guards some distinct arc $a' = yz$ in H, then G contains a cycle that defies Prop. 3.1, contrary to the assumption $G \in \mathcal{SPL}$. Therefore, x is no guard and $H \overset{\ell}{\Leftarrow} G$ a valid reduction; since by assumption $G \Leftarrow^\star \mathbf{P_1}$, we find $H \Leftarrow^\star \mathbf{P_1}$, which completes the proof. □

While Thm. 3.2 implies that membership of G in $\mathcal{SPL}$ can be decided by reducing G to the axiom-graph, it does not hint at how to do so. Actually, there is no need for a strategy, since the reduction-system exhibits unique normal-forms. To see this, we first show that reductions are locally confluent (see e.g. [Ohl02] for an introduction to confluence-properties of abstract rewriting systems).

Lemma 3.3. *Let G be a hammock. Then $G \Leftarrow H_1$ and $G \Leftarrow H_2$, implies the existence of a hammock J, satisfying $H_1 \Leftarrow^\star J$ and $H_2 \Leftarrow^\star J$.*

Proof. Let $G \overset{c_i}{\Leftarrow} H_i$ for $c_i \in \{s, p, \ell\}$. If $c_1, c_2 \in \{s, p\}$, the claim reduces to the equivalent property provided by Valdes et al. [VTL81] for acyclic digraphs; the generalization to non-acyclic hammocks is trivial. So let $c_1 = \ell$, and let $l = xx$ be the relevant loop in G. We assume that the subgraph of G that allows for $G \overset{c_2}{\Leftarrow} H_2$ contains an arc incident to x — otherwise, $G \overset{\ell}{\Leftarrow} H_1$ and $G \overset{c_2}{\Leftarrow} H_2$ take place in different regions of G and can be applied in any order, yielding the same graph.

- $c_2 = s$: Let y be the simple vertex to be removed, then G contains an xy- and a yz-arc (or a zy- and a yx-arc, which is symmetric). Applicability of ℓ-reduction implies that x is not a gate, so $z \neq y$. Also, x will not become a gate in H_2 due to s-reduction. Therefore, l-reduction is applicable to a in H_2, whereas s-reduction is applicable to y in H_1.

- $c_2 = p$: Let $G \overset{p}{\Leftarrow} H_2$ be valid due to parallel yz-arcs in G. Since we assume that ℓ-reduction is applicable to G in l, there is no loop parallel to l, i.e., at least one of y, z is distinct from x. Other than that, the argument is trivial.

- $c_2 = \ell$: If the two loops in G that permit the reductions share the vertex x, they are either parallel or identical. If they are parallel, ℓ-reduction is not applicable anyway, for x guards the 'other' loop. Since $G \overset{\ell}{\Leftarrow} H_1$ and $G \overset{\ell}{\Leftarrow} H_2$ are valid reductions, the loops must be identical, so $H_1 = H_2 = J$ and the statement is trivial.

□

Each reduction decreases the number of arcs or loops, while none introduces loops — hence every sequence of reductions eventually terminates. Any graph derived from G by exhaustive reduction is called *normal-form* of G and denoted $\mathrm{R}(G)$. Since reductions are locally confluent and terminating, we apply a well-known result from rewriting, namely Newman's Lemma [New42, Ohl02], with the following consequence.

Corollary 3.4. *Let G be a hammock, then $\mathrm{R}(G)$ is unique, hence*

$$G \in \mathcal{SPL} \quad \text{iff} \quad \mathrm{R}(G) = \mathbf{P_1}$$

A graph G that coincides with its normal-form, $G = \mathrm{R}(G)$, is called *reduced*.

4 Forbidden-Subgraph - Characterization of $\mathcal{SPL}$

We adapt the notion of topological minors that is well-known for undirected graphs, (see e.g. [Die06]) to our needs.

Definition 4. An *embedding* of F in G is an injection $e : V_F \to V_G$ satisfying that if $a = xy \in A_F$, then G contains an $e(x)e(y)$-path P_a, and that P_a and $P_{a'}$ are internally disjoint for distinct $a, a' \in A_F$.

If an embedding of F in G exists, we call F a *minor* of G *realized* by the embedding. We write $F \preccurlyeq G$ if F is a minor of G. If $F \preccurlyeq G$ does not hold then G is F*-free*; if $\mathcal{M}$ is a set of graphs and G is F-free for every $F \in \mathcal{M}$, then G is $\mathcal{M}$-free. It is easily seen that subdivisions allow for an equivalent characterization of minors:

Proposition 4.1. *$F \preccurlyeq G$* iff *G contains a* $\mathrm{D}F$

Let $F \preccurlyeq G$ be realized by e and $x \in V_F$, we call $e(x)$ a *peg* of F in G wrt. e; if G and e are known, we omit mentioning them. Observe that the in-/out-degree of a vertex does not exceed the in-/out-degree of its peg:

Proposition 4.2. *If e realizes $F \preccurlyeq G$, then $\mathrm{d}^-_F(x) \leq \mathrm{d}^-_G(e(x))$ and $\mathrm{d}^+_F(x) \leq \mathrm{d}^+_G(e(x))$.*

Let $F \preccurlyeq G$ be realized by e, a *bypass* of F in G wrt. e is a path from $e(x)$ to $e(y)$, s.t. xy is *not* an arc of F. An embedding of F in G is *bare*, if G contains no bypass of F wrt. to the embedding; we then write $M \sqsubseteq G$. Note that $F \preccurlyeq G$ might well be realized by various — in particular bare and non-bare — embeddings. Based on Prop. 4.1, we also call a $\mathrm{D}F$ in G bare, if G contains no bypass wrt. to the embedding realizing this $\mathrm{D}F$.

The existence of an xy-path is invariant under spl-expansion and -reduction, if x and y are not subject to the operation.

Proposition 4.3. *Let $G \Rightarrow H$ or $G \Leftarrow H$ and $\{x, y\} \subseteq V_G \cap V_H$, then G contains an xy-path* iff *H does.*

We start by excluding a family of graphs as minors of spl-graphs. Call a graph *bulky*, if it contains no loops, parallel arcs and simple vertices. Let $\boldsymbol{\mathcal{B}}$ denote the class of bulky graphs.

Lemma 4.4. *Every $G \in \mathcal{SPL}$ is $\boldsymbol{\mathcal{B}}$-free.*

Proof. Since all vertices of $\mathbf{P_1}$ are simple, Prop. 4.2 implies that $\mathbf{P_1}$ is $\boldsymbol{\mathcal{B}}$-free. Assume $G \in \mathcal{SPL}$ is $\boldsymbol{\mathcal{B}}$-free and let $G \Rightarrow H$. Consider any $F \in \boldsymbol{\mathcal{B}}$: since F is free of parallel arcs, and the existence of paths among vertices in $V_G \cap V_H$ is invariant under expansion (Prop. 4.3), $F \preccurlyeq H$ implies that a peg of F in H was introduced upon expansion. Hence in case $G \overset{p}{\Rightarrow} H$, F is not a minor of H, i.e., H is $\boldsymbol{\mathcal{B}}$-free. The same goes for $G \overset{s}{\Rightarrow} H$: as the new vertex in H is simple, but no vertex of F is, Prop. 4.2 implies that H is F-free and therefore $\boldsymbol{\mathcal{B}}$-free.
If $G \overset{\ell}{\Rightarrow} H$, let $a = xy$ be the relevant constriction of G and $l = zz$ the loop of H introduced by expansion. If $F \preccurlyeq H$ is realized by e, then $z = e(q)$ for some $q \in V_F$, as was discussed above. Let $H' = H \setminus l$: since F is free of loops, $F \preccurlyeq H'$ holds, too, and since q is not simple in F, z is not simple in H'. We actually find $\mathrm{d}^-_{H'}(z) \geq 2$ *and* $\mathrm{d}^+_{H'}(z) \geq 2$: if $\mathrm{d}^-_{H'}(z) = 0$, then $F \preccurlyeq G$ is realized by e', which is defined as e except

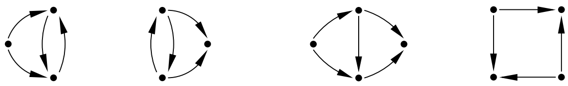

(a) C **(b) $\mathbf{C}^{\mathrm{R}}$** **(c) N** **(d) Q**

Figure 3: Bulky graphs constituting $\boldsymbol{\mathcal{F}}$.

that $e'(q) = y$ — contradicting the assumption that G is $\boldsymbol{\mathcal{B}}$-free. If $\mathrm{d}^{-}_{H'}(z) = 1$, there is exactly one arc entering z in H. Let this be $a' = z'z$, then $F \preccurlyeq G$ is realized by e'' which is as e except that, again, $e''(q) = y$, contradicting our assumption. A symmetric argument shows $\mathrm{d}^{+}_{H'}(z) \geq 2$. In fact, we have also shown that q, of which z is the peg, has in- and out-degree at least two.
But since $\mathrm{d}^{-}_{G}(x) = \mathrm{d}^{-}_{H'}(z)$ and $\mathrm{d}^{+}_{G}(y) = \mathrm{d}^{+}_{H'}(z)$, some $F' \in \boldsymbol{\mathcal{B}}$, constructed by splitting q in F satisfies $F' \preccurlyeq G$ — contradicting the assumption that G is $\boldsymbol{\mathcal{B}}$-free. □

The bulky graphs that are relevant for our purpose are given by $\boldsymbol{\mathcal{F}} = \{\mathbf{C}, \mathbf{C}^{\mathrm{R}}, \mathbf{N}, \mathbf{Q}\}$, shown in Fig. 3. Note that Valdes et al. proved that a an acyclic two-terminal graph is arc-series-parallel *iff* it is $\mathbf{N}$-free [VTL81].
The proof of Lem. 4.4 utilized that if G is $\boldsymbol{\mathcal{B}}$-free and $G \Rightarrow H$, then H is $\boldsymbol{\mathcal{B}}$-free. Put differently, expansion does not introduce a 'bulky subdivision' in H if none is present in G. Likewise, it can be shown in general that if $H \Leftarrow G$ and H is not $\boldsymbol{\mathcal{B}}$-free, then neither is H; however, there is a catch: the bulky minors of G need not be same as those of H. This already happens with $\boldsymbol{\mathcal{F}}$:

Lemma 4.5. *If $H \Leftarrow G$ for hammocks H and G, then*

i) $F \preccurlyeq H$ iff $F \preccurlyeq G$ *for $F \in \{\mathbf{C}, \mathbf{C}^{\mathrm{R}}, \mathbf{Q}\}$*

ii) $\mathbf{N} \preccurlyeq H$ *only if* $\mathbf{N} \preccurlyeq G$*, whereas*
$\mathbf{N} \preccurlyeq G$ *only if* $(\mathbf{N} \preccurlyeq G \vee \mathbf{C} \preccurlyeq G \vee \mathbf{C}^{\mathrm{R}} \preccurlyeq G)$

Proof. Actually, $F \preccurlyeq H$ *iff* $F \preccurlyeq G$ holds for all $F \in \boldsymbol{\mathcal{F}}$ in case of $H \overset{s}{\Leftarrow} G$ (easy) and $H \overset{p}{\Leftarrow} G$ (trivial). So let $H \overset{\ell}{\Leftarrow} G$ with loop $l = xx$ in H that allows for reduction, and $a = x_1x_2$ as the constriction arising in G. First, we show that $F \preccurlyeq H$ implies $F \preccurlyeq G$ for all $F \in \boldsymbol{\mathcal{F}}$. Let e realize $F \preccurlyeq H$; if x is not a peg of F in H, then e realizes $F \preccurlyeq G$ as well. If on the other hand, x is a peg, note that every $q \in V_F$ satisfies $\mathrm{d}^{-}_{F}(q) \leq 1$ or $\mathrm{d}^{+}_{F}(q) \leq 1$; it is easily verified that e' realizes $F \preccurlyeq G$, where $e'(r) = e(r)$ for $r \in V_F \setminus \{q\}$ and $e'(q) = x_1$, if $\mathrm{d}^{+}_{F}(q) \leq 1$ resp. $e'(q) = x_1$ otherwise. Conversely, starting from $F \preccurlyeq G$, we proceed by case distinction as in the claim.

i) If $F \preccurlyeq G$ for $F \in \{\mathbf{C}, \mathbf{C}^{\mathrm{R}}, \mathbf{Q}\}$

ii) If $\mathbf{N} \preccurlyeq G$, we distinguish whether one or both of x_1, x_2 are pegs of $\mathbf{N}$. If it is only one, $\mathbf{N} \preccurlyeq H$ is inferred similar to the converse direction; this is left to the reader. However, if both vertices are pegs, observe that, due to the in- and out-degrees of the x_i, the constriction x_1x_2 does not represent an arc of $\mathbf{N}$. By the same argument, the construction can neither be anti-parallel to an arc of $\mathbf{N}$, so x_1 and x_2 are pegs of the only two vertices of $\mathbf{N}$ that are not adjacent. It now easily follows that x_1 and x_2 lie on a cycle of G and further on a $\mathrm{D}\mathbf{C}$ or $\mathrm{D}\mathbf{C}^{\mathrm{R}}$ (Fig. 4). So in this case $\mathbf{C}$ or $\mathbf{C}^{\mathrm{R}}$ is a minor of G, hence also of H, as shown in the previous item.

□

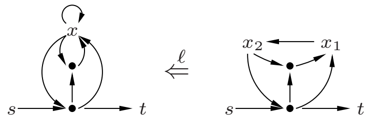

Figure 4: The graph $\mathbf{N}$ emerges as a subgraph due to ℓ-reduction of a hammock (H, s, t). However, both sides have $\mathbf{C}$ as a minor.

While Lem. 4.5 could be stated in greater detail wrt. $\mathbf{N}$, our primary interest is in $\boldsymbol{\mathcal{F}}$ as a whole. Still, $\boldsymbol{\mathcal{F}}$-freeness of a hammock does not suffice for membership in $\mathcal{SPL}$: for example, the hammock $\mathbf{\Phi}$, shown in Fig. 5a, is $\boldsymbol{\mathcal{F}}$-free, but not included in $\mathcal{SPL}$. The additional graphs necessary for the sought characterization are $\mathbf{\Phi}$, $\mathbf{\Psi}$, and $\mathbf{\Psi}^{\mathrm{R}}$, shown in Fig. 5.

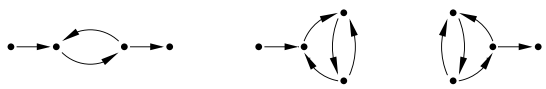

(a) $\mathbf{\Phi}$ **(b)** $\mathbf{\Psi}$ **(c)** $\mathbf{\Psi}^{\mathrm{R}}$

Figure 5: Graphs that do not allow for a bare embedding in any $G \in \mathcal{SPL}$.

Lemma 4.6. *Every $G \in \mathcal{SPL}$ is free of bare* $\mathbf{\Phi}$*-,* $\mathbf{\Psi}$*-, and* $\mathbf{\Psi}^{\mathrm{R}}$*-minors*

Proof. A bare embedding of one of $\mathbf{\Phi}$, $\mathbf{\Psi}$, or $\mathbf{\Psi}^{\mathrm{R}}$ would violate Prop. 3.1. □

On the other hand, each may well be a minor of certain spl-graphs, as the reader is invited to verify. In the absence of $\boldsymbol{\mathcal{F}}$-minors an invariance-result akin to Lem. 4.5 holds for bare subdivisions of these three graphs.

Lemma 4.7. *Let H be an $\boldsymbol{\mathcal{F}}$-free hammock and assume $H \Leftarrow G$, then $F \sqsubseteq H$* iff *$F \sqsubseteq G$ for $F \in \{\mathbf{\Phi}, \mathbf{\Psi}, \mathbf{\Psi}^{\mathrm{R}}\}$.*

Proof. Since $\mathbf{\Phi}$, $\mathbf{\Psi}$, and $\mathbf{\Psi}^{\mathrm{R}}$ all are free of parallel arcs, Prop. 4.3 provides the claim if all pegs occur in $V_G \cap V_H$; in particular, nothing needs to be done for $H \overset{p}{\Leftarrow} G$. We only prove the claim fully for $\mathbf{\Phi}$, the procedure is the same for $\mathbf{\Psi}$ and $\mathbf{\Psi}^{\mathrm{R}}$. In the following, let H be $\boldsymbol{\mathcal{F}}$-free.
Let $\mathbf{\Phi} \sqsubseteq H$ be realized by e. If $G \overset{s}{\Leftarrow} H$ removes a peg $x = e(q)$, q is one of the two simple vertices of $\mathbf{\Phi}$; here, let q be the unique vertex with $\mathrm{d}^-_{\mathbf{\Phi}}(q) = 0$ (the other case is symmetric). Since s-reduction is applicable due to x, an arc $a = yx$ exists in H, with y also occurring in G. Let e' be an embedding of $\mathbf{\Phi}$ in G, s.t. $e'(q) = y$ and e' as e for the other vertices. If e' is bare, the claim follows for $\mathbf{\Phi}$ and s-reduction, so assume it is not. Then G contains a bypass of $\mathbf{\Phi}$ wrt. e', which is necessarily a path *leaving* y, otherwise H would contain a bypass of $\mathbf{\Phi}$ wrt. e, contradicting the assumption that e is bare. We find $\mathbf{C} \preccurlyeq G$, if the other endpoint of the bypass is the peg of the vertex in $\mathbf{\Phi}$'s cycle that is not adjacent to q. If the bypass is from $e'(q)$ to the peg of the vertex with out-degree 0 in $\mathbf{\Phi}$, we get $\mathbf{Q} \preccurlyeq G$. In both cases Lem. 4.5 implies that H is not $\boldsymbol{\mathcal{F}}$-free, contradicting our assumption. Proving that s-reduction does not introduce new bare $\mathrm{D}\mathbf{\Phi}$'s is trivial.
Again let $\mathbf{\Phi} \sqsubseteq H$ be realized by e with peg $x \in V_H$. Considering $H \overset{\ell}{\Leftarrow} G$, let $a = xx$ be the loop that allows for reduction, and let x_1x_2 denote that constriction arising from it. As in the proof of Lem. 4.5 our argument is based on the facts that a is irrelevant for

(a) $Z = \{z\}$ **(b)** $|Z| \geq 2$

Figure 6: Cases distinguished in Prop. 4.10 if P_1 (dashed) and P_2 (dotted) are not internally disjoint.

the D$\mathbf{\Phi}$ in H and that $\mathrm{d}^-_G(x_1) = \mathrm{d}^-_{H\setminus a}(x)$ and $\mathrm{d}^+_G(x_2) = \mathrm{d}^-_{H\setminus a}(x)$ hold. Since every of $\mathbf{\Phi}$ has either in- or out-degree ≤ 1, we can construct an embedding e' of $\mathbf{\Phi}$ in G by assigning the role of x to either x_1 or x_2. □

In the remainder of this section, we will show that the properties proven in Lems. 4.4 and 4.6 are in fact sufficient for an characterization of spl-graphs via forbidden (bare) minors. First off, we need some preliminary propositions.

Proposition 4.8. *Let G be a reduced hammock with distinct arcs a_1, a_2 s.t. $h(a_1) = v = t(a_2)$. Then v is incident to a third proper arc.*

Proof. The vertex v is incident to a further arc a_3, otherwise G could be s-reduced. If a_3 is a loop, v must be a gate; then however, arcs different from the a_i connect v to the gated vertex. □

Proposition 4.9. *Let x and y be distinct vertices of a hammock (G, s, t). Then exactly one of the following is true in G*

1. *x dominates y*
2. *y dominates x*
3. *for some $z \in V_G \setminus \{x, y\}$ there are internally disjoint zx- and zy-paths in G*

Proof. If neither vertex dominates to other, let P_x and P_y be two shortest paths from s to x, resp. y. Since x and y are distinct, so are P_x and P_y. Let z denote that 'last' vertex, that occurs on both paths, then the subpaths of P_x and P_y that start from z, satisfy the claim. □

Proposition 4.10. *Let x and y be distinct vertices of a strong graph G, then there is a cycle $C \subseteq G$ and distinct $z_x, z_y \in V_C$, s.t. G contains an xz_x- and a yz_y-path that are disjoint.*

Proof. Since G is strong, let P_1 denote a shortest xy-path and P_2 a shortest yx-path in G. Consider the set of vertices belonging to both, $Z = (V_{P_1} \cap V_{P_2}) \setminus \{x, y\}$. If $Z = \emptyset$, P_1 and P_2 form a cycle, and the claim follows for $z_x = x$ and $z_y = y$; in this case, both claimed paths are empty. If $Z = \{z\}$, the claim follows for $z_x = x$ and $z_y = z$, or $z_x = x$ and $z_y = z$; here one of the claimed paths is empty (Fig. 6a). Finally, if $|Z| \geq 2$, let z_x and z_y be distinct elements of Z that are consecutive on the P_i, i.e., no other $z \in Z$ lies between z_x and z_y on P_1, resp. P_2. These vertices then satisfy the claim. □

Definition 5. A *kebab* is a connected graph consisting of three arc-disjoint subgraphs: a strong component B, called the *body*, and two nonempty vertex-disjoint paths S_1 and S_2, called the *spikes* of the kebab.

We further denote some unique vertices in a kebab: the endpoint of a spike that lies outside the body is called the *tip* of that spike, the endpoint that connects the spike to the body is called its *puncture*. A spike that enters the body of a kebab is called an *in-spike*, one that leaves the body is called an *out-spike*. If both spikes of a kebab K enter (leave) the body, K is also called an *in*-kebab (*out*-kebab), if one enters and the other leaves the body, K is called an *inout*-kebab.

Lemma 4.11. *Let G be an spl-reduced hammock containing a kebab. Then $F \preccurlyeq G$ for some $F \in \boldsymbol{\mathcal{F}}$ or $\mathbf{\Phi} \sqsubseteq G$.*

Proof. Let $G = (G, s, t)$ be a reduced hammock, since G contains at least one kebab, we choose a 'biggest' kebab $K \subseteq G$ as follows

1. the body of K is arc-maximal in G, in the sense that no kebab of G has a body with more arcs than K
2. the spikes of K are inclusion-maximal in G, i.e. they are not 'sub-spikes' of a bigger kebab with the same body as K.

We distinguish whether K is in an in-, an out- or an inout-kebab. The first and second case are symmetric, so we elaborate on the first and the last only. Throughout the proof, let B denote the body of K, S_1 and S_2 the spikes, with tips t_1 and t_2, and punctures p_1 and p_2, respectively.

1. K is an **in-kebab** (Fig. 7a).
Since (G, s, t) is a hammock, we apply Prop. 4.9 to t_1 and t_2 in G. The first two cases of the proposition are symmetric, so we distinguish by the second and third case only.

(a) If t_2 lies on every st_1-path, let P be a shortest t_2t_1-path in G. If P is disjoint with B, then P contains a segment P' from S_2 to S_1, that yields $\mathbf{C} \preccurlyeq G$ (Fig. 7b), proving the claim. So let P go through B, then the last segment of P is a (B, t_2)-path outside B. By the choice of K and P, it consist of a single arc $a = bt_1$ for $b \in V_B$ (Fig. 7c). According to Prop. 4.8, t_1 is incident to a further arc a', as G is reduced. Our choice of K requires that the other endpoint z of a' lies in K, since B, S_1 and a form a strong component bigger than B. It is now easy to see (from Fig. 7c), that $z \in V_{S_2}$ yields $\mathbf{C} \preccurlyeq G$ (regardless of a's orientation), and that $z \in V_B$ yields $\mathbf{C} \preccurlyeq G$ or $\mathbf{C}^{\mathrm{R}} \preccurlyeq G$ (depending on a's orientation), so let $z \in V_{S_1} \setminus \{p_1\}$. This leaves two possibilities: If $a' = zt_1$ we find $\mathbf{Q} \preccurlyeq G$, with pegs t_1,p_1,b and z (Fig. 7d). On the other hand, $a' = t_1z$ leads to a contradiction: Since G is p-reduced, there is at least one vertex z' between t_1 and z on S_1; omitting the t_1z'-segment of S_1 lets us identify an in-kebab with tips z' and t_2 and a body properly containing B (Fig. 7e), contradicting maximality of B.

(b) Let G contain a zt_1-path P_1 and a zt_2-path P_2 which are internally disjoint. If both P_i are disjoint with B, we find $\mathbf{C} \preccurlyeq G$ with help of Prop. 4.10 (Fig. 7f, where x_i denotes the first vertex on P_i that is also in S_i). If wlog. P_1 intersects B, let b denote the last vertex on P_1 that is in B and x the first vertex on P_1 that is in $V_{S_i} \setminus \{p_i\}$. If $x \neq t_1$, we find a kebab in G with a body containing B, contradicting our choice of K. As the claim was already proven for $x = t_1$ (see Fig. 7c), the statement follows for in-kebabs.

2. K is an **inout-kebab**, wlog with in-spike S_1.
Prop. 4.10 implies $\mathbf{\Phi} \sqsubseteq K$, realized by some embedding e, and with the tips of K

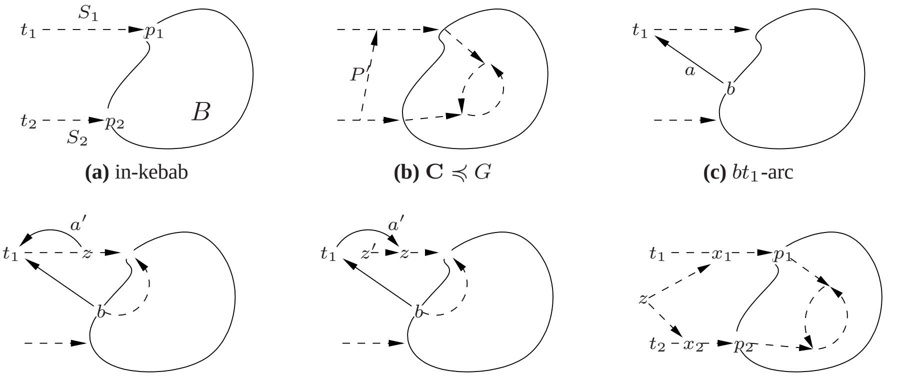

(a) in-kebab **(b)** $\mathbf{C} \preccurlyeq G$ **(c)** bt_1-arc

(d) $\mathbf{Q} \preccurlyeq G$ **(e)** $a' = t_1z$ **(f)** $\mathbf{C} \preccurlyeq G$

Figure 7: Cases occurring in the proof of Lem. 4.11 for K being an in-kebab. Solid arrows represent arcs, dashed arrows represent paths.

being pegs (Fig. 8a). If this is also true for G, we are done, so let P be a bypass of $\mathbf{\Phi}$ wrt. e in G. Then one endpoint of P is a tip of K: assume it is t_1 (the other case is symmetric) and distinguish by the orientation of P and P's other endpoint in K.

First assume that P is a t_1x-path: if $x \in V_B$, Prop. 4.10 yields $\mathbf{C} \preccurlyeq G$ (Fig. 8b); if $x \in V_{S_2}$, we find $\mathbf{Q} \preccurlyeq G$ with pegs t_1, x, p_1, and p_2 (Fig. 8c). Since the definition of a bypass implies $x \notin V_{P_1}$, these are all subcases relevant for a bypass leaving t_1.

Next, let P be an xt_1-path: since S_1 is maximal, the predecessor of t_1 on P must be a vertex of K. If $x \in V_B$ we find a situation similar to that shown in Fig. 7c, only that here, S_2 is an out-spike. Still, reasoning is equivalent and therefore omitted. If $x = t_2$, i.e., G contains a t_2t_1-arc (Fig. 8d), we apply Prop. 4.8 once more, finding new arcs a_i incident to each t_i. Let z_i denote the other endpoint of a_i: regardless of the a_i's orientations, if $z_i \notin V_K$ for either i, we find a bigger kebab (the details are left to the reader), so assume both z_i lie in K. We need to distinguish by the locations of the z_i in $V_K = V_{S_1} \cup V_{S_2} \cup V_B$ and the orientation of either a_i. First assume $z_i \in V_{S_i}$ for both i: if $t_1 = h(a_1)$ or $t_2 = t(a_2)$ we find $\mathbf{Q} \preccurlyeq G$; on the other hand, if $t_1 = t(a_1)$ and $t_2 = h(a_2)$, then, because G is p-reduced, there are vertices z_1 and z_2 s.t. either z_i lies between $t(a_i)$ and $h(a_i)$ on S_i, resulting in a bigger kebab in G (Fig. 8e) thus contradicting our assumption. Next, let $z_i \in V_B \setminus \{p_1, p_2\}$ and consider the orientation of a_1: if $a_1 = t_1z_1$, we find $\mathbf{C} \preccurlyeq G$ as in the case shown in Fig. 8b, whereas in case $a_1 = z_1t_1$, Prop. 4.10 implies $\mathbf{C}^{\mathrm{R}} \preccurlyeq G$ (Fig. 8f). A symmetric argument proves that $\mathbf{C}$ or $\mathbf{C}^{\mathrm{R}}$ is a minor of G, if $z_2 \in V_B \setminus \{p_1, p_2\}$. The remaining cases are those with $z_1 \in V_{S_2}$ and $z_2 \in V_{S_1}$ simultaneously, they are left to the reader.

□

We can now give a forbidden-subgraph characterization of spl-graphs. Since Valdes et al. [VTL81] showed that acyclic reduced hammocks that are not equal to $\mathbf{P_1}$ contain a D$\mathbf{N}$, we only proof the non-acyclic case. In the proof of the following lemma, we call a cycle *proper* if it is not merely a loop.

Lemma 4.12. *Let $G \neq \mathbf{P_1}$ be a reduced hammock with cycles. Then $F \preccurlyeq G$ for some $F \in \boldsymbol{\mathcal{F}}$ or $F' \sqsubseteq G$ for some $F' \in \{\mathbf{\Phi}, \mathbf{\Psi}, \mathbf{\Psi}^{\mathrm{R}}\}$.*

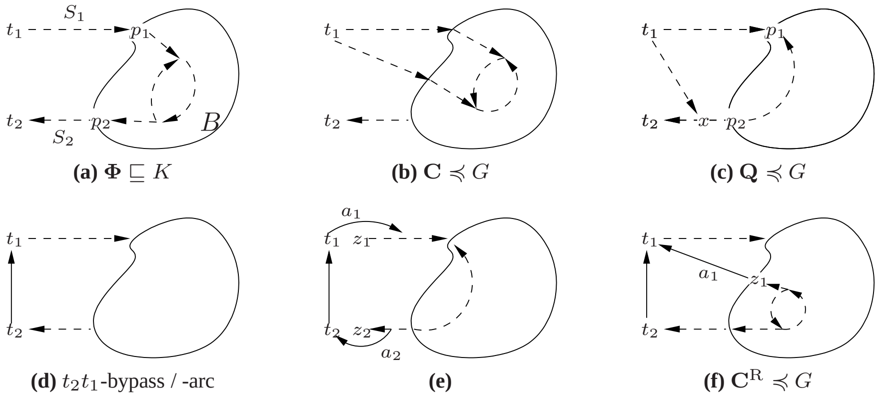

Figure 8: Cases occurring in the proof of Lem. 4.11 for K being an inout-kebab.

Proof. Let G be spl-reduced and not acyclic, then G contains proper cycles. This is because for a loop $a = xx \in A_G$, x guards a distinct arc $a' = yz$, for G is ℓ-reduced. If a' is a loop, then $x = y = z$, which can not be, as G is p-reduced. Hence G contains an xy- and a zx-path which form a proper cycle.
So let $C \subseteq G$ be a minimal proper cycle with distinct vertices x_1 and x_2. By Prop. 4.8 either x_i is incident to a further proper arc $a_i = x_iy_i$ or $a_i = y_ix_i$. We find $y_i \notin V_C$, otherwise G could be p-reduced or a_i would be a chord to C, contradicting C's minimality.
Now if $y_1 \neq y_2$, G contains a kebab with body C and spikes a_1 and a_2, and Lem. 4.11 provides the statement. If $y_1 = y_2$ we denote this vertex just y. If $a_i = yx_i$ for both i, we get $\mathbf{C} \preccurlyeq G$ while $a_i = x_iy$ yields $\mathbf{C}^{\mathrm{R}} \preccurlyeq G$, and the claim follows. So assume wlog. $a_1 = x_1y$ and $a_2 = yx_2$, and let $a_3 = yz$ or $a_3 = zy$ be the additional arc incident to y as assured by Prop. 4.8. Depending on the orientation of a_3, we find $\mathbf{C} \preccurlyeq G$ or $\mathbf{C}^{\mathrm{R}} \preccurlyeq G$, if $z \in V_C$ and $\mathbf{\Psi} \preccurlyeq G$ or $\mathbf{\Psi}^{\mathrm{R}} \preccurlyeq G$ otherwise. The first two cases yield the claim, while the latter two do not, so we proceed with $\mathbf{\Psi} \preccurlyeq G$, the other case is symmetric. The part of G 'found' so far is sketched in Fig. 9a. If this D$\mathbf{\Psi}$ is bare in G, the claim follows, so assume G contains a bypass. Except one, all possibilities lead immediately to $\mathbf{C} \preccurlyeq G$, $\mathbf{C}^{\mathrm{R}} \preccurlyeq G$ or a kebab; the exception is an yz-path, on which we elaborate.
Let P be the yz-bypass in G (Fig. 9b), once more Prop. 4.8 provides a new proper arc $a_4 = zz'$ or $a_4 = z'z$. If $z' \notin V_P$, we find a kebab (left to the reader), so assume $z' \in V_P$. However, $z' \neq z$ since a_4 is not a loop; if $z' = y$, we find $a_4 = yz$ (since a_4 can not be parallel to a_3), but then P is a bypass with several arcs — so a kebab with a final segment of P and one of a_1, a_2 as spikes is found. Finally, let z' be an inner vertex of P and denote the yz'-segment of P P_1 and the $z'z$-segment P_2. Depending on the orientation of a_4, either P_2 and a_4 form a cycle, or P_1, a_4 and a_3 do (Fig. 9c), and a kebab is immediately found in both cases. □

We have thus found a characterization of $\mathcal{SPL}$ by forbidden subgraphs.

Theorem 4.13. *Let G be a hammock, then $G \in \mathcal{SPL}$ iff G is $\boldsymbol{\mathcal{F}}$-free and no $F' \in \{\mathbf{\Phi}, \mathbf{\Psi}, \mathbf{\Psi}^{\mathrm{R}}\}$ is a bare minor of G.*

Proof. For $G \in \mathcal{SPL}$, Lem. 4.4 implies that G is $\boldsymbol{\mathcal{F}}$-free, and Lem. 4.6 states that none of $\{\mathbf{\Phi}, \mathbf{\Psi}, \mathbf{\Psi}^{\mathrm{R}}\}$ is a bare minor of G. Conversely, assume $G \notin \mathcal{SPL}$, hence

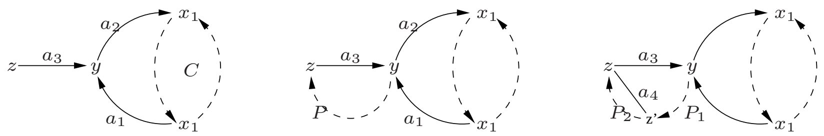

(a) $\mathbf{\Psi} \preccurlyeq G$ **(b)** yz-bypass of $\mathbf{\Psi}$ **(c)** $a_4 = zz'$ or $z'z$

Figure 9: Cases occurring in the proof of Lem. 4.12. Solid arrows represent arcs, dashed arrows represent paths. The arc a_4 in **(c)** is deliberately not drawn as an arrow.

$\mathrm{R}(G) \neq \mathbf{P_1}$ by Cor. 3.4. By Valdes' result and Lem. 4.12, we know $F \preccurlyeq \mathrm{R}(G)$ for some $F \in \boldsymbol{\mathcal{F}}$ and/or $F' \sqsubseteq \mathrm{R}(G)$ for some $F' \in \{\mathbf{\Phi}, \mathbf{\Psi}, \mathbf{\Psi}^{\mathrm{R}}\}$. If $G = \mathrm{R}(G)$, i.e., G is already reduced, the claim follows immediately; otherwise, induction on the length of the reduction using Lems. 4.5 and 4.7 provides the statement. □

5 SPL-Graphs and Regular Expressions

A *regular expression* (RE) over an alphabet is defined almost as usual: we allow for ε, $\bullet$, $+$, and *, but not $\varnothing$ — by the canonical semantics, this allows to express any language except the empty one. The set of REs over Σ is denoted $\mathrm{reg}(\Sigma)$. An *extended finite automaton* (*EFA*) is a 5-tuple $E = (Q, \Sigma, \delta, I, F)$, whose elements denote the set of states, the alphabet, the transition relation, the initial and the final states, respectively. These sets are all finite and satisfy $Q \cap \Sigma = \emptyset$, $\delta \subseteq Q \times \mathrm{reg}(\Sigma) \times Q$, $I \subseteq Q$, and $F \subseteq Q$. A *configuration* of E is a pair $Q \times \Sigma^*$; the relation $\vdash$ is defined on configurations of E as $(q, ww') \vdash (q', w')$, if $(q, \alpha, q') \in \delta$ and $w \in L(\alpha)$. The language accepted by E is

$$L(E) = \{w \mid (q_0, w) \vdash^* (q_f, \varepsilon) \text{ for } q_0 \in I,\ q_f \in F\}$$

The family of languages accepted by EFAs is exactly that of regular languages. Two EFAs accepting the same language are called *equivalent*. We further consider *finite automata* (FAs) which are nondeterministic and may have ε-transitions. By the above definition, an FA $A = (Q, \Sigma, \delta, I, F)$ is an EFA satisfying $\delta \subseteq Q \times (\Sigma \cup \{\varepsilon\}) \times Q$.

We treat EFA as graphs with REs as arc-labels and two distinguished sets of states; to this end, let $G(E)$ denote the graph underlying E. Formally, the graph underlying $E = (Q, \Sigma, \delta, I, F)$ is defined by $V_{G(E)} = Q$, $A_{G(E)} = \delta$, $t_{G(E)} : (q, \alpha, q') \mapsto q$, and $h_{G(E)} : (q, \alpha, q') \mapsto q'$. As $G(E)$ conveys only the graph-structural properties of E, the information about labels and initial and final states is generally considered separately.

Every EFA displays a compromise between the complexity of its transition-labels and that of its underlying graph; REs and FAs represent the extremes in this tradeoff: an RE can be considered as an EFA whose underlying graph is trivial, namely $\mathbf{P_1}$, while an FA is an EFA with trivial labels. Locally relaying information about a language between the graph-structure of an EFA and its labels is the basis of several conversions between REs to FAs. In the following, we investigate a few such conversions in either direction, in particular comparing the sizes of the in- and outputs. We do not give a precise definition of *size*, though, backed by the results of Ellul et al. [EKSW05], who showed that all (reasonable) such measures relate linerarly to another.

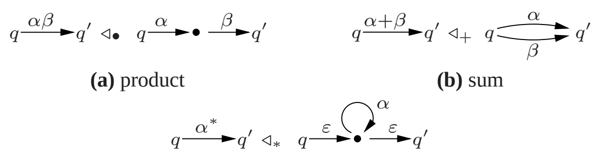

(a) product **(b)** sum

(c) star

Figure 10: Replacement of a transition (q, ρ, q') depending on ρ.

5.1 Expressions to Automata

The method given by Ott & Feinstein [OF61] is a rewriting-system on EFAs, where each rule replaces a transition according to its label; these rules are denoted $\triangleleft_\bullet$, $\triangleleft_+$, and $\triangleleft_*$, as shown in Fig. 10.
Given $\alpha \in \mathrm{reg}(\Sigma)$, the EFA $A^0_\alpha = (\{q_0, q_f\}, \Sigma, \{(q_0, \alpha, q_f)\}, \{q_0\}, \{q_f\})$, serves as the starting-point of the conversion; the algorithm then constructs a sequence of EFAs A^n_α, s.t. $A^i_\alpha \triangleleft_\mathrm{c} A^{i+1}_\alpha$ holds for some $\mathrm{c} \in \{\bullet, +, *\}$. Any such sequence terminates in an FA, denoted just A_α.

Lemma 5.1. *Every A_α satisfies $G(A_\alpha) \in \mathcal{SPL}$.*

Proof. Clearly, $G(A^0_\alpha) = \mathbf{P_1} \in \mathcal{SPL}$, so assume $G(A^i_\alpha) \in \mathcal{SPL}$ and let $A^i_\alpha \triangleleft_\mathrm{c} A^{i+1}_\alpha$ for some $\mathrm{c} \in \{\bullet, +, *\}$. The graph underlying A^i_α is s-expanded upon $A^i_\alpha \triangleleft_\bullet A^{i+1}_\alpha$ and p-expanded upon $A^i_\alpha \triangleleft_+ A^{i+1}_\alpha$. For $A^i_\alpha \triangleleft_* A^{i+1}_\alpha$ we find that $G(A^{i+1}_\alpha)$ can be derived from $G(A^i_\alpha)$ by two s-expansion, followed by an ℓ-expansion, i.e., $G(A^i_\alpha) \overset{s}{\Rightarrow}\overset{s}{\Rightarrow}\overset{\ell}{\Rightarrow} G(A^{i+1}_\alpha)$. Hence $G(A^n_\alpha) \in \mathcal{SPL}$ whenever A^n_α exists, and since $A_\alpha = A^k_\alpha$ for some k, the statement follows. □

Most methods that construct FAs from REs by manipulating graphs[1] work in a bottom-up-manner on the parse of the input. However each can be formulated in a top-down-manner, lending themselves to structural investigation as in Lem. 5.1. This reveals that the constructions by Sippu & Soisalon-Soininen [SSS88] and by Ilie & Yu [IY03] both yield FAs with spl-structure, whereas those by Thompson [Tho68] and Gulan & Fernau [GF08b] do not. Thompson's construction instroduces a $\mathrm{D}\mathbf{Q}$ for every Kleene-star in the input-expression; the construction by Gulan & Fernau introduce a $\mathrm{D}\mathbf{N}$ for certain products of sums. However, Thompson's method 'plays it safe' by introducing an excessive amount of ε-transitions connecting the subautomata — the bare $\mathrm{D}\mathbf{\Phi}$s result from this. The three other cited works gradually improve on Thompson's construction and each other wrt. lowering the size of the constructed automaton; in particular, Gulan & Fernau proved that their construction attains an optimal ratio between in- and output-sizes. Interestingly enough, the $\mathrm{D}\mathbf{N}$s allowed by their construction make for the only structural difference compared to the FAs provided by Ilie & Yu's method — this already suggests that $\mathbf{N}$-substructures allow for a certain conciseness of FAs as compared to regular expressions. In fact, a result by Korenblit & Levit [KL03] implies that $\mathrm{D}\mathbf{N}$s in the graph underlying an FA causes a quadratic blowup in the size of equivalent expressions.

[1] as opposed to algebraic methods, e.g. by derivatives [Brz64]; see [Wat94]

5.2 Automata to Expressions

An EFA E is called *normalized*, if $G(E)$ is a hammock; this is equivalent to requiring that E contains exactly one initial state q_0 and one final state q_f, and that every state lies on some path from q_0 to q_f. Any EFA can be transformed into an equivalent normalized EFA $\widehat{A}$ by removing the states and transitions that do not lie on some (I, F)-path, followed by adding a new initial state q_0 and a final state q_f with appropriate ε-transitions to I, resp. from F.

By augmenting the spl-reductions with labels, we get yet another rewriting-system on EFAs that is to be used with normalized EFAs only. The rules, $\triangleright_\bullet$, $\triangleright_+$, and $\triangleright_*$, are shown in Fig. 11, however, we restrict the applicability of star-reduction of a normalized EFA E to cases which allow for ℓ-reduction of the hammock $G(E)$.

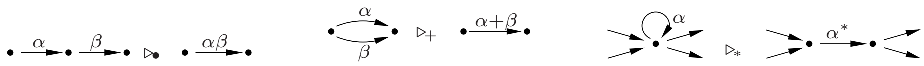

(a) product-reduction **(b)** sum-reduction **(c)** star-reduction

Figure 11: Labeled spl-reductions

The duplication of labels / subexpressions, which motivates the heuristics for state-elimination, is completely avoided by labeled spl-reduction. For this reason, it also is weaker than state-elimination, since $\triangleright_\bullet$ — the only rule actually eliminating states — is not applicable in the general case. On the other hand, let $\mathrm{R}_l(A)$ denote any EFA that is constructed by arbitrary but exhaustive application of labeled spl-reduction to $\widehat{A}$ for some FA A. We get

Lemma 5.2. $\mathrm{R}_l(A)$ *is unique modulo associativity and commutativity of labels.*

This is due to local confluence of spl-reduction, which carries over to labeled reduction. Hence $\mathrm{R}_l(A)$ is a regular expression (in the sense of being an EFA with trivial graph-structure) *iff* $G(\widehat{A}) \in \mathcal{SPL}$. For such FAs it is trivial to establish linear upper bounds on the size of the resulting RE depending on the size of the input FA. Applying the elimination-heuristic proposed by Han & Wood [HW07] to such fully reducible automata with n states provides these expressions in time $\mathcal{O}(n^2)$; this generalizes a result by Moreira & Reis who investigated optimal reduction-sequences on automata with (acyclic) series-parallel structure [MR09].

Regarding Lem. 5.1, it can also be shown that labeled reduction allows to convert an FA constructed by Ott & Feinstein's method back to the RE it originated from — if some simplifications of REs, like the removal of ε from products, are allowed for. Given that the graph underlying some *minimal* FA of a language is an spl-graph, we can then convert back an forth between a minimal FA and a minimal expression.

6 Concludision

The class $\mathcal{SPL}$ reflects the structural properties of graphs that can be encoded by regular expressions and vice versa: serial arcs are equivalent to products, parallel arcs are equivalent to sums, and loops are equivalent to iterations. Based on this, we argued (informally) that the sizes of expressions and their equivalent finite automata with spl-structure are linearly related. The forbidden-subgraph characterization of $\mathcal{SPL}$ has an interesting implication considering the contrapositive of above argument: Given an arbitrary automaton A, an exponential blowup cannot be avoided in the size of an expression equivalent to A; this study suggests that this blowup is caused by $\{\mathbf{C}, \mathbf{C}^{\mathrm{R}}, \mathbf{N}, \mathbf{Q}\}$ and $\{\mathbf{\Phi}, \mathbf{\Psi}, \mathbf{\Psi}^{\mathrm{R}}\}$ being (bare) minors of $\widehat{A}$.

References

[AH09] Jae-Hee Ahn and Yo-Sub Han. Implementation of state elimination using heuristics. In *CIAA 2009*, number 5642 in LNCS, pages 178–187, 2009.

[BJ63] J. A. Brzozowski and E. J. McCluskey Jr. Signal flow graph techniques for sequential circuit state diagrams. *IEEE Transactions on Electronic Computers*, 2:67–76, 1963.

[Brz64] Janusz A. Brzozowski. Derivatives of regular expressions. *Journal of the ACM*, 11:481–494, 1964.

[Die06] Reinhard Diestel. *Graph Theory*. Number 173 in Graduate Texts in Mathematics. Springer, 3 edition, 2006.

[DM04] Manuel Delgado and José Morais. Approximation to the smallest regular expression for a given regular language. Number 3317 in LNCS, pages 312–314. Springer, 2004.

[EKSW05] Keith Ellul, Bryan Krawetz, Jeffrey Shallit, and Ming-Wei Wang. Regular expressions: new results and open problems. *Journal of Automata, Languages and Combinatorics*, 10(4):407–437, 2005.

[EZ74] Andrzej Ehrenfeucht and Paul Zeiger. Complexity measures for regular expressions. In *STOC '74: Proceedings of the sixth annual ACM symposium on Theory of computing*, pages 75–79, New York, NY, USA, 1974. ACM Press.

[GF08a] Stefan Gulan and Henning Fernau. Local elimination-strategies in finite automata for shorter regular expressions. In *Theory and Practice of Computer Science: SOFSEM 2008*, volume 2, pages 46–57, 2008.

[GF08b] Stefan Gulan and Henning Fernau. Top-down construction of finite automata from regular expressions. In Ramesh Hariharan, Madhavan Mukund, and V. Vinay, editors, *28th International Conference on Foundations of Software Technology and Theoretical Computer Science*, pages 211–222, 2008.

[HW07] Yo-Sub Han and Derick Wood. Shorter regular expressions from finite-state automata. *Theoretical Computer Science*, 370:110–120, 2007.

[IY03] Lucian Ilie and Sheng Yu. Follow automata. *Information and Computation*, (186):140–162, 2003.

[KL03] Mark Korenblit and Vadim E. Levit. On algebraic expressions of series-parallel and fibonacci graphs. In Cristian S. Calude, Michael J. Dinneen, and Vincent Vajnovszki, editors, *DMTCS 2003*, number 2731 in LNCS, pages 215–224. Springer, 2003.

[Kle56] S. C. Kleene. *Representation of Events in Nerve Nets and Finite Automata*, pages 3–41. Annals of Mathematics Studies. 1956.

[McN] Robert McNaughton. Techniques for manipulating regular expressions. In John F. Hart and Satoru Takasu, editors, *Systems and Computer Science*, pages 27–41.

[MR09] Nelma Moreira and Rogério Reis. Series-parallel automata and short regular expressions. *Fundamenta Informaticae*, 91(3-4):611–629, 2009.

[New42] M.H.A. Newman. On theories with a combinatorial definition of "equivalence". *Annals of Mathematics*, 43(2), 1942.

[OF61] Gene Ott and Neil H. Feinstein. Design of sequential machines from their regular expressions. *Journal of the ACM*, 8(4):585–600, 1961.

[Ohl02] Enno Ohlebusch. *Advanced Topics in Term Rewriting*. Springer, 2002.

[SSS88] S. Sippu and E. Soisalin-Soininen. *Parsing Theory*. EATCS Monographs on Theoretical Computer Science. Springer, 1988.

[Tho68] Ken Thompson. Regular expression search algorithm. *Communications of the ACM*, 11(6):419–422, 1968.

[VTL81] Jacobo Valdes, Robert E. Tarjan, and Eugene L. Lawler. The recognition of series parallel digraphs. *SIAM Journal of Computation*, 11(2):298–313, 1981.

[Wat94] Bruce W. Watson. A taxonomy of finite automata construction algorithms. Technical Report Computing Science Note 93/43, Eindhoven University of Technology, May 1994.