Theory of Computing Systems manuscript No.
(will be inserted by the editor)

Series Parallel Digraphs with Loops

Graphs Encoded by Regular Expression

Stefan Gulan

Received: date / Accepted: date

Abstract In the conversion of finite automata to regular expressions, an exponential blowup in size can generally not be avoided. This is due to graph-structural properties of automata which cannot be directly encoded by regular expressions and cause the blowup combinatorially. In order to identify these structures, we generalize the class of arc-series-parallel digraphs beyond the acyclic case. The resulting digraphs are shown to be reversibly encoded by linear-sized regular expressions. Also, a characterization of this new class by a set of seven forbidden substructures is given. Automata that require expression of superlinear size must contain some of these substructures.

Keywords Digraphs, Regular Expressions, Forbidden Subgraphs

1 Introduction

A fundamental result in the theory of regular languages is the equivalent descriptive power of regular expressions and finite automata, as originally shown by Kleene [18]. While regular expressions come natural to humans as a way to describe such a language, automata are the objects of choice on the machine level. Consequently, converting between these two representations is of great practical importance. There are several linear-time algorithms to transform regular expressions into automata with size linear in that of the input; a detailed overview is given by Watson [25]. The converse construction, however, is considerably more troubling.

In particular, Ehrenfeucht & Zeiger [6] give a class of automata over a growing alphabet for which the size of any equivalent expression is exponential in the number

S. Gulan
Universität Trier
FB IV – Informatik
54286 Trier, Germany
Tel.: +46-651-201-4106
Fax: +49-651-201-3954
E-mail: gulan@uni-trier.de

of states. An automaton of this class contains a transition among each ordered pair of states, i.e., its structure is that of a complete digraph with loops. The smallest expression equivalent to such an n-state automaton contains at least 2^n literals.

This led Ellul et al. [7] to ask whether a similar blowup can be shown for automata over a fixed alphabet. The question was answered in the affirmative by Gelade & Neven [8] and independently by Gruber & Holzer [10]; the latter authors provided a proof for automata over a binary alphabet already. This mostly rules out alphabet size as a factor contributing to the exponential blowup, the modifier "mostly" giving credit to the fact that automata over a unary alphabet can be converted to expressions of quadratic size, using Chrobak's normal form [4, 7].

Intuitively, the exponential increase of expression- over automaton-size is due to the fact that automata are combinatorial objects, whereas expressions are terms, i.e., linear objects, that must resort to repeated subterms in order to reproduce language properties that ensue from the graph structure of an automaton. This was observed quite early by McNaughton [20], who remarks that "although every regular expression can be transformed into a graph that has the same structure, the converse is not true".

This motivates to investigate the interrelations between the (size of) expressions and the graph structure of corresponding automata. For example, automata constructed from expressions by means of a particular method admit structural properties that "reflect" both expression structure and the conversion method. The graph structure of automata constructed by Thompson's method has been analyzed by Giammaeresi et al. [9]. A structural characterization of Glushkov automata has been given by Caron & Ziadi [3]. See [25] for details of either construction.

Conversely, structurally restricted automata might allow for small equivalent expressions. Optimally, said restrictions might be exploited in the conversion process already. This was proposed for arc-series parallel substructures by Gulan&Fernau [12], and an efficient conversion to linear sized expressions was given Moreira & Reis [21] for automata that strictly adhere to this structure. The crucial feature of arc-series parallel graphs is their recursive definition by two composition rules that relate naturally to products and sums in regular expressions.

A prominent feature of arc-series parallel graphs among acyclic graphs is their characterization by a single forbidden (directed) minor [24], i.e., a structure that captures the complement of this class—among the acyclic graphs with a single source and sink, that is. Korenblit & Levit [19] conjecture that overlapping copies of this minor already cause a quadratic blowup in the size of expressions from acyclic automata with a single source and sink. Moreover, Gruber & Johannsen [11] prove that the blowup is already quasi-exponential for acylic finite automata (without source and sink restrictions), namely $n^{\Omega(\log n)}$.

However, the restriction to series parallel structures inherently confines formal treatment to finite languages. In order to accept infinite languages, automata need to contain cycles, which evades the class of arc-series-parallel digraphs.

This motivates our generalization of arc-series parallel digraphs beyond the acyclic case in Sec. 3. We augment the recursive definition of this class by adding a rule that produces loops; we refer to the resulting graphs as "series parallel loop graphs". This new class is shown to be effectively decidable within all graphs with a single source and sink vertex. Following that, we give two alternative characterizations of our new class.

In Sec. 4 we show that series parallel loop graphs are reversibly encoded by regular expressions. Moreover, we will see that every regular expression encodes a graph of this class. In fact, we give a bijection between our new graph class and the class expressions modulo certain (trivial) properties of expressions. Encoding and decoding are done in an automata-theoretic framework and are immediately applicable to convert among automata and expressions. We find linear bounds of the relative sizes of expressions and automata with series parallel structure in either direction.

In Sec. 5 we work out a forbidden subgraph characterization of our new class. We show that the absence of seven structures—or five modulo arc-reversal—suffices to imply that a graph is a series parallel loop graph. The characterization extends the one mentioned for arc-series parallel graphs.

The characterizations given in Secs. 4 and 5 can be read independently; the first touches matters of the complexity of regular language descriptors, while the second deals with structural graph theory. These topics are briefly discussed together in Sec. 6.

2 Preliminaries

The terms *set* and *class* are used interchangeably. Binary relations are written infixed. If R is a binary relation, then R^{-1} denotes its dual, and $R^{\star}$ denotes its reflexive transitive closure.

We consider finite directed graphs with loops and multiple arcs. These are canonically known as *directed pseudographs* but will be referred to just as *graphs*. Formally, a graph is a tuple (V,A,t,h) with *vertices* V, *arcs* A, *tail*-map $t : A \to V$ and *head*-map $h : A \to V$. If G is not given explicitly, then let $G = (V_G,A_G,t_G,h_G)$. The *(directional) dual* of G is the graph G^{R}, defined as $G^{\mathrm{R}} := (V_G,A_G,h_G,t_G)$. We say that G and G^{R} arise from another by *arc-reversal*.

Let $a \in A_G$ s.t. $t_G(a) = x$ and $h_G(a) = y$; we call a an *xy-arc* of G and write $a = xy \in A_G$ for short. We also call a an *out-arc* of x and an *in-arc* of y. If $a = xy \in A_G$, then x and y are *adjacent* in G, and x is a *predecessor* of y while y is a *successor* of x. If x is a vertex of a, i.e., $a = xy$ or $a = yx$, then a is *incident* to x.

The *in-degree* of $x \in V_G$, denoted $\mathrm{d}_G^-(x)$, is the number of in-arcs of x in G, and *out-degree* $\mathrm{d}_G^+(x)$ is the number out-arcs of x. A vertex x is *simple* in G if $\mathrm{d}_G^-(x) \leq 1$ and $\mathrm{d}_G^+(x) \leq 1$. A *constriction* of G is an xy-arc where $\mathrm{d}_G^+(x) = 1$ and $\mathrm{d}_G^-(y) = 1$. Distinct xy-arcs of a graph are *parallel* to each other. An xx-arc is called *x-loop* or just loop, while an arc that is no loop is a *proper* arc. If l is an x-loop, then *x carries l*.

The graph $G - x$ is derived from G by removing x from V_G, all arcs incident to x from A_G, and restricting the tail and head maps to arcs that are not incident to x. Likewise, removing an arc a from G yields the graph $G \setminus a$. Assuming that x is no vertex of G, adding x to G yields the graph $G + x$. If x and y are vertices of G, adding a new xy-arc to G yields the graph $G \cup xy$. If the removal of vertices and arcs from G yields F, then F is a *subgraph* of G, and we write $F \subseteq G$. If $F \subseteq G$, we say that G *contains F*.

To *merge* x and y in G yields the graph $G[x=y]$, where, informally, x and y are replaced by a new vertex z, and each arc incident to x or y is redirected to z. Formally, this is

$$\begin{aligned} G[x=y] :=& ((V_G \setminus \{x,y\}) \cup \{z\}, A_G, t', h'), \text{ where} \\ t' =& \{(a,v) \in t_G \mid v \notin \{x,y\}\} \cup \{(a,z) \mid (a,v) \in t_G, v \in \{x,y\}\}, \\ h' =& \{(a,v) \in h_G \mid v \notin \{x,y\}\} \cup \{(a,z) \mid (a,v) \in h_G, v \in \{x,y\}\}. \end{aligned}$$

To *split* x in G yields the graph $G \ll x \gg$, where, informally, x is replaced by vertices x_1 and x_2 and an x_1x_2-arc and each in-arc of x is redirected to x_1, while each out-arc of x is redirected to x_2. Formally, this is

$$\begin{aligned} G \ll x \gg :=& ((V_G \setminus \{x\}) \cup \{x_1, x_2\}, A_G, t', h') \cup x_1x_2, \text{ where} \\ t' =& \{(a,v) \in t_G \mid v \neq x\} \cup \{(a,x_2) \mid (a,x) \in t_G\}, \\ h' =& \{(a,v) \in h_G \mid v \neq x\} \cup \{(a,x_1) \mid (a,x) \in h_G\}. \end{aligned}$$

An (x,y)-*walk* W of *length* n in G is a sequence $W = a_1, \ldots, a_n$, $a_i \in A_G$, where $t(a_1) = x$, $h(a_i) = t(a_{i+1})$ for $1 \leq i < n$, and $h(a_n) = y$. We also say that W is a walk *from* x *to* y. The vertices x and y are the *endpoints* of W, while every $h(a_i)$, for $1 \leq i < n$, is an *internal* vertex of W. Notice that either endpoint might also be an internal vertex. An internal vertex of a walk *lies* on this walk, and a walk *passes through* its internal vertices.

An (x,y)-*path* in G is an (x,y)-walk s.t. neither x nor y is an internal vertex and every internal vertex occurs exactly once. An (F,F')-path for $F,F' \subseteq G$ is any (x,y)-path where $x \in V_F$ and $y \in V_{F'}$. A path in G is considered to be a subgraph of G. A subgraph of a path P that is a path itself is called a *segment* of P. Two paths $P_1, P_2 \subseteq G$ are *internally disjoint* if $V_{P_1} \cap V_{P_2}$ contains no internal vertex of either path. A *cycle* is an (x,x)-path, any (x,y)-path with $x \neq y$ is also called *proper*. Let $C \subseteq G$ be a cycle, then a *chord* of C is any $a = xy \in A_G \setminus A_C$ with $x,y \in V_C$.

The proper path of length $k \geq 1$ is denoted $\mathsf{P_k}$. In addition, the path of length zero, called the *empty* path is denoted P_0. A proper xy-arc induces a P_1 from x to y. To *subdivide* the arc $a = xy$ means to replace a by a P_2 from x to y. A *subdivision* of G is any graph that results from G by successively subdividing arbitrary arcs. Any subdivision of G is denoted $\mathrm{D}G$.

An *abstract rewriting system* (ARS) $\mathfrak{A}$ on a set U is a structure $\mathfrak{A} = \langle U, \rightarrow_1, \ldots, \rightarrow_n \rangle$, where each $\rightarrow_i$ is a binary relation on U. We call U the *universe*, and each $\rightarrow_i$ a *rule* of $\mathfrak{A}$. A *rewriting step* in $\mathfrak{A}$ is any member of any rule. If $x \rightarrow_i y$ is a rewriting step, we say that $\rightarrow_i$ *applies* to x, and that $\rightarrow_i$ *rewrites* x to y. More generally, a *rewriting* of x to y in $\mathfrak{A}$ is a sequence of rewriting steps

$$l_1 \rightarrow_{i_1} r_1, l_2 \rightarrow_{i_2} r_2, \ldots, l_n \rightarrow_{i_n} r_n,$$

where $l_1 = x$, $r_i = l_{i+1}$ for $1 \leq i < n$, and $r_n = y$. We say that x is $\rightarrow_i$-*normal*, if $\rightarrow_i$ does not apply to x. We further say that x is $\mathfrak{A}$-normal, or a normal form in $\mathfrak{A}$, if x is normal for every rule of $\mathfrak{A}$.

3 Series-Parallel Graphs with Loops

In this section we introduce the graph class which is the main subject of interest in this work. This class consists of graphs that are "reached" within an ARS by rewriting a smallest, explicitly given, member. A second ARS will be introduced to decide membership. For convenience, the universe of either ARS is restricted to a class which is now defined.

Definition 1 A *hammock* is a graph G with vertices *src* and *snk*, respectively called the *source* and *sink* of G, s.t. $\mathrm{d}_G^-(src) = 0$, $\mathrm{d}_G^+(snk) = 0$, and every vertex of G lies on a (src, snk)-path. The class of hammocks is denoted $\mathbf{H}$.

We write $(G, src, snk) \in \mathbf{H}$ to express that G is a hammock with source *src* and sink *snk*. If (G, src, snk) contains vertices x and y s.t. x lies on every (src, y)-path in G, then x *dominates* y, while y *is dominated* by x. Symmetrically, if x lies on every (y, snk)-path in G, x *co-dominates* y. If x dominates and co-dominates y, we say that x *guards* y, resp. that x is a *guard* of y. If x and y coincide, the guard is *trivial*.

More generally, if $F \subseteq G$ and $x \in V_G$ co-/dominates or guards every $y \in V_F$, then x co-/dominates or guards F in G. The domination relation is easily seen to be transitive and antisymmetric. It thus induces a partial order on the vertices of a hammock. A characterization for incomparability of vertices is given in the following;

Proposition 1 *Let H be a hammock with source src and distinct vertices x and y. Then exactly one of the following properties holds in H:*

1. *x dominates y*
2. *y dominates x*
3. *Some vertex z dominates both x and y, and H contains a (src, z)-path, a (z, x)-path and a (z, y)-path, which are pairwise internally disjoint.*

We consider three operations on hammocks which are presented in a relational fashion.

Definition 2 The relations $\overset{s}{\Rightarrow}$, $\overset{p}{\Rightarrow}$ and $\overset{\ell}{\Rightarrow}$, called *series expansion*, *parallel expansion* and *loop expansion* respectively, are defined as follows. Let $(G, src, snk) \in \mathbf{H}$ with $a = xy \in A_G$, then

- $G \overset{s}{\Rightarrow} H$, if H is obtained from G by subdividing a. Formally, $H = ((G \setminus a) + z) \cup \{xz, zy\}$.
- $G \overset{p}{\Rightarrow} H$, if H is obtained from G by adding a further xy-arc. Formally, $H = G \cup xy$.
- $G \overset{\ell}{\Rightarrow} H$, if a is a constriction where $x \neq src$ and $y \neq snk$, and H is obtained from G by merging x and y. Formally, $H = G[x = y]$.

The local changes in a graph are sketched in Fig. 1 for each expansion. Observe that each expansion is indeed defined on $\mathbf{H}$, and that the sources, resp. sinks, of G and H coincide in each case. We set

$$\Rightarrow := \overset{s}{\Rightarrow} \cup \overset{p}{\Rightarrow} \cup \overset{\ell}{\Rightarrow},$$

and write just $G \Rightarrow H$ if the particular kind of expansion from G to H is irrelevant.

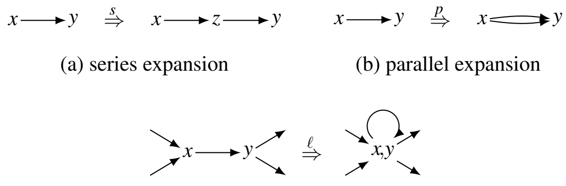

(a) series expansion

(b) parallel expansion

(c) loop expansion, for a constriction xy

Fig. 1: Changes in a graph upon expanding an xy-arc.

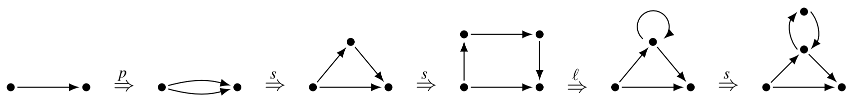

Fig. 2: Construction of an spl-graph from P_1 by a sequence of expansions.

Definition 3 The class of *series parallel loop graphs*, denoted **SPL**, is generated from P_1 by $\overset{s}{\Rightarrow}$, $\overset{p}{\Rightarrow}$ and $\overset{\ell}{\Rightarrow}$. It is the smallest set satisfying

- $\mathsf{P}_1 \in \mathbf{SPL}$, and
- if $G \in \mathbf{SPL}$ and $G \Rightarrow H$, then $H \in \mathbf{SPL}$.

For brevity we speak of *spl-graphs*. The graph P_1 is called the *axiom* of **SPL**. An example for the construction of an spl-graph from P_1 is shown in Fig. 2.

We observe that the connectivity among two vertices of a graph is invariant under expansion, as long as neither vertex is removed.

Proposition 2 *Assume $G \Rightarrow H$ where $x, y \in V_G \cap V_H$. Then G contains an (x,y)-path* iff *H does.*

A first property that relies on this fact is the following.

Proposition 3 *Every cycle in an spl-graph is guarded by exactly one of its vertices.*

Proof The claim is vacuously true for P_1, so suppose it is true for $(G, src, snk) \in \mathbf{SPL}$, and let $G \Rightarrow H$. It is easy to see that the claimed property carries over if H is derived from G by means of s- or p-expansion.

For $G \overset{\ell}{\Rightarrow} H$ let $a = xy$ be the constriction that is expanded in G and let z denote the merge vertex of x and y in H. We consider the cycles of H. The cycle that was introduced with expansion consists of z and the loop $l = zz$. This cycle is guarded by its only vertex and thus satisfies the claim. Let $v, z \in V_G$ be distinct and assume neither vertex is incident to a in G, which implies $v, z \in V_G \cap V_H$. We have already noticed that the source and the sink of G and H coincide. Thus, following Prop. 2, v guards z in H *iff* v guards z in G. Applying this to vertices of any cycle in H, this carries the inductive assumption from G to H. □

Corollary 1 $\mathbf{SPL} \subsetneq \mathbf{H}$

Fig. 3: A hammock that does not belong to **SPL**

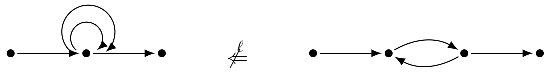

Fig. 4: Effect of "reducing" a loop on sight. Although the left-hand side is an spl-graph, the right-hand side cannot be further reduced to P_1.

Proof The expansions are defined on **H**, so $\mathbf{SPL} \subseteq \mathbf{H}$ follows from $\mathsf{P}_1 \in \mathbf{H}$ inductively. The hammock shown in Fig. 3 contains a cycle that defies Prop. 3, so the inclusion is proper. □

To decide the membership of G in **SPL**, we look for a sequence of expansions from P_1 to G. This is done inside a second ARS on **H**, wherein we construct such a sequence backwards.

Definition 4 The relations $\stackrel{s}{\Leftarrow}$, $\stackrel{p}{\Leftarrow}$ and $\stackrel{\ell}{\Leftarrow}$, called *series reduction*, *parallel reduction* and *loop reduction* respectively, are defined on **H** as follows. Let (G, src, snk) be a hammock, then

- $G \stackrel{s}{\Leftarrow} H$, if $y \in V_G$ is simple with predecessor x and successor z, and H is obtained from G by removing y and adding an xz-arc. Formally $H = (G - y) \cup xz$.
- $G \stackrel{p}{\Leftarrow} H$, if G contains parallel arcs a and a' and H is obtained from G by removing one of them. Formally $H = G \setminus a$.
- $G \stackrel{\ell}{\Leftarrow} H$, if G contains an x-loop l s.t. x is not a guard, no loop is parallel to l, and H is obtained from G by removing l and then splitting x. Formally $H = (G \setminus l) \ll x \gg$.

Reductions do not destroy the defining properties of a hammock, i.e., they are also relations on **H**. We abbreviate the reduction relations as s-, p- and ℓ-reduction. Reductions are formally investigated in the ARS

$$\mathfrak{R} := \langle \mathbf{H}, \stackrel{s}{\Leftarrow}, \stackrel{p}{\Leftarrow}, \stackrel{\ell}{\Leftarrow} \rangle.$$

Similar to expansions, we set

$$\Leftarrow := \stackrel{s}{\Leftarrow} \cup \stackrel{p}{\Leftarrow} \cup \stackrel{\ell}{\Leftarrow}$$

and write just $G \Leftarrow H$ if the particular type of reduction from G to H is irrelevant.

Notice that ℓ-reduction is restricted wrt. the vertex which carries the loop. In particular, this ensures that a loop is not ℓ-reduced in the presence of parallel loops, which might lead to false negatives (see Fig. 4). In that case, the restriction enforces that p-reduction is applied before ℓ-reduction "becomes applicable".

It is important to realize that the reduction relation is not the proper dual of the expansion relation. More precisely, we find

$$\stackrel{s}{\Leftarrow} = (\stackrel{s}{\Rightarrow})^{-1} \text{ and } \stackrel{p}{\Leftarrow} = (\stackrel{p}{\Rightarrow})^{-1}, \text{ whereas } \stackrel{\ell}{\Leftarrow} \subsetneq (\stackrel{\ell}{\Rightarrow})^{-1}.$$

This is due to the restriction on ℓ-reduction. Consider Fig. 4 again: there, the left hand side can not be ℓ-reduced to the right hand side. However, the hammock on the right *can* be ℓ-expanded to the one on the left. Because of this asymmetry, there is actually something to prove for the following result.

Proposition 4 *$G \in$ **SPL** iff P_1 is an $\mathfrak{R}$-normal form of G.*

Proof As stated before, $G \Leftarrow H$ implies $H \Rightarrow G$. From this, we infer that $G \Leftarrow^\star \mathsf{P}_1$ implies $\mathsf{P}_1 \Rightarrow^\star G$, meaning that $G \in \mathbf{SPL}$.

The converse direction is proven by structural induction. For $G = \mathsf{P}_1$ the claim holds. Assume that the claim is true for $G \in \mathbf{SPL}$ and consider $H \in \mathbf{SPL}$, derived via $G \Rightarrow H$. For $G \overset{s}{\Rightarrow} H$ and $G \overset{p}{\Rightarrow} H$, we find $H \overset{s}{\Leftarrow} G$ and $H \overset{p}{\Leftarrow} G$ immediately. For $G \overset{\ell}{\Rightarrow} H$, let $a = xy$ denote the constriction that is ℓ-expanded, i.e., assume $H = G[x = y]$. Further, let z denote the merge vertex of x and y, and let $l = zz$ denote the loop that is introduced with expansion. We need to show that ℓ-reduction is applicable to l in H, i.e., that no loop is parallel to l and that z is not a guard.

Since xy is a constriction in G, it follows immediately that no loop is parallel to l in H. For the sake of contradiction, now suppose that z is a guard in H. Then some vertex k is guarded by z and the two vertices lie on a cycle C in H. Following Prop. 3, C is guarded by exactly one of its vertices, which must be z. Consider the respective cycle C' in G that contains k, x and y. Since $G \in \mathbf{SPL}$, Prop. 3 yields that C' that is guarded by exactly one of its vertices. This guard must be x or y, any other guard would also be present for C in H (by virtue of Prop. 2), contradicting Prop. 3. However, x does not co-dominate y, while y does not dominate x in H, thus neither vertex guards the other. Therefore, none of its vertices guards C' at all, which contradicts Prop. 3 once again. Therefore, z is no guard, and $H \overset{\ell}{\Leftarrow} G$ holds.

We found that $G \Rightarrow H$ implies $H \Leftarrow G$ in each case. Since we assumed $G \Leftarrow^\star \mathsf{P}_1$, this also yields $H \Leftarrow^\star \mathsf{P}_1$. $\square$

So we can test membership in **SPL** by finding a reduction of a graph to the axiom of **SPL**. Conveniently, this requires no particular strategy, as each reduction eventually yields the same normal form.

Theorem 1 *The system $\mathfrak{R}$ is terminating and admits unique normal forms.*

Proof We first show that $\mathfrak{R}$ is terminating. Let $p(G)$ denote the sum of arcs and loops in $G \in \mathbf{H}$, i.e., set

$$p(G) := |A_G| + |\{l \mid l \in A_G \text{ and } t_G(l) = h_G(l)\}|.$$

Observe that that each loop of G is counted twice in $p(G)$. For $G \Leftarrow H$ we now find $p(H) < p(G)$. Since $p(\cdot) \in \mathbb{N}$, each reduction eventually terminates.

Next, we show that $\mathfrak{R}$ is locally confluent. That is, for reductions $G \Leftarrow H_1$ and $G \Leftarrow H_2$, we always find some J s.t. $H_1 \Leftarrow^\star J$ and $H_2 \Leftarrow^\star J$ holds. More specifically, let $G \overset{c_i}{\Leftarrow} H_i$ for $i \in \{1,2\}$. For $c_1, c_2 \in \{s, p\}$, this was shown by Valdes et al. [24] for acyclic hammocks; the proof carries over to the general case without effort. We thus fix $c_1 = \ell$ and let $H_1 = (G \setminus l) \ll q \gg$; recall that we denote the left and right split-vertices of q as q_1 and q_2, respectively. Now we distinguish by c_2.

- $c_2 = s$: For $H_2 = (G - y) \cup xz$ we find $q \neq y$, since in a hammock, a simple vertex cannot carry a loop. If $q \notin \{x, z\}$, we apply s-reduction to y in H_1 and ℓ-reduction in q to H_2. Either operation yields J, i.e.

$$(H_1 - y) \cup xz = J = (H_2 \setminus l) \ll q \gg .$$

  If $q = x$ (or $q = z$, which is shown analogous), we denote this vertex q. Then H_2 actually is $H_2 = (G - y) \cup qz$ and J can be derived by s-reducing H_1 in y and ℓ-reducing H_2 in l.
- $c_2 = p$: Let a and a' denote the parallel arcs of G that allow for reduction. Since ℓ-reduction applies to l in G, no arc is parallel to l, so $l \notin \{a, a'\}$. In this case, the order of the two operations can clearly be swapped without changing the resulting graph J.
- $c_2 = \ell$: let l' be the q'-loop in H that allows for reduction to H_2. Since ℓ-reduction applies to l and l' in H these loops are not parallel, so $q \neq q'$ follows. Thus ℓ-reduction applies to l' in H_1 and to l in H_2, and either reduction yields J.

From the properties that $\mathfrak{R}$ is terminating and locally confluent, the claim follows as an application of Newman's Lemma [22,16]. ☐

As a corollary, we get a stronger variant of Prop. 4.

Theorem 2 $G \in \mathbf{SPL}$ iff $\mathsf{R}(G) = \mathsf{P}_1$.

4 Encoding by Regular Expressions

Syntax and semantics of regular expressions (REs) follow Hopcroft & Ullman's textbook [14] except that we do not allow for $\emptyset$ in REs. As for notation, L_r denotes the language described by the RE r and $\text{reg}(\Sigma)$ denotes the set of all REs over an alphabet Σ. An RE is *simplified*, if it does not contain ε as a factor. any RE r can be converted to a simplified expression $\text{simp}(r)$, satisfying $L_{\text{simp}(r)} = L_r$, by replacing every subexpression $r\varepsilon$ or εr with just r. The size of r is given by the frequency of letters and ε in r, this value is called the *alphabetic width* of r and denoted $\text{alph}(r)$.

An *extended finite automaton* (EFA) over Σ 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 $Q \cap \Sigma = \emptyset$, $\delta \subseteq Q \times \text{reg}(\Sigma) \times Q$, $I \subseteq Q$, and $F \subseteq Q$. The relation $\vdash_E$ is defined on $Q \times \Sigma^*$ as $(q, ww') \vdash_E (q', w')$, if $(q, r, q') \in \delta$ and $w \in L_r$. The language accepted by E is

$$L(E) := \{w \mid (q_i, w) \vdash_E^* (q_f, \varepsilon) \text{ for } q_i \in I, q_f \in F\}.$$

Two EFAs are *equivalent* if they accept the same language. An EFA is *normalized* if $|I| = |F| = 1$ and the initial and final state are distinct; any EFA can normalized by adding a new initial (final) state and ε-transitions from (to) this new initial (final) state to (from) the original ones. The EFA E is *trim*, if for every state q of E there is a word $w = w_1 w_2 \in L(E)$ s.t. $(q_i, w_1) \vdash_E^* (q, \varepsilon)$ and $(q, w_2) \vdash_E^* (q_f, \varepsilon)$ hold for some $q_i \in I$ and $q_F \in F$. Any EFA can be converted to a trim equivalent EFA by removing all states that do not meet this requirement and adjusting the transition relation. A *nondeterministic finite automaton* with ε-transitions (FA) is an EFA whose transition relation is restricted to $\delta \subseteq Q \times (\Sigma \cup \{\varepsilon\}) \times Q$.

The graph *underlying* E is $\mathrm{G}(E) := (Q, \delta, t, h)$ where $t : (p,r,q) \mapsto p$ and $t : (p,r,q) \mapsto q$. Whenever we speak of the structure of an EFA, we are actually referring to its underlying graph. The following is easily verified:

Proposition 5 *An EFA E is trim and normalized* iff $\mathrm{G}(E)$ *is a hammock.*

An EFA displays a compromise between the complexity of its structure and the complexity of its transition labels. REs and FAs represent the extremes in this tradeoff: an RE can be considered as an EFA whose structure is trivial, namely P_1, while an FA is an EFA with trivial labels. Locally relaying information about a language between the structure of an EFA and its labels is the basis of several conversions between REs and FAs.

In this section, we consider conversions between regular expressions and trim normalized EFAs. According to Prop. 5 the latter are exactly the EFAs with hammock structure. Conversely, each hammock can be interpreted as an EFA by interpreting the set of arcs as an alphabet, and labeling every arc with itself.

Definition 5 Let (G, src, snk) be a hammock. The *automaton interpretation* of G is the EFA $\mathrm{A}(G) := (V_G, A_G, \delta_G, src, snk)$, where $\delta_G := \{(t(a), a, h(a)) \mid a \in A_G\}$.

Proposition 6 *Let* $(G, src, snk) \in \mathbf{H}$. *Then the following properties hold:*

1. $\mathrm{G}(\mathrm{A}(G)) = G$
2. $L(\mathrm{A}(G)) = \{w \mid w \text{ is a } (src, snk)\text{-walk in } G\}$

4.1 Expressions to Automata

We consider a fragment of the ARS proposed in [13]. Let E be an EFA with transition $\tau = (p, r, q)$ where r contains operators, then τ can be replaced depending on the root operator of r. If r is a product or a sum, the replacement is determined, while if r is an iteration, the out-degree of p and the in-degree of q are also considered. The rules of this ARS, denoted $\triangleleft_{\bullet}$, $\triangleleft_{+}$, and $\triangleleft_{*1}$ to $\triangleleft_{*4}$, are shown in Fig. 5.

In order to convert an RE into an FA, we identify $r \in \mathrm{reg}(\Sigma)$ with the EFA

$$A_r^0 := (\{q_i, q_f\}, \Sigma, \{(q_i, r, q_f)\}, \{q_i\}, \{q_f\}),$$

which trivially satisfies $L(A_r^0) = L_r$. The language accepted by an EFA is invariant under each rewriting, hence exhaustive application of $\triangleleft_{\bullet}$, $\triangleleft_{+}$ and $\triangleleft_{*i}$ yields a sequence $A_r^0, A_r^1, \ldots$ of equivalent EFAs terminating in an FA which we denote A_r. We find that each EFA in this sequence is structurally an spl-graph.

Lemma 1 *For all i,* $\mathrm{G}(A_r^i) \in \mathbf{SPL}$.

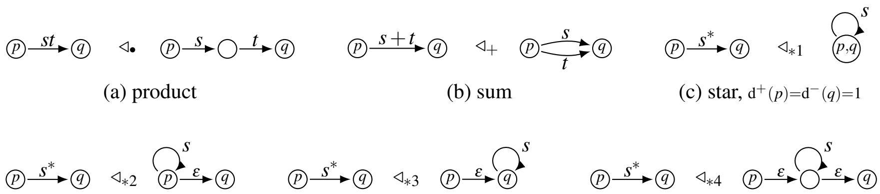

(a) product (b) sum (c) star, $d^+(p)=d^-(q)=1$

(d) star, $d^+(p)=1, d^-(q)>1$ (e) star, $d^+(p)>1, d^-(q)=1$ (f) star, $d^+(p)>1, d^-(q)>1$

Fig. 5: Replacing the transition (p,r,q) depending on r, $d^+(p)$, and $d^-(q)$. Applying the rule $\lhd_{*1}$ merges p and q. Applying $\lhd_{\bullet}$ and $\lhd_{*4}$ introduces a new state.

Proof The graph underlying A_r^0 satisfies $G(A_r^0) = \mathsf{P}_1$, which is the axiom of **SPL**. Assume $G(A_r^i) \in \mathbf{SPL}$ and consider the possible rewritings from A_r^i to A_r^{i+1}, where the transition $\tau = (p,r,q)$ is replaced.

For $\lhd_{\bullet}$, $\lhd_{+}$, and $\lhd_{*1}$, we find that $G(A_r^i)$ is s-, p-, and ℓ-expanded, respectively (observe that for $\lhd_{*1}$, the degree conditions on p and q state that τ is a constriction in $G(A_r^i)$). In each case, $G(A_r^{i+1}) \in \mathbf{SPL}$ follows.

Among the remaining cases, consider $A_r^i \lhd_{*2} A_r^{i+1}$. The difference between $G(A_r^i)$ and $G(A_r^{i+1})$ consists of an additional p-loop in the latter graph. This can be achieved with s- and ℓ-expansion, as follows. First, apply s-expansion to τ in $G(A_r^i)$. This replaces (the arc) τ with a new vertex z and two arcs $\tau_1 = pz$ and $\tau_2 = zq$. Since $\lhd_{*2}$ assumes $d^+(p) = 1$, while $d^-(z) = 1$ by construction, we find that τ_1 is a constriction. We may apply ℓ-expansion to τ_1 in this intermediate graph to get a graph that is isomorphic to $G(A_r^{i+1})$. Therefore, we get $G(A_r^i) \overset{s}{\Rightarrow}\overset{\ell}{\Rightarrow} Q \cong G(A_r^{i+1})$. Since Q was derived from an spl-graph by expansion, $Q \in \mathbf{SPL}$ follows, thus also $G(A_r^{i+1}) \in \mathbf{SPL}$.

A symmetric argument realizes the inductive step for $\lhd_{*3}$. For $\lhd_{*4}$ we need to expand $G(A_r^i)$ *twice* by s-expansion, to get a constriction that allows for ℓ-expansion. The details are left to the reader for this case. This shows that $G(A_r^{i+1}) \in \mathbf{SPL}$ holds, too, which concludes the induction. □

As a special case of Lem. 1, it follows that the graph underlying A_r is an spl-graph, too. It was shown in [13] that A_r is unique. We define the map α from simplified expressions to spl-FAs, by setting $\alpha(r) := A_r$ for $r = \mathrm{simp}(r)$.

4.2 Automata to Expressions

The spl-reductions are augmented to handle expression-labeled arcs which yields a second rewriting-system on EFAs. We assume trim normalized EFAs, whose underlying graphs are hammocks. Recall that any EFA can be converted to an equivalent trim normalized EFA. The labeled counterparts of $\overset{s}{\Leftarrow}$, $\overset{p}{\Leftarrow}$ and $\overset{\ell}{\Leftarrow}$ are denoted $\rhd_{\bullet}$, $\rhd_{+}$, and $\rhd_{*}$, respectively; they are shown in Fig. 6. Labeled reduction of an EFA requires that its underlying graph meets the preconditions for the corresponding unlabeled reduction, as defined in Def. 4. Observe that the language accepted by an EFA is invariant under labeled reduction.

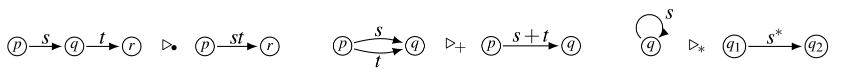

(a) labeled s-reduction, for simple q

(b) labeled p-reduction

(c) labeled ℓ-reduction, for non-guard q

Fig. 6: Labeled spl-reductions

Exhaustive reduction of a trim normalized EFA E terminates in an equivalent EFA which we denote $\mathrm{R}_l(E)$. The next result is a straightforward extension of Thm. 1 to labeled reductions on expression-labeled hammocks.

Proposition 7 *Let E be a trim normalized EFA. Then the EFA $\mathrm{R}_l(E)$ is unique up to associativity and commutativity in labels.*

In particular, for $\mathrm{G}(E) \in \mathbf{SPL}$ we find $\mathrm{G}(\mathrm{R}_l(E)) = \mathsf{P}_1$, so the only label of $\mathrm{R}_l(E)$ is an RE r with $L_r = L(E)$. By Prop. 7 this RE is unique up to trivialities, so we define a map β from EFAs with spl-structure to REs by first computing the unique label of $\mathrm{R}_l(E)$, which is then simplified. That is, we set $\beta(E) := \mathrm{simp}(r)$, where r is the label of $\mathrm{R}_l(E)$. For $G \in \mathbf{SPL}$, we call $\beta(\mathrm{A}(G))$ the *encoding* of G.

Lemma 2 *Let A be an spl-FA. Then $L(A)$ is denoted by an RE r with* $\mathrm{alph}(r) \leq |\delta_A|$.

Proof The sum of alphabetic widths of all labels in an EFA is invariant under labeled reduction (cf. Fig. 6). Let r_1 denote the unique label of $\mathrm{R}_l(A)$ for spl-FA A. Since every transition of A is labeled by a letter or ε, we find $\mathrm{alph}(r_1) = |\delta_A|$. We remove all ε-factors from r_1 to get $r = \mathrm{simp}(r_1)$: clearly, $\mathrm{alph}(r) \leq \mathrm{alph}(r_1)$ holds. Since $L_r = L_{r_1} = L(A)$, the claim follows. □

The performance of a conversion from FAs to REs is usually measured by comparing the size of an expressions relative to the number of *states* in an automaton. For spl-FAs, we get the following result:

Theorem 3 *Let A be an spl-FA. Then $L(A)$ is denoted by an RE r with*

$$\mathrm{alph}(r) < 4|Q_A|(|\Sigma_A| + 1).$$

Proof We bound the number of transitions in A relative to the number of states. We first assume that $\mathrm{G}(A)$ is p-normal and treat parallel transitions afterwards.

To bound the number of arcs in a p-normal spl-graph, we first consider how such a graph can be constructed from P_1. Obviously, P_1 is p-normal. Assume that G is p-normal, then H is p-normal for $G \overset{s}{\Rightarrow} H$ and $G \overset{\ell}{\Rightarrow} H$. Now consider an "intermediate" graph H', derived by $G \overset{p}{\Rightarrow} H'$. The only way to expand H' to a p-normal graph is by series expansion of either arc of the unique pair of parallel arcs.

For $\mathrm{c} \in \{s, p, \ell\}$, let $\#_{\mathrm{c}}$ denote the number of c-expansions in the construction of a p-normal H from P_1. The preceding discussion implies that $\#_p \leq \#_s$. Counting the number of arcs and vertices introduced or removed by each expansion (cf. Fig. 1), starting with $|A_{\mathsf{P}_1}| = 1$ and $|V_{\mathsf{P}_1}| = 2$, we arrive at

$$|A_H| = 1 + \#_s + \#_p \leq 1 + 2\#_s \text{ and}$$
$$|V_H| = 2 + \#s - \#_\ell.$$

Solving the second equation for $\#_s$ and substituting the result in the first inequation gets us

$$|A_H| \leq 2|V_H| + 2\#_\ell - 3 < 2(|V_H| + \#_\ell).$$

To bound $\#_\ell$, consider a vertex z that carries a loop introduced by ℓ-expansion at some point in the rewriting $\mathsf{P}_1 \Rightarrow^\star H$. On its first appearance, $\mathrm{d}^-(z) \geq 2$ and $\mathrm{d}^+(z) \geq 2$ hold. Since no later expansion decreases either degree, z is never incident to a constriction and thus still present in H. Each ℓ-expansion gives rise to a unique such vertex, which implies that $\#_\ell \leq |V_H|$.

We arrive at $|A_H| < 4|V_H|$ for a p-normal spl-graph H. Getting back to arbitrary $H \in \mathbf{SPL}$, observe that it is irrelevant for H where its parallel arcs emerge in the expansion. So we may assume wlog. that the p-normal "skeleton" of H is derived first, followed by a sequence of p-expansions.

While for arbitrary spl-graphs, the number of arcs is unbounded, this is different for spl-FAs. For any FA over Σ, the number of parallel transitions from p to q, for fixed p, q, is bounded by $|\Sigma_A| + 1$. For an spl-FA A we correct the derived bound for $\mathrm{G}(A)$, by that factor, to find

$$|\delta_A| < 4|Q_A|(|\Sigma_A| + 1).$$

The claim now follows with Lem. 2. □

4.3 Duality of the Conversions

The presented conversions between REs and εNFAs with spl-structure are duals modulo associativity and commutativity of the regular operators. These are trivial matters, so we write $r = r'$, if the REs r and r' identical up to these two equivalences.

Theorem 4

1. *Let r be a simplified RE. Then $r = \beta(\alpha(r))$*
2. *Let A be an spl-FA. Then $A = \alpha(\beta(A))$*

Proof The ARSs that realize α and β are both locally confluent. To prove our claim, it is thus sufficient to show that each rewriting step in either ARS can be reverted in the other ARS. This is immediately clear for product and sum replacements, which are reverted by labeled s-reduction and p-reduction, respectively, and vice versa. Labeled ℓ-reduction and the first star replacement, $\lhd_{*1}$, also mutually revert another.

No $\lhd_{*i}$ is reverted by labeled reduction, for $i \in \{2,3,4\}$. This is due to the fact that ε-transitions are introduced by each $\lhd_{*i}$. However, it is irrelevant for $\beta(A)$ if simplifying the label of $\mathrm{R}_l(A)$ is realized as defined, or by successively simplifying its subexpression when they appear as labels in the course of the rewriting. Therefore, $A \lhd_{*2} B$, applied to $\tau = (p, s^*, q)$, is reverted by $B \rhd_* \rhd\!. A'$, followed by simplifying the label of the transition $\tau' = (p, s^*\varepsilon, q)$. The argument is analogous for $\lhd_{*2}$ and $\lhd_{*3}$. □

Thus the encoding of labeled spl-graphs by simplified expressions is unique and reversible, and every simplified expression encodes a labeled spl-graph. Hence every RE over a non-empty alphabet encodes an spl-graph and every spl-graph can be encoded. Informally, we state

Corollary 2 $G \in \mathbf{SPL}$ iff *G can be encoded by an RE.*

5 Characterizing SPL by Forbidden Minors

We extend the concept of a topological minor, which is well-known for undirected graphs, to the directed case. For undirected graphs, this relates to subgraphs which are subdivisions, as opposed to the more general notion of a minor, which also incorporates the merging of adjacent vertices (see, e.g. [5]). As this distinction is not necessary for us, we drop the modifier "topological".

Definition 6 An *embedding* of F in G is an injection $e : V_F \to V_G$ satisfying

1. for $a = xy \in A_F$, G contains an $(e(x), e(y))$-path P_a, and
2. for distinct $a, a' \in A_F$, the paths P_a and $P_{a'}$ are internally disjoint.

If e is an embedding of F in G, we write $e : F \preccurlyeq G$. We write just $F \preccurlyeq G$ if some e with $e : F \preccurlyeq G$ exists and call F a *minor* of G. We say that $F \preccurlyeq G$ is *realized* by e, and for $x \in V_F$ we call $e(x)$ the *peg* of x in G wrt. e.

If $\mathbf{M}$ is a set of graphs and $M \preccurlyeq G$ for some $M \in \mathbf{M}$ we also say that G has a minor in $\mathbf{M}$. If M is no minor of G, then G is *M-free*, and if G is M-free for each $M \in \mathbf{M}$, then G is $\mathbf{M}$-free. An equivalent characterization of minors is given by means of subdivisions.

Proposition 8 $F \preccurlyeq G$ iff *G contains a* $\mathrm{D}F$

This allows us to choose the notion that fits our purposes. We observe that the in- and out-degree of a peg is no less than the respective degree of the peg's preimage.

Proposition 9 *Let $e : F \preccurlyeq G$ and $x \in V_F$. Then* $\mathrm{d}^-_F(x) \le \mathrm{d}^-_G(e(x))$ *and* $\mathrm{d}^+_F(x) \le \mathrm{d}^+_G(e(x))$.

A stricter notion of embedding models the absence of an arc in a minor.

Definition 7 Let $e : F \preccurlyeq G$ and assume $xy \notin A_F$. An (x,y)-*bypass* wrt. e is an $(e(x), e(y))$-path in G that does not pass through any other peg of F. If G contains no bypass wrt. e, we call e a *bare embedding* of F in G, and F a *bare minor* of G, denoted $e : F \preccurlyeq_{\mathrm{b}} G$.

If F is a bare minor of G, then G contains a $\mathrm{D}F$ whose pegs are not connected by bypasses in G. We call this $\mathrm{D}F$ a bare $\mathrm{D}F$.

5.1 Effects of Expansion and Reduction on Minors

We study the effects of spl-expansion and -reduction on two sets of graphs. We will find that these sets are, in a sense, "minor stable" under spl-operations.

Definition 8 A *bulky* graph contains no simple vertices, parallel arcs or loops. The class of bulky graphs is denoted $\mathbf{B}$.

Notice that a vertex x is not simple *iff* $\mathrm{d}^-(x) \geq 2$ or $\mathrm{d}^+(x) \geq 2$. Bulky graphs are essentially defined by prohibiting the characteristic features that come with spl-expansion. We thus find that $\mathbf{B}$-minors are not introduced by expansion alone, i.e., coming from an otherwise $\mathbf{B}$-free graph.

Lemma 3 *Assume that G is $\mathbf{B}$-free and let $G \Rightarrow H$. Then H is $\mathbf{B}$-free.*

Proof Let G be $\mathbf{B}$-free and assume $G \Rightarrow H$. For the sake of contradiction, suppose $e : B \preccurlyeq H$ for some $B \in \mathbf{B}$. Since G is $\mathbf{B}$-free, the DB in H is the result of expansion. We consider each expansion separately.

1. $G \overset{s}{\Rightarrow} H$: Let z be the vertex that is introduced by expansion and notice that z is simple. Since G is B-free, a DB in H requires that $z = e(q)$ for some $q \in V_B$. But q is not simple, i.e., $\mathrm{d}^-_B(q) \geq 2$ or $\mathrm{d}^+_B(q) \geq 2$, so this contradicts Prop. 9.
2. $G \overset{p}{\Rightarrow} H$: Every peg of B in H is also a vertex in G. If the mere addition of a parallel arc yields a DB, this implies that B contains parallel arcs, which it does not.
3. $G \overset{\ell}{\Rightarrow} H$: Let $a = xy$ denote the constriction in G that allows ℓ-expansion and let $l = zz$ be the loop introduced by merging x and y, i.e., let z denote the merge vertex of x and y. Again, z is a peg, so $z = e(q)$ for some $q \in V_B$. Otherwise a DB would be present in G already, as implied by Prop. 2, but contradictory to the assumption that G is $\mathbf{B}$-free.
   Consider the graph $H' := H \setminus l$, derived from H by removal of the "new" loop. As B is free of loops, l does not belong to the DB in H. Consequently, H' contains the same DB as H, with $z = e(q)$. The advantage of H' over H is that $\mathrm{d}^-_{H'}(z) = \mathrm{d}^-_G(x)$ and $\mathrm{d}^+_{H'}(z) = \mathrm{d}^+_G(y)$ hold.
   We show $\mathrm{d}^-_B(q) \geq 2$ by rejecting the other possibilities. First, suppose $\mathrm{d}^-_B(q) = 0$. Then, only the out-arcs of z belong to $e(B)$ in H'. From $\mathrm{d}^+_{H'}(z) = \mathrm{d}^+_G(y)$ follows that G contains a DB, too. This is realized by $e' : V_B \to V_G$, which is defined as e except that $e'(q) = y$. This contradicts $\mathbf{B}$-freeness of G. If we suppose $\mathrm{d}^-_B(q) = 1$, the argument is similar. In that case, the only path in $e(B)$ that enters z in H' can be realized in G (for a slightly different DB, namely, a subdivision of $e(B)$) by using the constriction a. With e' defined as before, $e' : B \preccurlyeq G$ follows, contradicting the assumption that G is $\mathbf{B}$-free.
   Thus follows $\mathrm{d}^-_B(q) \geq 2$, and a symmetric argument shows $\mathrm{d}^+_B(q) \geq 2$. Now let B' denote the graph derived from B by splitting q into q_1 and q_2. Then $\mathrm{d}^-_{B'}(q_1) \geq 2$ and $\mathrm{d}^+_{B'}(q_2) \geq 2$, while the degrees of the other vertices in B and B' agree in these graphs. Therefore B' is bulky, too. But since H contains a DB it now follows that G contains a DB', contradicting the assumption that G is $\mathbf{B}$-free.

□

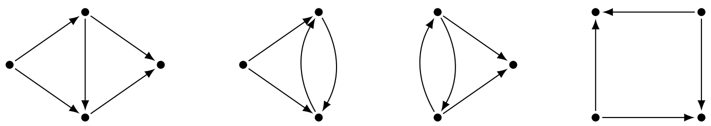

(a) N (b) C (c) C^R (d) Q

Fig. 7: Crucial subset of bulky graphs, $\mathbf{F} = \{\mathsf{N}, \mathsf{C}, \mathsf{C}^\mathsf{R}, \mathsf{Q}\}$.

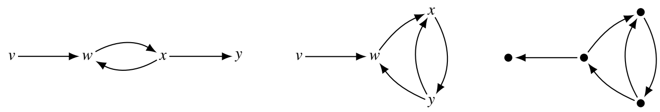

(a) Φ (b) Ψ (c) Ψ^R

Fig. 8: Set of graphs $\mathbf{K} = \{\Phi, \Psi, \Psi^\mathsf{R}\}$. The vertices of Φ and Ψ will always be referred to as shown.

Corollary 3 *Every spl-graph is* **B**-*free.*

Proof Since P_1, the axiom of **SPL**, is **B**-free, the claim follows inductively. □

For the purpose of characterizing **SPL**, it suffices to consider four bulky graphs, which constitute the set

$$\mathbf{F} := \{\mathsf{N}, \mathsf{C}, \mathsf{C}^\mathsf{R}, \mathsf{Q}\},$$

shown in Fig. 7. Observe that **F** is closed under arc-reversal: C and C^R are mutually dual, whereas both N and Q are self-dual.

It was shown by Valdes [23] that that the acyclic spl-graphs are characterized by means of N alone.

Theorem 5 **(Valdes)** *Let* $H \in \mathbf{H}$ *be acyclic. Then* $H \in \mathbf{SPL}$ iff $\mathsf{N} \not\preccurlyeq H$.

We make use of Thm. 5 whenever acylic graphs are considered. Getting back to the general case, we observe that the absence of **F**-minors is not sufficient to identify a hammock as a member of **SPL**. For example, the hammock shown in Fig. 3 is **F**-free, yet, as was already shown, it is no spl-graph, as it violates Prop. 3.

To get an adequate obstruction set characterization of **SPL**, we resort to bare embeddings of non-bulky graphs. The three additional graphs that will be part of the sought characterization constitute the set

$$\mathbf{K} := \{\Phi, \Psi, \Psi^\mathsf{R}\},$$

shown in Fig. 8. Observe that **K** is closed under arc-reversal, too. We find that spl-graphs are free of bare minors in **K**.

Lemma 4 *Every spl-graph is free of bare* **K**-*minors.*

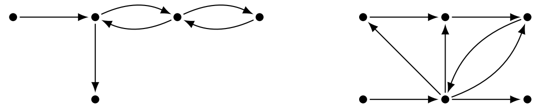

(a) $\Phi \preccurlyeq G$ (b) $\Psi \preccurlyeq G'$

Fig. 9: Members of **SPL** with minors in **K**.

Proof Let $(G, src, snk) \in \mathbf{SPL}$, and let $e : K \preccurlyeq G$ for $K \in \{\Phi, \Psi\}$. For Ψ^{R}, the proof is symmetric to that for Ψ. We refer to the vertices of Φ and Ψ as showin in Fig. 8.

Since w and x lie on a cycle in Φ, their pegs $e(w)$ and $e(x)$ lie on a cycle C in G. Then Prop. 3 states that C is guarded by a unique $g \in V_C$. Moreover, G contains an $(e(x), e(y))$-path. But since g co-dominates $e(x)$, it also co-dominates $e(y)$: thus G contains an $(e(y), g)$-path P, and since g lies on a cycle with $e(w)$ and $e(x)$, G further contains $(e(y), e(w))$- and $(e(y), e(v))$-paths. But either path is a bypass, since $yw, yx \notin A_\Phi$. It follows that e is not bare.

For Ψ, notice that this graph has two cycles that share an xy-arc. The cycle $C \subseteq G$ that contains the pegs of w, x, y is guarded by a unique $g \in V_C$. In particular, g guards each vertex on every $(e(x), e(y))$-path, as well as on every $(e(y), e(x))$-path. Hence g guards the "image" of the inner cycle of Ψ which contains $e(x)$ and $e(y)$ but not $e(w)$. It follows that g is also a vertex on this inner cycle in G. Now since g dominates $e(w)$ it also dominates $e(v)$, so G contains paths from $e(w)$, $e(x)$, and $e(y)$, to $e(v)$. Since v has no in-arc in Ψ, the embedding e is not bare. □

Informally, **K** consists of the prototypical graphs with cycles that do not satisfy Prop. 3. Adding certain arcs to Φ, Ψ, and Ψ^{R} would mend this deficiency. Regarding the bare minor relation, such arcs—resp. their subdivisions—remain absent in the disguise of prohibited bypasses. Examples for spl-graph with (non-bare) minors in **K** are given in Fig. 9. The reader might want to verify that these graphs are indeed spl-graphs.

Next, we consider the effect of reduction on **F**-free hammocks. We will find that the members of **F** "behave" quite differently in that respect. While the presence or absence of a DC or DC$^{\mathsf{R}}$ in a hammock is invariant under reduction, this is not the case for N and Q. For N, we find that a DN might be removed by reduction, in which case, however, a DC and a DC$^{\mathsf{R}}$ must be present. For Q, a similar property holds, which also involves Ψ and Ψ^{R}.

In the following lemma, these properties are stated in a contrapositive manner. Therein, we consider a hammock H and its reduct G, and trace an **F**-minor of G back to an **F**-minor or a bare **K**-minor of H.

Lemma 5 *Let $H \Leftarrow G$, then an* **F***-minor of G implies an* **F***-minor or a bare* **K***-minor of H, as follows:*

1. *if $F \preccurlyeq G$, then $F \preccurlyeq H$ for $F \in \{\mathsf{C}, \mathsf{C}^{\mathsf{R}}\}$*
2. *if $\mathsf{N} \preccurlyeq G$, then $(\mathsf{N} \preccurlyeq H \vee (\mathsf{C} \preccurlyeq H \wedge \mathsf{C}^{\mathsf{R}} \preccurlyeq H))$*
3. *if $\mathsf{Q} \preccurlyeq G$, then $(\mathsf{Q} \preccurlyeq H \vee \mathsf{C} \preccurlyeq H \vee \mathsf{C}^{\mathsf{R}} \preccurlyeq H \vee \Psi \preccurlyeq_{\mathsf{b}} H \vee \Psi^{\mathsf{R}} \preccurlyeq_{\mathsf{b}} H)$*

Proof We prove the claim separately for each reduction. For s- and p-reduction we show a stronger property for bulky graphs. The argument is downright trivial for these two reductions. For ℓ-reduction, however, we need to consider several subcases that lead to the details of the claim.

- $H \overset{s}{\Leftarrow} G$: Let $B \in \mathbf{B}$ be arbitrary and observe that H is a DG. Therefore, if G contains a DB, so does H. But since H is B-free by assumption, G is B-free, too.
- $H \overset{p}{\Leftarrow} G$: Let $B \in \mathbf{B}$ be arbitrary. Removing an arc from H does certainly not *introduce* any subgraph at all. This goes for a DB in particular.
- $H \overset{\ell}{\Leftarrow} H$: Let $l = xx$ denote the loop in H that allows for reduction, and let $a = x_1x_2$ be the constriction in G that results from splitting x. Let $e : F \preccurlyeq G$ for some $F \in \mathbf{F}$ distinguish by which of the x_i are pegs of F in G. As a visual aid we sketch the relevant parts of H and G below:

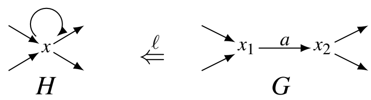

  If neither x_i is a peg, then each peg of F occurs in both G and H. It follows from Prop. 2 that $e : F \preccurlyeq H$.
  Next assume $x_1 = e(q)$, for some $q \in V_F$, while x_2 is no peg (or vice versa, which is symmetric). Let $e' : V_F \to V_H$ be as e, except that q is mapped to x in H. Now H contains an $(e'(q_1), e'(q_2))$-path for $q_1q_2 \in A_F$, where $q \neq q_i$, and these paths are pairwise internally disjoint. This follows from Prop. 2 and the fact the e is an embedding. Moreover, H contains an $(e'(q_1), x)$-path for each q_1q-arc of F, and an $(x, e'(q_2))$-path for each qq_1-arc. Again we find that these paths are internally disjoint. Therefore, H contains a DF and $e' : F \preccurlyeq H$.
  Finally, let both x_1 and x_2 be pegs of F in G, with preimages q_1 and q_2, respectively. Then Prop. 9 requires that $\mathrm{d}^+_F(q_1) \leq 1$ and $\mathrm{d}^-_F(q_2) \leq 1$ hold. We proceed by case distinction for F as in the claim.

  1. Assume that $F = \mathsf{C}$: there is only one vertex q_2 that satisfies $\mathrm{d}^-_{\mathsf{C}}(q_2) \leq 1$ (cf. Fig. 7b). The vertex q_1 might be either vertex of the cycle of C, as each satisfies $\mathrm{d}^+_{\mathsf{C}}(q_1) \leq 1$. We fix one of these two vertices as q_1 and denote the other as y. Since C contains a q_1y-arc, G must contain an $(e(q_1), e(y))$-path. But the only out-arc of $e(q_1)$, which is a, is an in-arc of $e(q_2)$. In other words, $e(q_2)$ lies on every path from $e(q_1)$ to $e(y)$. Consequently, the paths in G that represent the arcs of C are not internally disjoint, which contradicts the assumption that e is an embedding. This conclusion follows symmetrically for $F = \mathsf{C}^{\mathsf{R}}$.
  2. For $F = \mathsf{Q}$, the degree restrictions on q_1 and q_2 are satisfied by two vertices each. The choices for the q_i lead to symmetric cases, due to the symmetry of Q. We find that Q contains a q_2q_1-arc, therefore G contains a path from x_2 to x_1. Observe that $H = G[x_1 = x_2]$, with merge vertex x. Thus the arcs that constitute the (x_2, x_1)-path in G form a cycle C in H, with $x \in V_C$. Since ℓ-reduction is applicable in x, C is not merely a further x-loop, i.e., C contains at least one more vertex y. The situation in G and H is sketched below; we denote the remaining pegs of Q in G as k_1 and k_2; these vertices are also present in H.

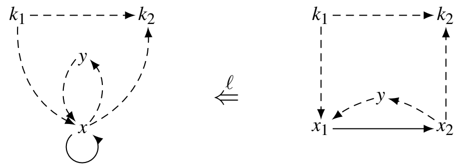

Since x does not guard C, there is a further vertex on C that is not guarded—i.e., not dominated or co-dominated—by x. We assume wlog. that this vertex is y and that x does not dominated y (the case where x does not co-dominate y is symmetric). Now Prop. 1 implies that either y dominates x, or some z dominates x and y, and H contains internally disjoint (z,x)- and (z,y)-paths.

(a) Assume that y dominates x, then y also dominates k_1, since H contains a (k_1,x)-path that does not pass through y. So there is also a (y,k_1)-path in H. If any such path does not pass through k_2, we find a DC^{R} in H with pegs x, y, and k_2.
On the other hand, if each (y,k_1)-path passes through k_2, then there is a (k_2,k_1)-path in H. We may assume that this path is internally disjoint with the (k_1,k_2)-path sketched above. Otherwise, we choose some k_1' that is an intersection vertex of the two paths and argue with k_1' instead of k_1. So k_1 and k_2 lie on a cycle. Moreover, y then also dominates k_2. Hence, there is a path from the source of H to y that does not intersect with any of the paths we are considering right now. Together, these paths form a $\mathsf{D}\Psi$ with the source of H being a peg. If this $\mathsf{D}\Psi$ is bare, the claim follows. Otherwise, any bypass to the $\mathsf{D}\Psi$ implies the existence of a DC or a DC^{R}.

(b) Assume that there is some z with internally disjoint (z,x)- and (z,y)-paths in H. Since x and y lie on C, the paths from z enter C in distinct vertices. Thus follows $\mathsf{C} \preccurlyeq H$.

3. For $F = \mathsf{N}$, consider an appropriate DN in G. Since a is a constriction, it does not represent the trivial subdivision of any arc of N, for N is free of constrictions. By the same argument, it follows that a is not anti-parallel to an arc of N.
Therefore, q_1 and q_2 are not adjacent in N, which is only satisfied by the source and sink of N (notice that N is a hammock). The degrees of x_1 and x_2 and Prop. 9 determine that q_2 is the source of N and that q_1 is the sink. Thus G contains a "back arc" in its DN subgraph, as sketched below:

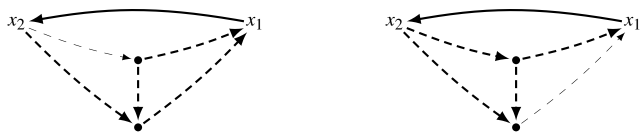

This yields both a DC and a DC^{R}, i.e., $\mathsf{C} \preccurlyeq G$ and $\mathsf{C}^{\mathsf{R}} \preccurlyeq G$. As we have shown in the first case of the overall proof, $\mathsf{C} \preccurlyeq H$ and $\mathsf{C}^{\mathsf{R}} \preccurlyeq H$ follow.

We have shown the for reductions $H \stackrel{s}{\Leftarrow} G$ and $H \stackrel{p}{\Leftarrow} G$, $F \preccurlyeq G$ implies $F \preccurlyeq G$ for $F \in \mathbf{F}$. Considering the reduction $H \stackrel{\ell}{\Leftarrow} G$, a $\mathbf{F}$-minor of G implies an $\mathbf{F}$-minor of H or a bare $\mathbf{K}$-minor of H as broken down in the claim. Thus the claim holds and the proof is concluded. □

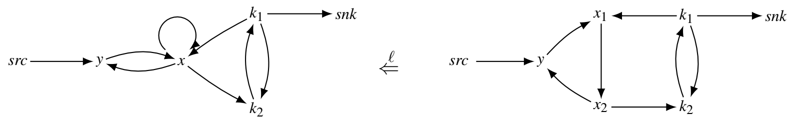

Fig. 10: A DQ emerging with ℓ-reduction. Either side contains a bare DΨ.

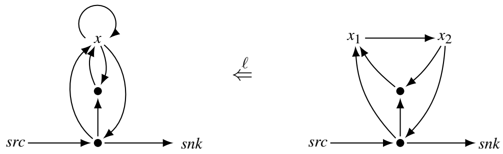

Fig. 11: A DN emerging with ℓ-reduction. Either side contains a DC and a DC^{R}.

Examples for the emergence of a DQ and a DN upon ℓ-reduction are show in Figs. 10 and 11, respectively.

Next we consider the effects of reduction on hammocks that *do* have a minor in **F**. We find that minors, resp. subdivisions, from **F** are not removed upon reduction.

Lemma 6 *Let $H \Leftarrow G$, then $F \preccurlyeq H$ implies $F \preccurlyeq G$ for $F \in \mathbf{F}$.*

Proof We know that a reduction allows for the complementary expansion, i.e., $H \Leftarrow G$ implies $G \Rightarrow H$. Assume now that G is **F**-free and reconsider Lem. 3, resp. its proof: therein we have shown that $B \preccurlyeq G$ implies $B \preccurlyeq H$ for bulky B and $G \overset{s}{\Rightarrow} H$ or $G \overset{p}{\Rightarrow} H$. Since $\mathbf{F} \subseteq \mathbf{B}$, this holds for $B \in \mathbf{F}$ as well. For $G \overset{\ell}{\Rightarrow} H$, we have shown that $B \preccurlyeq G$ yields $B' \preccurlyeq H$, for bulky B, B', and indirectly, that $\mathrm{d}^-_B(v) < 2$ or $\mathrm{d}^+_B(v) < 2$ for some $v \in V_B$ implies $B = B'$. This degree property is satisfied by each $F \in \mathbf{F}$, so $F \preccurlyeq G$ implies $F \preccurlyeq H$ for ℓ-reduction, too.

Thus, if $H \Leftarrow G$, the assumption that H contains a minor in **F**, while G is **F**-free, is contradictory. Equivalently, the claim follows. □

As for minors in **F**, we study the effects of spl-reductions on minors in **K**. We first investigate if—or rather "how"—such minors can be the result of a reduction.

Lemma 7 *Let $H \Leftarrow G$, then $K \preccurlyeq_{\mathrm{b}} G$ implies $K \preccurlyeq_{\mathrm{b}} H \vee \mathsf{C} \preccurlyeq H \vee \mathsf{C}^{\mathsf{R}} \preccurlyeq H$ for $K \in \mathbf{K}$.*

Proof Observe that each $K \in \mathbf{K}$ is free of parallel arcs. Let $e : K \preccurlyeq_{\mathrm{b}} G$ and recall that if all pegs of K in G also occur in H, i.e., if $e(q) \in V_H \cap V_G$ for $q \in V_K$, then the connectivity among the pegs of K is invariant under reduction (Prop. 2). This holds for paths representing arcs of K, as well as for bypasses. In other words, a bypass is not removed upon reduction if its endpoints occur in $V_H \cap V_G$, so a nonbare embedding does not "become" bare. It is thus sufficient to consider cases where pegs are introduced with $G \Leftarrow H$. Notice that the removal of a peg is irrelevant to this proof.

Since $V_H = V_G$ holds for $H \overset{p}{\Leftarrow} G$, nothing needs to be done for this case, following the discussion above. Regarding $H \overset{s}{\Leftarrow} G$, a peg is removed at most, so the statement follows for this case, too.

It remains to consider the reduction $H \overset{\ell}{\Leftarrow} G$, where the number of vertices increases. To this end, let $l = qq$ be the loop in H that allows for reduction, and let $a = q_1q_2$ denote the constriction that emerges from reduction. We need to consider the cases where at least one of the q_i is a peg of K.

Let $q_1 = e(k)$ and let q_2 not be a peg. Since $\mathrm{d}^+_G(q_1) = 1$, at most one arc leaves k in K. Then $K \preccurlyeq H$ is realized by e', which is as e, except that $e'(k) = q$, the in-arcs of k in K are mapped to in-paths of q and the sole out-arc, if it exists, is mapped to an out-path of q. Since e is bare, no bypass possibly intersects the out-path of q_1 by assumption, in particular not in q_2. Hence no bypass is accidentally "removed" by using e' as an embedding, and $e' : K \preccurlyeq_{\mathrm{b}} H$ follows.

If q_1 and q_2 are both pegs wrt. e, some additional effort is required. In this case, we treat Φ and Ψ explicitly; for Ψ^{R}, the argument is symmetric to that for Ψ. In the following, the vertices of Φ and Ψ are denoted as in the proof of Lem. 4.

1. We start with $K = \Phi$. Since v and y are not adjacent in Φ, the case where the q_i are pegs of these two vertices immediately yields a as a bypass in G. As a bare embedding, e maps at least one of w and x to a q_i. We assume wlog. $q_1 = e(w)$ and examine the preimage of q_2.
   If $q_2 = e(v)$ or $q_2 = e(y)$, it follows immediately that e is not a bare embedding, since a constitutes a bypass.
   If $q_2 = e(x)$, then G contains a nonempty (q_2, q_1)-path P, i.e., an $(e(x), e(w))$-path. The length of this path at least two: otherwise, the arc that constitutes P would be parallel to l in G, meaning that ℓ-reduction is not applicable. So let z be an internal vertex of P, then z occurs in H, too; more specifically, H contains a (q,z)- and a (z,q)-path. Since ℓ-reduction is applicable to l, the vertex q does not guard z in H. This means that H contains a (src, z)-path or a (z, snk)-path that does not pass through q. We proceed with the first possibility and let P_z denote this path. The part of H and G that we are considering looks as follows:

   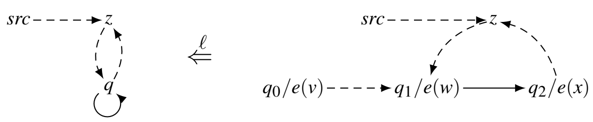

   Let P_{q_0} by a (src, q_0)-path in G, where q_0 is the peg of v. Notice that if $q_0 = src$ holds in G, then $e' : \Phi \preccurlyeq_{\mathrm{b}} H$ holds for $e'(w) = z$, $e'(x) = q$ and letting e' otherwise be identical to e. In that case, the claim follows, so assume $q_0 \neq src$, which holds for both G and H. If P_{q_0} passes through neither q_1 nor q_2, we find $\mathsf{C} \preccurlyeq G$ and, following Lem. 5, also $\mathsf{C} \preccurlyeq H$. On the other hand, if P passes through q_1 or q_2, then G contains a bypass to the uncovered $\mathsf{D}\Phi$. This contradicts the assumption that e is a bare embedding. If we choose P_z as a (z, snk)-path that does not pass through q, and P_{q_0} as a (q_0, snk)-path, a symmetric argument yields either $\mathsf{C}^{\mathrm{R}} \preccurlyeq H$ or that e is not bare. This concludes the argument for Φ.
2. Now let $K = \Psi$. Since e is bare by assumption, we find $q_2 \neq e(v)$ immediately, because v has no in-arc in Ψ, yet q_1 is a peg, too. If we assume $q_1 = e(v)$, we construct an embedding e' that realizes $\Phi \preccurlyeq_{\mathrm{b}} H$: e' maps the preimage of q_2 to q and v to any predecessor of q; otherwise e' is as e. Since v has no in-arcs and e

is bare, e' is bare too. This covers the case that either q_i is the peg of v, and we consider the remaining cases.
Let us first assume $q_1 = e(x)$ and $q_2 = e(y)$. This case is akin to the nontrivial case for Φ, which we discussed at length above. Similarly, we find that G contains a (q_2, q_1)-path, which, due to ℓ-reduction being applicable to H, consists of at least two arcs. Again, let z denote a vertex on this path and notice that q does not guard z in H. As before, assume wlog. that H contains a path from its source to z that does not pass through q. The relevant parts of H and G are shown below:

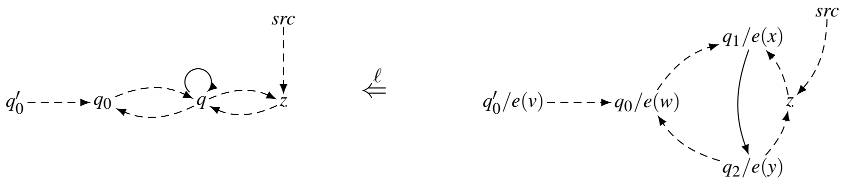

Now let $e(v) = q'_0$ and consider a (src, q'_0)-path P in G. If none of q_0, q_1, or q_2 lies on P, then $\mathsf{C} \preccurlyeq G$ follows, which further yields $\mathsf{C} \preccurlyeq H$. Otherwise, a segment of P constitutes a bypass to the DΨ realized by e, contradicting the assumption that e is bare. Again, if we consider paths to *snk* in G, instead of paths from *src*, we find $\mathsf{C}^{\mathsf{R}} \preccurlyeq H$ symmetrically.
The final case for Ψ is that one q_i is the peg of w while the other is the peg of either x or y. The argument is symmetric if x and y are swapped; we pursue the argument for x. First, let $q_1 = e(w)$ and $q_2 = e(x)$. Since we have $\mathrm{d}^-_{\Psi}(x) = 2$, Prop. 9 implies $\mathrm{d}^-_G(q_2) \geq 2$. This, however, is infeasible with ℓ-reduction, which produces a constriction, i.e., $\mathrm{d}^-_G(q_2) = 1$. On the other hand, if we assume $q_1 = e(x)$ and $q_2 = e(w)$, then the $q_1 q_2$-arc already represents an $(e(x), e(w))$-bypass, so e is not bare.

This concludes the proof. □

Next, we examine the effects of reducing a hammock that has a bare **K**-minor. We find that there are at least some "traces" of this minor in the emerging graph.

Lemma 8 *Let $H \Leftarrow G$, then $K \preccurlyeq_{\mathrm{b}} H$ implies $K \preccurlyeq_{\mathrm{b}} G$ or $F \preccurlyeq G$ for $K \in \mathbf{K}$, $F \in \mathbf{F}$, as follows:*

- *if $\Phi \preccurlyeq_{\mathrm{b}} H$, then $(\Phi \preccurlyeq_{\mathrm{b}} G \vee \mathsf{C} \preccurlyeq G \vee \mathsf{Q} \preccurlyeq G)$*
- *if $\Psi \preccurlyeq_{\mathrm{b}} H$, then $(\Psi \preccurlyeq_{\mathrm{b}} G \vee \mathsf{C} \preccurlyeq G \vee \mathsf{N} \preccurlyeq G)$*
- *if $\Psi^{\mathsf{R}} \preccurlyeq_{\mathrm{b}} H$, then $(\Psi^{\mathsf{R}} \preccurlyeq_{\mathrm{b}} G \vee \mathsf{C}^{\mathsf{R}} \preccurlyeq G \vee \mathsf{N} \preccurlyeq G)$*

Proof As for Lem. 7, the statement follows immediately from Prop. 2 if all pegs of K occur in $V_H \cap V_G$. In particular, nothing needs to be done for $H \overset{p}{\Leftarrow} G$. In the following, assume $e : K \preccurlyeq_{\mathrm{b}} F$ for $K \in \mathbf{K}$. Again, we only consider Φ and Ψ.

For $H \overset{\ell}{\Leftarrow} G$ let $l = xx$ be the loop that is reduced in H and let $a = x_1 x_2$ denote the constriction introduced in G. Observe that each $k \in V_K$ satisfies $\mathrm{d}^-_K(k) \leq 1$ or $\mathrm{d}^+_K(k) \leq 1$. Thus, if $x = e(k)$, we define $e' : V_K \to V_G$ with either $e'(k) = x_1$ or $e'(k) = x_2$, depending on the degrees of k. The connectivity among the pegs of K in G then is the same as in H, so $e' : K \preccurlyeq_{\mathrm{b}} G$ follows.

In the case $H \overset{s}{\Leftarrow} G$, let z denote the simple vertex of H that allows for reduction. Assume that z is a peg of K. Since z is simple in H, Prop. 9 implies that the preimage of z satisfies $\mathrm{d}_K^-(e^{-1}(z)) \leq 1$ and $\mathrm{d}_K^+(e^{-1}(z)) \leq 1$.

There are two vertices with this property in Φ and one in Ψ. We choose v as the preimage of z for either graph. Observe that the in-degree of v is zero. Since z is simple, H contains a unique predecessor z' of z. Moreover, since e is a bare embedding, H contains no path from any other peg to z. This implies that z' is not also be peg of K wrt. e, and that there is no path from any other peg of K to z'. The embedding e' of K into H is now defined by mapping v to z' and being otherwise as e. If e' is bare, the claim follows. Otherwise, we observe that since no bypass wrt. e enters z, no bypass wrt. e' enters z'. So if e' is not bare, a bypass must leave $e'(v) = z'$; the connectivity among all other pegs does not differ in e and e'.

There are two possibilities for how a bypass like that in $K = \Phi$ in H, they are sketched below:

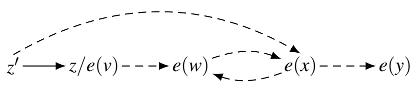

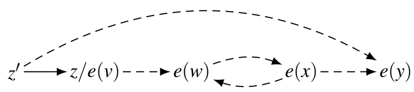

In the case on the left, we find a DC, while on the right, we find a DQ. Thus follows the claim for $K = \Phi$. For $F = \Psi$, the possibilities in H are as follows:

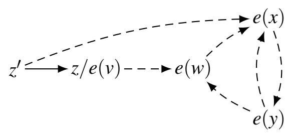

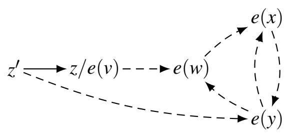

In either case, we find a DC, while in the case on the right, we also find a DN. This proves the claim for Ψ, and the statement for Ψ^{R} follows by symmetry. □

We have thus found that $\mathbf{F} \cup \mathbf{K}$ behaves stable wrt. being (bare) minors in hammocks under spl-operations. Although such minors are exchanged for another in certain cases, they never appear by manipulating **F**- and **K**-free graph; likewise, they never disappear altogether. By omitting the finer details of the previous lemmas, we arrive at a cornerstone result regarding the characterization of **SPL** by forbidden minors.

Theorem 6 *Let $H_1, H_2 \in \mathbf{H}$ s.t. $H_1 \Rightarrow H_2$ or $H_1 \Leftarrow H_2$. Then H_1 is free of* **F**-*minors and bare* **K**-*minors* iff *H_2 is.*

Most importantly, we can associate (bare) minors in an arbitrary hammock with (bare) minors in its spl-normal form.

Corollary 4 *Let $H \in \mathbf{H}$, then H is free of* **F**-*minors and bare* **K**-*minors* iff $\mathrm{R}(H)$ *is.*

It remains to show that **F** and **K** sufficiently describe the structures that are absent in spl-graphs. In other words, we need to prove that $\mathrm{R}(H) \neq \mathsf{P}_1$ implies a (bare) minor of these classes in $\mathrm{R}(H)$ and therefore in H. This will be done in the following subsection.

5.2 Sufficiency of **F** and **K**

Definition 9 A *kebab* is a graph consisting of three arc-disjoint subgraphs: a strong component B, called the *body*, and two nonempty, disjoint paths S_1 and S_2, called the *spikes* of the kebab. Exactly one endpoint of either spike is a vertex of B.

Moreover, the endpoint connecting a spike to the body is the *puncture* of this spike, the other endpoint is its *tip*. A spike that enters the body of a kebab is an *in-spike*, one that leaves the body is an *out-spike*. Kebabs are further differentiated wrt. the orientation of their spikes. A kebab with two in-spikes is an *in*-kebab, one with two out-spikes is an *out*-kebab. If one spike is an in-spike and the other an out-spike, we call this kebab an *inout*-kebab. An in-kebab is sketched in Fig. 12a.

In order to prove an essential lemma about kebabs, we need two technical results.

Proposition 10 *Let $(H, src, snk) \in \mathbf{H}$ be spl-normal and let $v \in V_H \setminus \{src, snk\}$ with $vv \notin A_H$. Then v is incident to at least three nonparallel proper arcs.*

Proof Fix a (src, v)-path and a (v, snk)-path in H, and let $a_1 = uv$ and $a_2 = vw$ be the in- and out-arcs of v on these paths. Obviously, a_1 and a_2 are distinct and proper. If v does not carry a loop, it must be incident to a third arc a_3, since H is s-normal. Since H is also p-normal, the a_i are pairwise nonparallel. □

Proposition 11 *Let G be a strong graph with distinct vertices x and y. Then there is a cycle $C \subseteq G$ with distinct vertices $z_x, z_y \in V_C$, s.t. G contains an (x, z_x)-path and a (y, z_y)-path which are disjoint.*

Proof Since G is strong, it contains an (x, y)-path P_1 and a (y, x)-path P_2. The set of crossing vertices Z is defined as the set of internal vertices shared by the P_i, $Z := (V_{P_1} \cap V_{P_2}) \setminus \{x, y\}$. Three cases need to be considered.

- If $Z = \emptyset$, then P_1 and P_2 form a cycle already, and we choose $z_x = x$ and $z_y = y$. In this case, P_x and P_y are both empty.
- If $Z = \{z\}$, the claim follows for $z_x = x$ and $z_y = z$, or $z_x = x$ and $z_y = z$. In this case, either P_x or P_y is empty, but not both.
- Finally if $|Z| \geq 2$, let z_x and z_y be consecutive crossing vertices, i.e., assume that no $z \in Z$ lies between z_x and z_y on P_1, resp. P_2. These vertices satisfy the claim.

□

We are now set to prove that a hammock that is not contained in **SPL** necessarily has a minor in **F** or a bare minor in **K**. The work that is involved in the proof is divided over the following three lemmas, and the pieces are put together in Thm. 7.

Lemma 9 *Let $G \in \mathbf{H}$ be spl-normal. If G contains an in-kebab or an out-kebab, then $\mathsf{C} \preccurlyeq G$, $\mathsf{C}^\mathsf{R} \preccurlyeq G$ or $\mathsf{Q} \preccurlyeq G$ holds.*

Proof We prove the claim for an in-kebab, the argument is symmetric for an out-kebab. Let $G = (G, src, snk)$ and choose K as an arc-maximal kebab of G, with body B and spikes S_1 and S_2 with tips t_i and punctures p_i of S_i (Fig. 12a). By Prop. 1 exactly one of the following holds for the tips of K in G:

a) t_1 dominates t_2,
b) t_2 dominates t_1, or
c) some $x \in V_G$ dominates either t_i, and G contains internally disjoint (x,t_1)- and (x,t_2)-paths.

The cases a) and b) are symmetric and we distinguish cases b) and c).

b) If t_2 dominates t_1, let P be a shortest (t_2,t_1)-path in G. If P is disjoint with B, then P contains an (S_2,S_1)-subpath. With Prop. 11 we find $\mathsf{C} \preccurlyeq G$ (Fig. 12b).
So let P and B intersect, then P contains a "terminal" segment P': the (B,t_1)-subpath that is internally disjoint with B. Our choice of K implies that P' consists of a single arc. First, P' is internally disjoint with S_1: otherwise, we could add arcs from P' to K, to get a kebab with more arcs. Now if P' contains several arcs, we remove its initial arc and use the remainder to extend S_1 to a longer in-spike. Either assumption contradicts our choice of K, thus P' consists of a single bt_1-arc, for some $b \in B$ (Fig. 12c).
We show that t_1 is incident to a third proper arc. Suppose that this claim is false. Since G is s- and p-normal, Prop. 10 implies that t_1 carries a loop. As G is also ℓ-normal, t_1 guards some distinct vertex. Consider the strong subgraph B', consisting of S_1, B, and the bt_1-arc (cf. Fig. 12c). If t_1 guards any vertex of B', then t_1 guards all of B'; consequently, there is a (src,t_1)-path which is otherwise disjoint with B'. On the other hand, if all vertices guarded by t_1 lie outside B', then t_1 reaches such a vertex through a path that is internally disjoint with B'. In either case, t_1 is incident to a third proper arc, contradicting the assumption that it is not.
Let a' denote this arc, i.e., let $a' = t_1z$ or $a' = zt_1$. Our choice of K implies $z \in V_K$. To see this, notice that B' properly contains B. If z lies outside K, we find a kebab K' with body B', in-spike S_2, and a', considered as a path, as a second in- or out-spike. Since $A_{K'}$ properly includes A_K this contradicts our choice of K.
It remains to locate z in K. With help of Prop. 11, we see (from Fig. 12c) that $z \in V_{S_2}$ yields $\mathsf{C} \preccurlyeq G$, and that $z \in V_B$ yields $\mathsf{C} \preccurlyeq G$ or $\mathsf{C}^{\mathsf{R}} \preccurlyeq G$ (depending on a's orientation).
So let $z \in V_{S_1} \setminus \{p_1\}$ and consider the orientation of a'. For $a' = zt_1$ we find $\mathsf{Q} \preccurlyeq G$, with pegs t_1, p_1, b, and z (Fig. 12d). On the other hand, $a' = t_1z$ leads to a contradiction: Since G is p-normal, at least one vertex z' lies between t_1 and z on S_1; omitting the (t_1,z')-segment of S_1 yields an in-kebab with tips z' and t_2 and a body properly containing B (Fig. 12e). Again, this contradicts our choice of K.

c) Let x dominate t_1 and t_2 s.t. G contains internally disjoint (x,t_1)- and (x,t_2)-paths P_1 and P_2, respectively. If neither P_i intersects B, we find $\mathsf{C} \preccurlyeq G$ with help of Prop. 11; this is shown in Fig. 12f, where x_i denotes the "first" vertex on P_i that is also in S_i.
If, say, P_1 intersects B, let b be the last vertex on P_1 that is in B and let x be the first vertex on P_1 that is in $V_{S_i} \setminus \{p_i\}$. If $x \neq t_1$, we find a kebab whose body contains B properly, contrary to our choice of K. For $x = t_1$, the claim was proven in the previous case already (cf. Fig. 12c).

Lemma 10 *Let $G \in \mathbf{H}$ be spl-normal. If G contains an inout-kebab, then $\mathsf{C} \preccurlyeq G$, $\mathsf{C}^{\mathsf{R}} \preccurlyeq G$, $\mathsf{Q} \preccurlyeq G$ or $\Phi \preccurlyeq_{\mathrm{b}} G$ holds.*

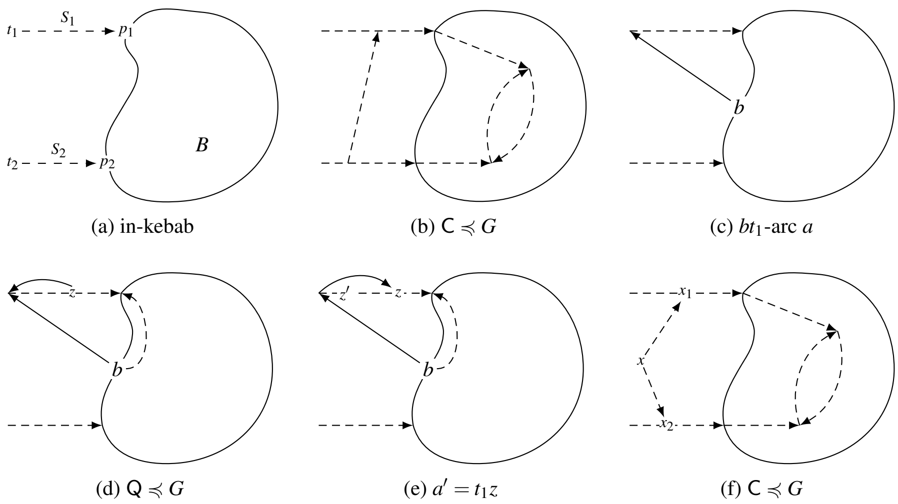

Fig. 12: Cases occurring in the proof of Lem. 9 for an an in-kebab. Solid arrows represent arcs, dashed arrows represent paths.

Proof Let $G = (G, src, snk)$ be as stated and let K be an inout-kebab in G s.t. the body B of K is arc-maximal among all inout-kebabs contained in G. Let S_1 denote the in-spike and S_2 the out-spike of K, and let t_i and p_i denote the tip and puncture of spike S_i.

From Prop. 11 follows $e : \Phi \preccurlyeq_b K$, for some e, where the tips of K are pegs of Φ (Fig. 13a). If e also realizes $\Phi \preccurlyeq_b G$, the claim follows already. Otherwise G contains a bypass P. The structure of Φ requires that at least one tip of K is an endpoint of P. Choosing among t_1 and t_2 leads to symmetric cases; we proceed with t_1. Then P is either a (t_1, q)-path or a (q, t_1)-path, where q is a second peg of Φ wrt. e.

– First let P be a (t_1, q)-path and distinguish by the preimage of q.
  If q is the peg of one of Φ's "inner" vertices, w or x, this leaves only $q = e(x)$ as a possibility. It follows that P is a $(t_1, e(x))$-path that deviates at some point from the path with segments S_1 and the $(p_1, e(x))$-path in B (as drawn in Fig. 13a). Following Prop. 11, we find $\mathsf{C} \preccurlyeq G$ (Fig. 13b).
  For $q = e(y)$ we find subcases that lead to different minors from **F**. We consider the cycle C of the DΦ in G (cf. Fig. 13a), and ask whether P and C are disjoint.
  If so, we find $\mathsf{Q} \preccurlyeq G$ immediately (Fig. 13c), so assume P and C intersect. Since G is a bypass of the found DΦ, at least one of $e(w)$, $e(x)$ does not lie on P. If P does not pass through $e(w)$, then an initial segment of P is a (t_1, C)-path, which yields $\mathsf{C} \preccurlyeq G$, similar to the case shown in Fig. 13b. If P does not pass through $e(x)$, we find $\mathsf{C}^\mathsf{R} \preccurlyeq G$ by taking P's terminal (C, t_2)-segment into account (Fig. 13d).
– Next, let P be a (q, t_1)-path and let k denote the predecessor of t_1 on P. Our choice of K requires $k \in V_K$: otherwise, S_1 could be augmented to a longer spike, a contradiction. We distinguish by the location of k in K.

The case where $k \in V_B$ was already dealt with in the proof of Lem. 9 (replace b with k in Figs. 12c–e). The fact that this time S_2 is an *out*-spike, is irrelevant for the argument. This case yields $\mathsf{Q} \preccurlyeq G$.
Next, suppose $k \in V_{S_2} \setminus \{p_2, t_2\}$, i.e., an inner vertex of S_2. Observe that B, the (p_2, k)-segment of P, the kt_1-arc, and S_1 form a strong graph $B' \subseteq G$ that properly contains B. Clearly, $src \notin V_{B'}$, since $\mathrm{d}^-(q) \geq 1$ for all $q \in V_{B'}$. Hence there is a nonempty (src, B')-path in G. This path and the nonempty (k, t_2)-segment of S_2 constitute a pair of spikes wrt. the body B'. The resulting inout-kebab contradicts our choice of K, so $k \notin V_{S_2} \setminus \{p_2, t_2\}$ follows.
Finally, if $k = t_2$, then P consists of a mere t_2t_1-arc. In this case, K and the t_2t_1-arc form a strong subgraph of G. Let B' denote this subgraph and let P_{src} and P_{snk} denote the shortest (src, B)-path and (B, snk)-path, respectively. Since B' is strong, $src, snk \notin V_{B'}$ follows, thus P_{src} and P_{snk} are nonempty. We also find that P_{src} and P_{snk} are internally disjoint; otherwise we would find a contradiction to K being an inout-kebab with arc-maximal body.
Let b_{src} and P_{snk} denote the endpoint of P_{src} and P_{snk} in B', respectively. If b_{src} and b_{snk} are distinct, G contains an inout-kebab with body B', in-spike P_{src} and out-spike P_{snk}. However, since B' properly contains B, this would contradict our choice of K. Thus $b_{src} = b_{snk}$ follows, and we refer to this vertex as b. We distinguish cases by the location of b in B'.

- If $b \in V_B \setminus \{p_1\}$, then G contains an in-kebab, with body B and in-spikes P_{src} and S_1. Notice that S_1 and P_{src} are indeed disjoint, since P_{src} is a shortest (src, B')-path and S_1 is contained in B'. The claim follows with Lem. 9.
- If $b = p_1$, G contains an out-kebab with body B and out-spikes P_{snk} and S_2; the claim follows with Lem. 9 again.
- If $b \in V_{S_1} \setminus \{p_1\}$, let a denote the third proper arc incident to t_2, as guaranteed by Prop. 10, and let x denote the vertex adjacent to t_2 by a. If x lies outside B', then we find a kebab with body B', and spikes a and P_{src} or P_{snk}, depending on the orientation of a. The found kebab must be an in- or out-kebab—else it would contradict our choice of K—so the claim follows with Lem. 9. This includes the case $x \in V_{P_{src}}$ and (by symmetry) $x \in V_{P_{snk}}$.
  It remains to consider the location of x in B'. First let $x \in V_B \setminus \{p_2\}$. If $a = xt_2$, then $\mathsf{C}^{\mathsf{R}} \preccurlyeq G$ follows with help of Prop. 11 (similar to the case shown in Fig. 13d). If $a = t_2x$, we find an inout-kebab whose body consists of B, S_2 and a, and with the (b, p_1)-segment of S_1 and the t_2t_1-arc as the in- and out-spike, respectively (Fig. 13e). This contradicts our choice of K.
  Next, let $x \in V_{S_1} \setminus \{p_1\}$. For $a = t_2x$, notice that $x \neq t_1$, since G is p-normal. Taking B, S_2 and the (x, p_1)-segment of S_1 as body of an inout-kebab with the t_2t_1-arc as out-spike and P_{src} as in-spike contradicts our choice of K. For $a = xt_2$ we find $\mathsf{Q} \preccurlyeq G$ (Fig. 13f).
  For $x \in V_{S_2}$ the possible orientations of a yield an in-kebab or $\mathsf{Q} \preccurlyeq G$ in a similar manner.
- For $b = p_2$ and $b \in V_{S_2} \setminus \{p_2\}$ the arguments are symmetric to the two preceding cases.

This exhausts all possible cases and the proof is completed. □

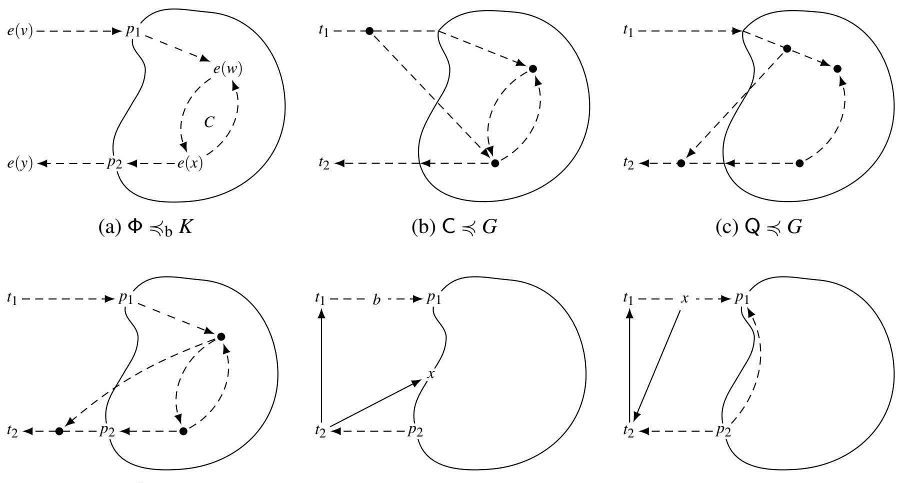

(a) $\Phi \preccurlyeq_{\mathsf{b}} K$ (b) $\mathsf{C} \preccurlyeq G$ (c) $\mathsf{Q} \preccurlyeq G$

(d) $\mathsf{C}^{\mathsf{R}} \preccurlyeq G$ (e) bigger inout-kebab (f) $\mathsf{Q} \preccurlyeq G$

Fig. 13: Cases occurring in the proof of Lem. 10 for an inout-kebab.

Lemma 11 *Let G be an spl-normal hammock with cycles. Then $F \preccurlyeq G$ for some $F \in \mathbf{F}$ or $K \preccurlyeq_{\mathsf{b}} G$ for some $K \in \mathbf{K}$.*

Proof Let (G, src, snk) be spl-normal and non-acyclic. We first show that G contains a proper cycle, i.e., one that is not merely a loop. This is seen as follows: if G is free of loops, it must contain a proper cycle. Otherwise, consider a loop $a = xx \in A_G$. Since G is ℓ-normal, x guards a distinct vertex y. So G contains an (x,y)- and a (y,x)-path; therefore, x and y lie on a proper cycle each (not necessarily the same one, though).

Let C be a smallest proper cycle in G, then $src, snk \notin V_C$ since every vertex of C has in- and out-degree at least one. Thus G contains a shortest (src, C)-path P_{src} and a shortest (C, snk)-path P_{snk}. Let x and y denote the vertices of C where P_{src} enters C and where P_{snk} leaves C, respectively. We distinguish whether x and y coincide.

1. If $x \neq y$, we further distinguish whether P_{src} and P_{snk} are disjoint; the two cases are sketched below.

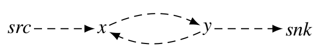

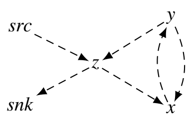

If P_{src} and P_{snk} are disjoint, we have found an inout-kebab with body C, in-spike P_{src}, and out-spike P_{snk}. This can be seen above on the left. In that case, the claim follows by virtue of Lem. 10.

If P_{src} and P_{snk} intersect, let z denote the shared vertex lying "closest" to C, as shown on the right above. Notice that every vertex shared by P_{src} and P_{snk} lies outside C, since P_{src} and P_{snk} are shortest paths and we assume $x \neq y$. We have

found a DΨ with pegs *src*, z, x, and y in G. If this subdivision is bare, the claim follows. Otherwise, let P be a bypass to the DΨ and consider its endpoints.
First, we note that P does not enter *src*, since $\mathrm{d}^-(src) = 0$. If P is a (src, x)-bypass, an inout-kebab occurs in G, and if P is a (src, y)-bypass, we find $\mathsf{C} \preccurlyeq G$. If P is from x to z, or from z to y, $\mathsf{C}^{\mathrm{R}} \preccurlyeq G$ follows. This exhausts all possibilities, and the proof is complete for $x \neq y$.

2. If $x = y$, we refer to this vertex as x_0 and enumerate the vertices of C as $V_C = \{x_0, \ldots, x_n\}$. Since C is a proper cycle, we have $n > 0$. We show that each x_i is adjacent to a vertex outside C.
First, x_i is incident to an arc $a_i \neq A_C$, since G is s-normal. Second, a_i is neither parallel to an arc of C, since G is p-normal, nor is it a chord of C, since C is a smallest cycle of G. This means that $a_i = x_i k_i$, resp. $a_i = k_i x_i$, for $k_i \neq V_C$ or $k_i = x_i$. In the first case, we found a vertex outside C and adjacent to x_i. Otherwise, a_i is a loop, i.e., $k_i = x_i$. Since G is ℓ-normal, this implies that x_i guards a distinct vertex z. We find $z \neq V_C$, as every vertex of C reaches *snk* or is reached from *src* via x_0, without passing through x_i. Consequently, no vertex of C is guarded by x_i, so the first internal vertex on a (x_i, z)-path lies outside C. This concludes the second case.
We have thus found a kebab in G, with body C, one of P_{src} or P_{snk} as one spike, and the arc between x_i and its successor or predecessor outside C as the other spike. The claim now follows with Lem. 9 or Lem. 10.

Each case implies a (bare) minor of G as claimed; since the cases are exhaustive, the claim follows. □

We combine the results found in this section to finally get a characterization of **SPL** by means of forbidden minors.

Theorem 7 *Let G be a hammock. Then*

$$G \in \mathbf{SPL} \quad \text{iff} \quad F \not\preccurlyeq G \text{ for each } F \in \mathbf{F} \text{ and } K \not\preccurlyeq_{\mathrm{b}} G \text{ for each } K \in \mathbf{K}.$$

Proof Let $G \in \mathbf{SPL}$, then G is a hammock. From Lem. 3 follows that G is $\mathbf{F}$-free, while Lem. 4 states that G is free of bare $\mathbf{K}$-minors. Therefore, minor-freeness as claimed is necessary.

To prove sufficiency, assume $G \notin \mathbf{SPL}$ for a hammock G. Then $\mathrm{R}(G) \neq \mathsf{P}_1$ follows with Thm. 2. Recall that $\mathrm{R}(G)$ is a hammock, too. If $\mathrm{R}(G)$ is acyclic, Valdes' theorem (Thm. 5) states $\mathsf{N} \preccurlyeq \mathrm{R}(G)$ which implies the claim. Otherwise $\mathrm{R}(G)$ contains cycles, so Lem. 11 yields $F \preccurlyeq \mathrm{R}(G)$ for some $F \in \mathbf{F}$ or $K \preccurlyeq_{\mathrm{b}} \mathrm{R}(G)$ for some $K \in \mathbf{K}$. If G is spl-normal, i.e., $G = \mathrm{R}(G)$, the claim follows immediately. Otherwise, induction on the length of the reduction from G to $\mathrm{R}(G)$, using Thm. 6 in each step, provides the claim. □

Lastly, we observe that the given characterization of **SPL** by forbidden minors is minimal wrt. the minor-notions we employ. In other words, no member of **F** or **K** can be omitted. This is shown by giving hammocks that have exactly one (bare) minor in **F** or **K**.

Lemma 12

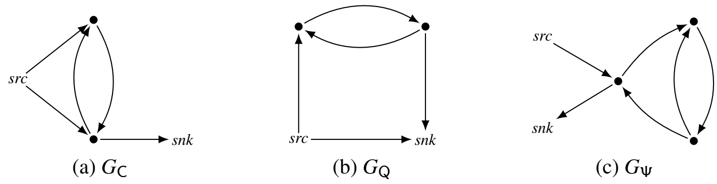

Fig. 14: Hammocks with exclusive (bare) minor C, Q, and Ψ.

1. *For each* $F \in \mathbf{F}$ *there is some* $G_F \in \mathbf{H}$ *s.t.* $F \preccurlyeq G_F$*, while* $F' \not\preccurlyeq G_F$ *for* $F' \in \mathbf{F} \setminus \{F\}$ *and* $K \not\preccurlyeq_{\mathsf{b}} G_F$ *for* $K \in \mathbf{K}$.
2. *For each* $K \in \mathbf{K}$ *there is some* $G_K \in \mathbf{H}$ *s.t.* $K \preccurlyeq_{\mathsf{b}} G_K$*, while* $F \not\preccurlyeq G_K$ *for* $F \in \mathbf{F}$ *and* $K' \not\preccurlyeq_{\mathsf{b}} G_K$ *for* $K' \in \mathbf{K} \setminus \{K\}$.

Proof We have $G_{\mathsf{N}} = \mathsf{N}$ and $G_{\Phi} = \Phi$. The graphs G_{C}, G_{Q} and G_{Ψ} are shown in Fig. 14. Symmetry further yields $G_{\mathsf{C}^{\mathsf{R}}} = G_{\mathsf{C}}^{\mathsf{R}}$ and $G_{\Psi^{\mathsf{R}}} = G_{\Psi}^{\mathsf{R}}$.

6 Conclusions

The class of series parallel loop graphs, **SPL**, extends that of arc-series parallel graphs beyond the acyclic case. These graphs have been characterized in three ways: by a recursive constriction, by an encoding through expressions and—within the class of hammocks—by a set of seven forbidden substructures.

The forbidden subgraph characterization contributes to the structural theory of directed graphs. Few such results are known for digraph classes, as opposed to the undirected case, where these characterizations are met rather often [2]. It is worth mentioning that spl-graphs can further be defined as graphs of treewidth two, if this notion is extended to digraphs with multiple arcs and loops. Among simple undirected graphs, treewidth two captures the class of series parallel graphs—the favourable algorithmic properties of this class [1,17]) might generalize to spl-graphs seamlessly.

The definitory and the expression characterization are tightly coupled, since the recursive structure of an spl-graph is immediately apparent in the parse of its associated expression (modulo associativity and commutativity). With a litte extra effort, this correspondence can be exploited in the conversion among REs and εFAs. Let us refer to εFAs with spl-structure as *spl-FAs*.

We have shown that any spl-FA can be converted to an expression that is linear in the size of the FA. Consequently, a trim normalized FA that requires expressions of superlinear size cannot have spl-structure. The forbidden minor characterization of **SPL** further dictates that such an FA contains structures from **F** or **K**. How to get a converse kind of result is unclear as yet. Any fixed FA with features from **F** and **K** can be converted to a linear-sized expression, given the right factor. The same goes for growing FAs with a fixed number of isolated occurrences of such forbidden structures. This suggests to devise a quantitative notion of “overlapping” or “nested” minors (or subdivisions). Notice that the graphs used by Ehrenfeucht & Zeiger to show exponential blowup trivially contain any smaller graph as a minor.

Let us finally compare spl-FAs with FAs that result from the constructions by Thompson and Glushkov.

Thompson-FAs and spl-FAs are both linear in the size of their origin expressions, owing to the presence of ε-transitions. It is easy to see that a Thompson-FA can also be "deconstructed" in a straightforward manner to the expression it originates from. However, spl-FAs contains considerably fewer ε-transitions: a Thompson-FA contains one ε-transition per product, four ε-transitions per sum, and three per Kleene-star in the original expression, while an spl-FA contains at most two for a star and none else. This yields a smaller FA and allows for a faster computation of the ε-closure in possible follow-up constructions.

Comparing Glushkov-FAs and spl-FAs by size is futile, as the former FAs are ε-free. This requires an $\Omega(n\ \log n)$ increase over expression size in the general case [15]. In the converse direction, their structural characterization, suggests that Glushkov-FAs can be converted to linear sized expressions [3]. An encoding of unlabeled digraphs with the appropriate structure should be applicable, by labeling each arc with its head vertex. However, it is unclear whether these transformations are invertible.

It might be further be insightful to devise forbidden minor characterizations of the graphs underlying Thompson-FAs and Glushkov-FAs, as these minors provide absent structures "at a glance".

References

1. Bodlaender, H.L., deFluiter, B.: Parallel algorithms for series parallel graphs and graphs with treewidth two. Algorithmica **29**(4), 534–559 (2001)
2. Brandstädt, A., Le, V.B., Spinrad, J.P.: Graph Classes: A Survey. No. 3 in Monographs on Discrete Mathematics and Applications (1999)
3. Caron, P., Ziadi, D.: Characterization of Glushkov automata. Theoretical Computer Science **233**(1-2), 75–90 (2000)
4. Chrobak, M.: Finite automata and unary languages. Theoretical Computer Science **47**, 149–158 (1986)
5. Diestel, R.: Graph Theory, 3 edn. No. 173 in Graduate Texts in Mathematics. Springer (2006)
6. Ehrenfeucht, A., Zeiger, P.: Complexity measures for regular expressions. Journal of Computer and System Sciences **12**, 134–146 (1976)
7. Ellul, K., Krawetz, B., Shallit, J., Wang, M.W.: Regular expressions: new results and open problems. Journal of Automata, Languages and Combinatorics **10**(4), 407–437 (2005)
8. Gelade, W., Neven, F.: Succinctness of the complement and intersection of regular expressions. In: STACS, pp. 325–336 (2008)
9. Giammarresi, D., Ponty, J.L., Wood, D., Ziadi, D.: A characterization of Thompson digraphs. Discrete Applied Mathematics **134**, 317–337 (2004)
10. Gruber, H., Holzer, M.: Finite automata, digraph connectivity, and regular expression size. In: ICALP 2008, Part II, no. 5126 in LNCS, pp. 39–50. Springer (2008)
11. Gruber, H., Johannsen, J.: Optimal lower bounds on regular expression size using communication complexity. In: FOSSAC 2008, no. 4962 in LNCS, pp. 273–286 (2008)
12. Gulan, S., Fernau, H.: Local elimination-strategies in finite automata for shorter regular expressions. In: Theory and Practice of Computer Science: SOFSEM 2008, vol. 2, pp. 46–57 (2008)
13. Gulan, S., Fernau, H.: An optimal construction of finite automata from regular expressions. In: 28th FSTTCS, pp. 211–222 (2008)
14. Hopcroft, J.E., Ullman, J.D.: Introduction to Automata Theory, Languages and Computation. Addison-Wesley (1979)

15. Hromkovič, J., Seibert, S.: Translating regular expressions into small epsilon-free nondeterministic finite automata. Journal of Computer and System Sciences **62**, 565–588 (2001)
16. Huet, G.: Confluent reductions: Abstract properties and applications to term rewriting systems. Journal of the ACM **27**(4), 797–821 (1980)
17. Jakoby, A., Tantau, T.: Logspace algorithms for computing shortest and longest paths in series-parallel graphs. In: 27th FSTTCS, pp. 216–227 (2007)
18. Kleene, S.C.: Representation of Events in Nerve Nets and Finite Automata, pp. 3–41. Annals of Mathematics Studies (1956)
19. Korenblit, M., Levit, V.E.: On algebraic expressions of series-parallel and Fibonacci graphs. In: DMTCS 2003, no. 2731 in LNCS, pp. 215–224. Springer (2003)
20. McNaughton, R.: Techniques for manipulating regular expressions. In: Systems and Computer Science, pp. 27–41. University of Toronto Press in association with the University of Western Ontario Toronto (1967)
21. Moreira, N., Reis, R.: Series-parallel automata and short regular expressions. Fundamenta Informaticae **91**(3-4), 611–629 (2009)
22. Newman, M.: On theories with a combinatorial definition of "equivalence". Annals of Mathematics **43**(2) (1942)
23. Valdes, J.: Parsing flowcharts and series-parallel graphs. Ph.D. thesis, Stanford University (1978). STAN-CS-78-682
24. Valdes, J., Tarjan, R.E., Lawler, E.L.: The recognition of series parallel digraphs. SIAM Journal on Computing **11**(2), 298–313 (1981)
25. Watson, B.W.: A taxonomy of finite automata construction algorithms. Tech. Rep. Computing Science Note 93/43, Eindhoven University of Technology (1994)