POSIX Lexing with Derivatives of Regular Expressions

Fahad Ausaf[1], Roy Dyckhoff[2], and Christian Urban[3]

[1] King's College London
`fahad.ausaf@icloud.com`
[2] University of St Andrews
`roy.dyckhoff@st-andrews.ac.uk`
[3] King's College London
`christian.urban@kcl.ac.uk`

Abstract. Brzozowski introduced the notion of derivatives for regular expressions. They can be used for a very simple regular expression matching algorithm. Sulzmann and Lu cleverly extended this algorithm in order to deal with POSIX matching, which is the underlying disambiguation strategy for regular expressions needed in lexers. Their algorithm generates POSIX values which encode the information of *how* a regular expression matches a string—that is, which part of the string is matched by which part of the regular expression. In this paper we give our inductive definition of what a POSIX value is and show (*i*) that such a value is unique (for given regular expression and string being matched) and (*ii*) that Sulzmann and Lu's algorithm always generates such a value (provided that the regular expression matches the string). We show that (*iii*) our inductive definition of a POSIX value is equivalent to an alternative definition by Okui and Suzuki which identifies POSIX values as least elements according to an ordering of values. We also prove the correctness of Sulzmann's bitcoded version of the POSIX matching algorithm and extend the results to additional constructors for regular expressions.

Keywords: POSIX matching, Derivatives of Regular Expressions, Isabelle/HOL

1 Core of the proof

This paper builds on previous work by Ausaf and Urban using regular expression'd bit-coded derivatives to do lexing that is both fast and satisfies the POSIX

---

[*] This paper is a revised and expanded version of [2]. Compared with that paper we give a second definition for POSIX values introduced by Okui Suzuki [10,11] and prove that it is equivalent to our original one. This second definition is based on an ordering of values and very similar to, but not equivalent with, the definition given by Sulzmann and Lu [13]. The advantage of the definition based on the ordering is that it implements more directly the informal rules from the POSIX standard. We also prove Sulzmann & Lu's conjecture that their bitcoded version of the POSIX algorithm is correct. Furthermore we extend our results to additional constructors of regular expressions.

specification. In their work, a bit-coded algorithm introduced by Sulzmann and Lu was formally verified in Isabelle, by a very clever use of flex function and retrieve to carefully mimic the way a value is built up by the injection funciton.

In the previous work, Ausaf and Urban established the below equality:

Lemma 1. *If $v : r \backslash s$ then Some $(flex\ r\ id\ s\ v) = decode\ (retrieve\ (r^{\uparrow} \backslash\!\backslash s)\ v)\ r$.*

This lemma establishes a link with the lexer without bit-codes.
With it we get the correctness of bit-coded algorithm.

Lemma 2. *$lexer_b\ r\ s = lexer\ r\ s$*

However what is not certain is whether we can add simplification to the bit-coded algorithm, without breaking the correct lexing output.

The reason that we do need to add a simplification phase after each derivative step of *blexer* is because it produces intermediate regular expressions that can grow exponentially. For example, the regular expression $(a + aa)^*$ after taking derivative against just 10 as will have size 8192.

Therefore, we insert a simplification phase after each derivation step, as defined below:

Lemma 3. *$blexer\_\_simp\ r\ s \overset{def}{=}$ if $nullable_b\ (bders\_\_simp\ (r^{\uparrow})\ s)$ then decode $(mkeps_b\ (bders\_\_simp\ (r^{\uparrow})\ s))\ r$ else None*

The simplification function is given as follows:

$$bsimp\ (ASEQ\ bs\ r_1\ r_2)\ \overset{def}{=}\ bsimp\_\_ASEQ\ bs\ (bsimp\ r_1)\ (bsimp\ r_2)$$
$$bsimp\ (AALTs\ bs1.0\ rs)\ \overset{def}{=}\ bsimp\_\_AALTs\ bs1.0\ (distinctBy\ (flts\ (map\ bsimp\ rs))\ erase\ \varnothing)$$
$$bsimp\ AZERO\ \qquad\quad \overset{def}{=}\ AZERO$$

And the two helper functions are:

$$bsimp\_\_AALTs\ bs_1\ [r]\ \qquad\qquad \overset{def}{=}\ bsimp\_\_ASEQ\ bs_1\ (bsimp\ r)\ (bsimp\ r2.0)$$
$$bsimp\_\_AALTs\ bs1.0\ [r]\ \qquad\qquad \overset{def}{=}\ bsimp\_\_AALTs\ bs1.0\ (distinctBy\ (flts\ (map\ bsimp\ rs))\ erase\ \varnothing)$$
$$bsimp\_\_AALTs\ bs1.0\ (v\!::\!vb\!::\!vc)\ \overset{def}{=}\ AZERO$$

This might sound trivial in the case of producing a YES/NO answer, but once we require a lexing output to be produced (which is required in applications like compiler front-end, malicious attack domain extraction, etc.), it is not straightforward if we still extract what is needed according to the POSIX standard.

By simplification, we mean specifically the following rules:

$$\overline{ASEQ\ bs\ AZERO\ r_2 \rightsquigarrow AZERO}$$

$$\overline{ASEQ\ bs\ r_1\ AZERO \rightsquigarrow AZERO}$$

$$\overline{ASEQ\ bs\ (AONE\ bs_1)\ r_1 \rightsquigarrow fuse\ (bs\ @\ bs_1)\ r_1}$$

$$\frac{r_1 \rightsquigarrow r_2}{ASEQ\ bs\ r_1\ r_3 \rightsquigarrow ASEQ\ bs\ r_2\ r_3}$$

$$\frac{r_3 \rightsquigarrow r_4}{ASEQ\ bs\ r_1\ r_3 \rightsquigarrow ASEQ\ bs\ r_1\ r_4}$$

$$\frac{r \rightsquigarrow r'}{AALTs\ bs\ (rs_1\ @\ [r]\ @\ rs_2) \rightsquigarrow AALTs\ bs\ (rs_1\ @\ [r']\ @\ rs_2)}$$

$$\overline{AALTs\ bs\ (rs_a\ @\ [AZERO]\ @\ rs_b) \rightsquigarrow AALTs\ bs\ (rs_a\ @\ rs_b)}$$

$$\overline{AALTs\ bs\ (rs_a\ @\ [AALTs\ bs_1\ rs_1]\ @\ rs_b) \rightsquigarrow AALTs\ bs\ (rs_a\ @\ map\ (fuse\ bs_1)\ rs_1\ @\ rs_b)}$$

And these can be made compact by the following simplification function:
where the function $bsimp_AALTs$

The core idea of the proof is that two regular expressions, if "isomorphic" up to a finite number of rewrite steps, will remain "isomorphic" when we take the same sequence of derivatives on both of them. This can be expressed by the following rewrite relation lemma:

Lemma 4. $(r\backslash\!\backslash s) \rightsquigarrow* bders\_simp\ r\ s$

This isomorphic relation implies a property that leads to the correctness result: if two (nullable) regular expressions are "rewritable" in many steps from one another, then a call to function $bmkeps$ gives the same bit-sequence :

Lemma 5. If $r1.0 \rightsquigarrow* r2.0$ and $nullable_b\ r1.0$ then $mkeps_b\ r1.0 = mkeps_b\ r2.0$.

Given the same bit-sequence, the decode function will give out the same value, which is the output of both lexers:

Lemma 6. $lexer_b\ r\ s \stackrel{def}{=}$ if $nullable_b\ (r^{\uparrow}\backslash\!\backslash s)$ then $decode\ (mkeps_b\ (r^{\uparrow}\backslash\!\backslash s))\ r$ else None

Lemma 7. $blexer\_simp\ r\ s \stackrel{def}{=}$ if $nullable_b\ (bders\_simp\ (r^{\uparrow})\ s)$ then $decode\ (mkeps_b\ (bders\_simp\ (r^{\uparrow})\ s))\ r$ else None

And that yields the correctness result:

Lemma 8. $lexer\ r\ s = blexer\_simp\ r\ s$

The nice thing about the above

2   Additional Simp Rules?

One question someone would ask is: can we add more "atomic" simplification/rewriting rules, so the simplification is even more aggressive, making the intermediate results smaller, and therefore more space-efficient? For example, one might want to do open up alternatives who is a child of a sequence:

ASEQ bs (AALTs bs1 rs) r ⤳? AALTs (bs @ bs1) (map (λr′. ASEQ [] r′ r) rs)

This rule allows us to simplify $(a + b) \cdot c + a \cdot c$ into $a \cdot c + b \cdot c$, which cannot be done under the rrewrite rule because only alternatives which are children of another alternative can be spilled out.

Now with this rule we have some examples where we get even smaller intermediate derivatives than not having this rule: *(a+aa)* → *(1+1a)(a+aa)* → *(1+a)(a+aa)**

3   Introduction

Brzozowski [3] introduced the notion of the *derivative* $r \backslash c$ of a regular expression r w.r.t. a character c, and showed that it gave a simple solution to the problem of matching a string s with a regular expression r: if the derivative of r w.r.t. (in succession) all the characters of the string matches the empty string, then r matches s (and *vice versa*). The derivative has the property (which may almost be regarded as its specification) that, for every string s and regular expression r and character c, one has $cs \in L(r)$ if and only if $s \in L(r \backslash c)$. The beauty of Brzozowski's derivatives is that they are neatly expressible in any functional language, and easily definable and reasoned about in theorem provers—the definitions just consist of inductive datatypes and simple recursive functions. A mechanised correctness proof of Brzozowski's matcher in for example HOL4 has been mentioned by Owens and Slind [12]. Another one in Isabelle/HOL is part of the work by Krauss and Nipkow [8]. And another one in Coq is given by Coquand and Siles [4].

If a regular expression matches a string, then in general there is more than one way of how the string is matched. There are two commonly used disambiguation strategies to generate a unique answer: one is called GREEDY matching [5] and the other is POSIX matching [1,9,10,13,14]. For example consider the string *xy* and the regular expression $(x + y + xy)^{\star}$. Either the string can be matched in two 'iterations' by the single letter-regular expressions x and y, or directly in one iteration by *xy*. The first case corresponds to GREEDY matching, which first matches with the left-most symbol and only matches the next symbol in case of a mismatch (this is greedy in the sense of preferring instant gratification to delayed repletion). The second case is POSIX matching, which prefers the longest match.

In the context of lexing, where an input string needs to be split up into a sequence of tokens, POSIX is the more natural disambiguation strategy for what

programmers consider basic syntactic building blocks in their programs. These building blocks are often specified by some regular expressions, say r_{key} and r_{id} for recognising keywords and identifiers, respectively. There are a few underlying (informal) rules behind tokenising a string in a POSIX [1] fashion:

- *The Longest Match Rule* (or *"Maximal Munch Rule"*): The longest initial substring matched by any regular expression is taken as next token.
- *Priority Rule:* For a particular longest initial substring, the first (leftmost) regular expression that can match determines the token.
- *Star Rule:* A subexpression repeated by * shall not match an empty string unless this is the only match for the repetition.
- *Empty String Rule:* An empty string shall be considered to be longer than no match at all.

Consider for example a regular expression r_{key} for recognising keywords such as *if*, *then* and so on; and r_{id} recognising identifiers (say, a single character followed by characters or numbers). Then we can form the regular expression $(r_{key} + r_{id})^\star$ and use POSIX matching to tokenise strings, say *iffoo* and *if*. For *iffoo* we obtain by the Longest Match Rule a single identifier token, not a keyword followed by an identifier. For *if* we obtain by the Priority Rule a keyword token, not an identifier token—even if r_{id} matches also. By the Star Rule we know $(r_{key} + r_{id})^\star$ matches *iffoo*, respectively *if*, in exactly one 'iteration' of the star. The Empty String Rule is for cases where, for example, the regular expression $(a^\star)^\star$ matches against the string *bc*. Then the longest initial matched substring is the empty string, which is matched by both the whole regular expression and the parenthesised subexpression.

One limitation of Brzozowski's matcher is that it only generates a YES/NO answer for whether a string is being matched by a regular expression. Sulzmann and Lu [13] extended this matcher to allow generation not just of a YES/NO answer but of an actual matching, called a [lexical] *value*. Assuming a regular expression matches a string, values encode the information of *how* the string is matched by the regular expression—that is, which part of the string is matched by which part of the regular expression. For this consider again the string *xy* and the regular expression $(x + (y + xy))^\star$ (this time fully parenthesised). We can view this regular expression as tree and if the string *xy* is matched by two Star 'iterations', then the x is matched by the left-most alternative in this tree and the y by the right-left alternative. This suggests to record this matching as

$$Stars~[Left~(Char~x),~Right~(Left~(Char~y))]$$

where *Stars*, *Left*, *Right* and *Char* are constructors for values. *Stars* records how many iterations were used; *Left*, respectively *Right*, which alternative is used. This 'tree view' leads naturally to the idea that regular expressions act as types and values as inhabiting those types (see, for example, [7]). The value for matching *xy* in a single 'iteration', i.e. the POSIX value, would look as follows

$$Stars\ [Seq\ (Char\ x)\ (Char\ y)]$$

where *Stars* has only a single-element list for the single iteration and *Seq* indicates that xy is matched by a sequence regular expression.

Sulzmann and Lu give a simple algorithm to calculate a value that appears to be the value associated with POSIX matching. The challenge then is to specify that value, in an algorithm-independent fashion, and to show that Sulzmann and Lu's derivative-based algorithm does indeed calculate a value that is correct according to the specification. The answer given by Sulzmann and Lu [13] is to define a relation (called an "order relation") on the set of values of r, and to show that (once a string to be matched is chosen) there is a maximum element and that it is computed by their derivative-based algorithm. This proof idea is inspired by work of Frisch and Cardelli [5] on a GREEDY regular expression matching algorithm. However, we were not able to establish transitivity and totality for the "order relation" by Sulzmann and Lu. There are some inherent problems with their approach (of which some of the proofs are not published in [13]); perhaps more importantly, we give in this paper a simple inductive (and algorithm-independent) definition of what we call being a *POSIX value* for a regular expression r and a string s; we show that the algorithm by Sulzmann and Lu computes such a value and that such a value is unique. Our proofs are both done by hand and checked in Isabelle/HOL. The experience of doing our proofs has been that this mechanical checking was absolutely essential: this subject area has hidden snares. This was also noted by Kuklewicz [9] who found that nearly all POSIX matching implementations are "buggy" [13, Page 203] and by Grathwohl et al [6, Page 36] who wrote:

> *"The POSIX strategy is more complicated than the greedy because of the dependence on information about the length of matched strings in the various subexpressions."*

Contributions: We have implemented in Isabelle/HOL the derivative-based regular expression matching algorithm of Sulzmann and Lu [13]. We have proved the correctness of this algorithm according to our specification of what a POSIX value is (inspired by work of Vansummeren [14]). Sulzmann and Lu sketch in [13] an informal correctness proof: but to us it contains unfillable gaps.[4] Our specification of a POSIX value consists of a simple inductive definition that given a string and a regular expression uniquely determines this value. We also show that our definition is equivalent to an ordering of values based on positions by Okui and Suzuki [10].

We extend our results to ??? Bitcoded version??

---

[4] An extended version of [13] is available at the website of its first author; this extended version already includes remarks in the appendix that their informal proof contains gaps, and possible fixes are not fully worked out.

4 Preliminaries

Strings in Isabelle/HOL are lists of characters with the empty string being represented by the empty list, written [], and list-cons being written as $\_ :: \_$. Often we use the usual bracket notation for lists also for strings; for example a string consisting of just a single character c is written $[c]$. We use the usual definitions for *prefixes* and *strict prefixes* of strings. By using the type *char* for characters we have a supply of finitely many characters roughly corresponding to the ASCII character set. Regular expressions are defined as usual as the elements of the following inductive datatype:

$$r := \mathbf{0} \mid \mathbf{1} \mid c \mid r_1 + r_2 \mid r_1 \cdot r_2 \mid r^\star$$

where $\mathbf{0}$ stands for the regular expression that does not match any string, $\mathbf{1}$ for the regular expression that matches only the empty string and c for matching a character literal. The language of a regular expression is also defined as usual by the recursive function L with the six clauses:

$$
\begin{array}{rll}
\text{(1)} & L(\mathbf{0}) & \stackrel{\text{def}}{=} \varnothing \\
\text{(2)} & L(\mathbf{1}) & \stackrel{\text{def}}{=} \{[]\} \\
\text{(3)} & L(c) & \stackrel{\text{def}}{=} \{[c]\} \\
\text{(4)} & L(r_1 \cdot r_2) & \stackrel{\text{def}}{=} L(r_1) \,@\, L(r_2) \\
\text{(5)} & L(r_1 + r_2) & \stackrel{\text{def}}{=} L(r_1) \cup L(r_2) \\
\text{(6)} & L(r^\star) & \stackrel{\text{def}}{=} (L(r))\star
\end{array}
$$

In clause *(4)* we use the operation $\_ \,@\, \_$ for the concatenation of two languages (it is also list-append for strings). We use the star-notation for regular expressions and for languages (in the last clause above). The star for languages is defined inductively by two clauses: (i) the empty string being in the star of a language and (ii) if s_1 is in a language and s_2 in the star of this language, then also $s_1 \,@\, s_2$ is in the star of this language. It will also be convenient to use the following notion of a *semantic derivative* (or *left quotient*) of a language defined as

$$Der\ c\ A \stackrel{def}{=} \{s \mid c :: s \in A\}\,.$$

For semantic derivatives we have the following equations (for example mechanically proved in [8]):

$$
\begin{array}{lll}
Der\ c\ \varnothing & \stackrel{\text{def}}{=} \varnothing \\
Der\ c\ \{[]\} & \stackrel{\text{def}}{=} \varnothing \\
Der\ c\ \{[d]\} & \stackrel{\text{def}}{=} if\ c = d\ then\ \{[]\}\ else\ \varnothing \\
Der\ c\ (A \cup B) & \stackrel{\text{def}}{=} Der\ c\ A \cup Der\ c\ B & \qquad (1) \\
Der\ c\ (A \,@\, B) & \stackrel{\text{def}}{=} (Der\ c\ A \,@\, B) \cup (if\ [] \in A\ then\ Der\ c\ B\ else\ \varnothing) \\
Der\ c\ (A\star) & \stackrel{\text{def}}{=} Der\ c\ A \,@\, A\star
\end{array}
$$

Brzozowski's derivatives of regular expressions [3] can be easily defined by two recursive functions: the first is from regular expressions to booleans (implementing a test when a regular expression can match the empty string), and the second takes a regular expression and a character to a (derivative) regular expression:

$$
\begin{aligned}
nullable\ (\mathbf{0}) &\overset{\text{def}}{=} False \\
nullable\ (\mathbf{1}) &\overset{\text{def}}{=} True \\
nullable\ (c) &\overset{\text{def}}{=} False \\
nullable\ (r_1 + r_2) &\overset{\text{def}}{=} nullable\ r_1 \lor nullable\ r_2 \\
nullable\ (r_1 \cdot r_2) &\overset{\text{def}}{=} nullable\ r_1 \land nullable\ r_2 \\
nullable\ (r^\star) &\overset{\text{def}}{=} True
\end{aligned}
$$

$$
\begin{aligned}
\mathbf{0}\backslash c &\overset{\text{def}}{=} \mathbf{0} \\
\mathbf{1}\backslash c &\overset{\text{def}}{=} \mathbf{0} \\
d\backslash c &\overset{\text{def}}{=} if\ c = d\ then\ \mathbf{1}\ else\ \mathbf{0} \\
(r_1 + r_2)\backslash c &\overset{\text{def}}{=} (r_1\backslash c) + (r_2\backslash c) \\
(r_1 \cdot r_2)\backslash c &\overset{\text{def}}{=} if\ nullable\ r_1\ then\ (r_1\backslash c) \cdot r_2 + (r_2\backslash c)\ else\ (r_1\backslash c) \cdot r_2 \\
(r^\star)\backslash c &\overset{\text{def}}{=} (r\backslash c) \cdot r^\star
\end{aligned}
$$

We may extend this definition to give derivatives w.r.t. strings:

$$
\begin{aligned}
r\backslash[] &\overset{\text{def}}{=} r \\
r\backslash(c :: s) &\overset{\text{def}}{=} (r\backslash c)\backslash s
\end{aligned}
$$

Given the equations in (1), it is a relatively easy exercise in mechanical reasoning to establish that

Proposition 1.
(1) nullable r if and only if $[] \in L(r)$, and
(2) $L(r\backslash c) = Der\ c\ (L(r))$.

With this in place it is also very routine to prove that the regular expression matcher defined as

$$
match\ r\ s \overset{def}{=} nullable\ (r\backslash s)
$$

gives a positive answer if and only if $s \in L(r)$. Consequently, this regular expression matching algorithm satisfies the usual specification for regular expression matching. While the matcher above calculates a provably correct YES/NO answer for whether a regular expression matches a string or not, the novel idea of Sulzmann and Lu [13] is to append another phase to this algorithm in order to calculate a [lexical] value. We will explain the details next.

5   POSIX Regular Expression Matching

There have been many previous works that use values for encoding *how* a regular expression matches a string. The clever idea by Sulzmann and Lu [13] is to define a function on values that mirrors (but inverts) the construction of the derivative on regular expressions. *Values* are defined as the inductive datatype

$$v := \textit{Empty} \mid \textit{Char } c \mid \textit{Left } v \mid \textit{Right } v \mid \textit{Seq } v_1 \; v_2 \mid \textit{Stars } vs$$

where we use *vs* to stand for a list of values. (This is similar to the approach taken by Frisch and Cardelli for GREEDY matching [5], and Sulzmann and Lu for POSIX matching [13]). The string underlying a value can be calculated by the *flat* function, written $|\_|$ and defined as:

$$\begin{array}{llll}
|\textit{Empty}| & \stackrel{\text{def}}{=} [] & |\textit{Seq } v_1 \; v_2| & \stackrel{\text{def}}{=} |v_1| \; @ \; |v_2| \\
|\textit{Char } c| & \stackrel{\text{def}}{=} [c] & |\textit{Stars } []| & \stackrel{\text{def}}{=} [] \\
|\textit{Left } v| & \stackrel{\text{def}}{=} |v| & |\textit{Stars } (v :: vs)| & \stackrel{\text{def}}{=} |v| \; @ \; |\textit{Stars } vs| \\
|\textit{Right } v| & \stackrel{\text{def}}{=} |v|
\end{array}$$

We will sometimes refer to the underlying string of a value as *flattened value*. We will also overload our notation and use $|vs|$ for flattening a list of values and concatenating the resulting strings.

Sulzmann and Lu define inductively an *inhabitation relation* that associates values to regular expressions. We define this relation as follows:[5]

$$\frac{}{\textit{Empty} : \mathbf{1}} \qquad\qquad \frac{}{\textit{Char } c : c}$$

$$\frac{v_1 : r_1}{\textit{Left } v_1 : r_1 + r_2} \qquad\qquad \frac{v_2 : r_1}{\textit{Right } v_2 : r_2 + r_1}$$

$$\frac{v_1 : r_1 \qquad v_2 : r_2}{\textit{Seq } v_1 \; v_2 : r_1 \cdot r_2} \qquad\qquad \frac{\forall\, v \in vs. \; v : r \,\wedge\, |v| \neq []}{\textit{Stars } vs : r^\star}$$

where in the clause for *Stars* we use the notation $v \in vs$ for indicating that v is a member in the list *vs*. We require in this rule that every value in *vs* flattens to a non-empty string. The idea is that *Stars*-values satisfy the informal Star Rule (see Introduction) where the * does not match the empty string unless this is the only match for the repetition. Note also that no values are associated with the regular expression $\mathbf{0}$, and that the only value associated with the regular expression $\mathbf{1}$ is *Empty*. It is routine to establish how values "inhabiting" a regular expression correspond to the language of a regular expression, namely

---

[5] Note that the rule for *Stars* differs from our earlier paper [2]. There we used the original definition by Sulzmann and Lu which does not require that the values $v \in vs$ flatten to a non-empty string. The reason for introducing the more restricted version of lexical values is convenience later on when reasoning about an ordering relation for values.

Proposition 2. $L(r) = \{|v| \mid v : r\}$

Given a regular expression r and a string s, we define the set of all *Lexical Values* inhabited by r with the underlying string being s:[6]

$$LV \; r \; s \; \stackrel{def}{=} \; \{v \mid v : r \wedge |v| = s\}$$

The main property of *LV r s* is that it is alway finite.

Proposition 3. *finite* $(LV \; r \; s)$

This finiteness property does not hold in general if we remove the side-condition about $|v| \neq []$ in the *Stars*-rule above. For example using Sulzmann and Lu's less restrictive definition, $LV \; (\mathbf{1}^\star) \; []$ would contain infinitely many values, but according to our more restricted definition only a single value, namely $LV \; (\mathbf{1}^\star) \; [] = \{Stars \; []\}$.

   If a regular expression r matches a string s, then generally the set $LV \; r \; s$ is not just a singleton set. In case of POSIX matching the problem is to calculate the unique lexical value that satisfies the (informal) POSIX rules from the Introduction. Graphically the POSIX value calculation algorithm by Sulzmann and Lu can be illustrated by the picture in Figure 1 where the path from the left to the right involving *derivatives/nullable* is the first phase of the algorithm (calculating successive Brzozowski's derivatives) and *mkeps/inj*, the path from right to left, the second phase. This picture shows the steps required when a regular expression, say r_1, matches the string $[a, b, c]$. We first build the three derivatives (according to a, b and c). We then use *nullable* to find out whether the resulting derivative regular expression r_4 can match the empty string. If yes, we call the function *mkeps* that produces a value v_4 for how r_4 can match the empty string (taking into account the POSIX constraints in case there are several ways). This function is defined by the clauses:

$$
\begin{aligned}
mkeps \; \mathbf{1} \quad &\stackrel{def}{=} \; Empty \\
mkeps \; (r_1 \cdot r_2) \quad &\stackrel{def}{=} \; Seq \; (mkeps \; r_1) \; (mkeps \; r_2) \\
mkeps \; (r_1 + r_2) \quad &\stackrel{def}{=} \; if \; nullable \; r_1 \; then \; Left \; (mkeps \; r_1) \; else \; Right \; (mkeps \; r_2) \\
mkeps \; (r^\star) \quad &\stackrel{def}{=} \; Stars \; []
\end{aligned}
$$

Note that this function needs only to be partially defined, namely only for regular expressions that are nullable. In case *nullable* fails, the string $[a, b, c]$ cannot be matched by r_1 and the null value *None* is returned. Note also how this function makes some subtle choices leading to a POSIX value: for example if an alternative regular expression, say $r_1 + r_2$, can match the empty string and furthermore r_1 can match the empty string, then we return a *Left*-value. The *Right*-value will only be returned if r_1 cannot match the empty string.

---

[6] Okui and Suzuki refer to our lexical values as *canonical values* in [10]. The notion of *non-problematic values* by Cardelli and Frisch [5] is related, but not identical to our lexical values.

$$r_1 \xrightarrow{\_\backslash a} r_2 \xrightarrow{\_\backslash b} r_3 \xrightarrow{\_\backslash c} r_4 \; nullable$$

$$\downarrow mkeps$$

$$v_1 \xleftarrow[inj \; r_1 \; a]{} v_2 \xleftarrow[inj \; r_2 \; b]{} v_3 \xleftarrow[inj \; r_3 \; c]{} v_4$$

Fig. 1. The two phases of the algorithm by Sulzmann & Lu [13], matching the string $[a, b, c]$. The first phase (the arrows from left to right) is Brzozowski's matcher building successive derivatives. If the last regular expression is *nullable*, then the functions of the second phase are called (the top-down and right-to-left arrows): first *mkeps* calculates a value v_4 witnessing how the empty string has been recognised by r_4. After that the function *inj* "injects back" the characters of the string into the values.

The most interesting idea from Sulzmann and Lu [13] is the construction of a value for how r_1 can match the string $[a, b, c]$ from the value how the last derivative, r_4 in Fig. 1, can match the empty string. Sulzmann and Lu achieve this by stepwise "injecting back" the characters into the values thus inverting the operation of building derivatives, but on the level of values. The corresponding function, called *inj*, takes three arguments, a regular expression, a character and a value. For example in the first (or right-most) *inj*-step in Fig. 1 the regular expression r_3, the character c from the last derivative step and v_4, which is the value corresponding to the derivative regular expression r_4. The result is the new value v_3. The final result of the algorithm is the value v_1. The *inj* function is defined by recursion on regular expressions and by analysing the shape of values (corresponding to the derivative regular expressions).

(1) $\quad inj \; d \; c \; (Empty) \qquad\qquad\qquad \overset{\text{def}}{=} \; Char \; d$

(2) $\quad inj \; (r_1 + r_2) \; c \; (Left \; v_1) \qquad\quad \overset{\text{def}}{=} \; Left \; (inj \; r_1 \; c \; v_1)$

(3) $\quad inj \; (r_1 + r_2) \; c \; (Right \; v_2) \qquad \overset{\text{def}}{=} \; Right \; (inj \; r_2 \; c \; v_2)$

(4) $\quad inj \; (r_1 \cdot r_2) \; c \; (Seq \; v_1 \; v_2) \qquad \overset{\text{def}}{=} \; Seq \; (inj \; r_1 \; c \; v_1) \; v_2$

(5) $\quad inj \; (r_1 \cdot r_2) \; c \; (Left \; (Seq \; v_1 \; v_2)) \overset{\text{def}}{=} \; Seq \; (inj \; r_1 \; c \; v_1) \; v_2$

(6) $\quad inj \; (r_1 \cdot r_2) \; c \; (Right \; v_2) \qquad \overset{\text{def}}{=} \; Seq \; (mkeps \; r_1) \; (inj \; r_2 \; c \; v_2)$

(7) $\quad inj \; (r^\star) \; c \; (Seq \; v \; (Stars \; vs)) \quad \overset{\text{def}}{=} \; Stars \; (inj \; r \; c \; v :: vs)$

To better understand what is going on in this definition it might be instructive to look first at the three sequence cases (clauses *(4) – (6)*). In each case we need to construct an "injected value" for $r_1 \cdot r_2$. This must be a value of the form *Seq* _ _. Recall the clause of the *derivative*-function for sequence regular expressions:

$$(r_1 \cdot r_2)\backslash c \overset{\text{def}}{=} \; if \; nullable \; r_1 \; then \; (r_1\backslash c) \cdot r_2 + (r_2\backslash c) \; else \; (r_1\backslash c) \cdot r_2$$

Consider first the *else*-branch where the derivative is $(r_1 \backslash c) \cdot r_2$. The corresponding value must therefore be of the form *Seq* v_1 v_2, which matches the left-hand side in clause *(4)* of *inj*. In the *if*-branch the derivative is an alternative, namely $(r_1 \backslash c) \cdot r_2 + (r_2 \backslash c)$. This means we either have to consider a *Left*- or *Right*-value. In case of the *Left*-value we know further it must be a value for a sequence regular expression. Therefore the pattern we match in the clause *(5)* is *Left* (*Seq* v_1 v_2), while in *(6)* it is just *Right* v_2. One more interesting point is in the right-hand side of clause *(6)*: since in this case the regular expression r_1 does not "contribute" to matching the string, that means it only matches the empty string, we need to call *mkeps* in order to construct a value for how r_1 can match this empty string. A similar argument applies for why we can expect in the left-hand side of clause *(7)* that the value is of the form *Seq* v (*Stars* vs)—the derivative of a star is $(r \backslash c) \cdot r^{\star}$. Finally, the reason for why we can ignore the second argument in clause *(1)* of *inj* is that it will only ever be called in cases where $c = d$, but the usual linearity restrictions in patterns do not allow us to build this constraint explicitly into our function definition.[7]

The idea of the *inj*-function to "inject" a character, say c, into a value can be made precise by the first part of the following lemma, which shows that the underlying string of an injected value has a prepended character c; the second part shows that the underlying string of an *mkeps*-value is always the empty string (given the regular expression is nullable since otherwise *mkeps* might not be defined).

Lemma 9.
(1) If $v : r \backslash c$ then $|inj\ r\ c\ v| = c :: |v|$.
(2) If nullable r then $|mkeps\ r| = []$.

Proof. Both properties are by routine inductions: the first one can, for example, be proved by induction over the definition of *derivatives*; the second by an induction on r. There are no interesting cases.   □

Having defined the *mkeps* and *inj* function we can extend Brzozowski's matcher so that a value is constructed (assuming the regular expression matches the string). The clauses of the Sulzmann and Lu lexer are

$$\begin{aligned}
lexer\ r\ []\quad &\overset{\text{def}}{=}\ if\ nullable\ r\ then\ Some\ (mkeps\ r)\ else\ None \\
lexer\ r\ (c :: s)\ &\overset{\text{def}}{=}\ case\ lexer\ (r \backslash c)\ s\ of \\
&\qquad None \Rightarrow None \\
&\qquad |\ Some\ v \Rightarrow Some\ (inj\ r\ c\ v)
\end{aligned}$$

If the regular expression does not match the string, *None* is returned. If the regular expression *does* match the string, then *Some* value is returned. One important virtue of this algorithm is that it can be implemented with ease in any

---

[7] Sulzmann and Lu state this clause as $inj\ c\ c\ (Empty) \overset{\text{def}}{=}\ Char\ c$, but our deviation is harmless.

$$\frac{}{([], \mathbf{1}) \to Empty}P\mathbf{1} \qquad \frac{}{([c], c) \to Char\ c}Pc$$

$$\frac{(s, r_1) \to v}{(s, r_1 + r_2) \to Left\ v}P{+}L \qquad \frac{(s, r_2) \to v \qquad s \notin L(r_1)}{(s, r_1 + r_2) \to Right\ v}P{+}R$$

$$\frac{\begin{array}{c}(s_1, r_1) \to v_1 \qquad (s_2, r_2) \to v_2 \\ \nexists s_3\ s_4.a.\ s_3 \neq [] \wedge s_3 @ s_4 = s_2 \wedge s_1 @ s_3 \in L(r_1) \wedge s_4 \in L(r_2)\end{array}}{(s_1 @ s_2, r_1 \cdot r_2) \to Seq\ v_1\ v_2}PS$$

$$\frac{}{([], r^{\star}) \to Stars\ []}P[]$$

$$\frac{\begin{array}{c}(s_1, r) \to v \qquad (s_2, r^{\star}) \to Stars\ vs \qquad |v| \neq [] \\ \nexists s_3\ s_4.a.\ s_3 \neq [] \wedge s_3 @ s_4 = s_2 \wedge s_1 @ s_3 \in L(r) \wedge s_4 \in L(r^{\star})\end{array}}{(s_1 @ s_2, r^{\star}) \to Stars\ (v{::}vs)}P\star$$

Fig. 2. Our inductive definition of POSIX values.

functional programming language and also in Isabelle/HOL. In the remaining part of this section we prove that this algorithm is correct.

The well-known idea of POSIX matching is informally defined by some rules such as the Longest Match and Priority Rules (see Introduction); as correctly argued in [13], this needs formal specification. Sulzmann and Lu define an "ordering relation" between values and argue that there is a maximum value, as given by the derivative-based algorithm. In contrast, we shall introduce a simple inductive definition that specifies directly what a *POSIX value* is, incorporating the POSIX-specific choices into the side-conditions of our rules. Our definition is inspired by the matching relation given by Vansummeren [14]. The relation we define is ternary and written as $(s, r) \to v$, relating strings, regular expressions and values; the inductive rules are given in Figure 2. We can prove that given a string s and regular expression r, the POSIX value v is uniquely determined by $(s, r) \to v$.

Theorem 1.
(1) If $(s, r) \to v$ then $s \in L(r)$ and $|v| = s$.
(2) If $(s, r) \to v$ and $(s, r) \to v'$ then $v = v'$.

Proof. Both by induction on the definition of $(s, r) \to v$. The second parts follows by a case analysis of $(s, r) \to v'$ and the first part. □

We claim that our $(s, r) \to v$ relation captures the idea behind the four informal POSIX rules shown in the Introduction: Consider for example the rules $P{+}L$ and $P{+}R$ where the POSIX value for a string and an alternative regular expression, that is $(s, r_1 + r_2)$, is specified—it is always a *Left*-value, *except* when the string to be matched is not in the language of r_1; only then it is a *Right*-value (see the side-condition in $P{+}R$). Interesting is also the rule for sequence regular expressions (PS). The first two premises state that v_1 and v_2 are the POSIX

values for (s_1, r_1) and (s_2, r_2) respectively. Consider now the third premise and note that the POSIX value of this rule should match the string $s_1 @ s_2$. According to the Longest Match Rule, we want that the s_1 is the longest initial split of $s_1 @ s_2$ such that s_2 is still recognised by r_2. Let us assume, contrary to the third premise, that there *exist* an s_3 and s_4 such that s_2 can be split up into a non-empty string s_3 and a possibly empty string s_4. Moreover the longer string $s_1 @ s_3$ can be matched by r_1 and the shorter s_4 can still be matched by r_2. In this case s_1 would *not* be the longest initial split of $s_1 @ s_2$ and therefore *Seq* v_1 v_2 cannot be a POSIX value for $(s_1 @ s_2, r_1 \cdot r_2)$. The main point is that our side-condition ensures the Longest Match Rule is satisfied.

A similar condition is imposed on the POSIX value in the $P\star$-rule. Also there we want that s_1 is the longest initial split of $s_1 @ s_2$ and furthermore the corresponding value v cannot be flattened to the empty string. In effect, we require that in each "iteration" of the star, some non-empty substring needs to be "chipped" away; only in case of the empty string we accept *Stars* $[]$ as the POSIX value. Indeed we can show that our POSIX values are lexical values which exclude those *Stars* that contain subvalues that flatten to the empty string.

Lemma 10. *If* $(s, r) \to v$ *then* $v \in LV\ r\ s$.

Proof. By routine induction on $(s, r) \to v$. □

Next is the lemma that shows the function *mkeps* calculates the POSIX value for the empty string and a nullable regular expression.

Lemma 11. *If nullable* r *then* $([], r) \to mkeps\ r$.

Proof. By routine induction on r. □

The central lemma for our POSIX relation is that the *inj*-function preserves POSIX values.

Lemma 12. *If* $(s, r \backslash c) \to v$ *then* $(c :: s, r) \to inj\ r\ c\ v$.

Proof. By induction on r. We explain two cases.

- Case $r = r_1 + r_2$. There are two subcases, namely (a) $v = Left\ v'$ and $(s, r_1 \backslash c) \to v'$; and (b) $v = Right\ v'$, $s \notin L(r_1 \backslash c)$ and $(s, r_2 \backslash c) \to v'$. In (a) we know $(s, r_1 \backslash c) \to v'$, from which we can infer $(c :: s, r_1) \to inj\ r_1\ c\ v'$ by induction hypothesis and hence $(c :: s, r_1 + r_2) \to inj\ (r_1 + r_2)\ c\ (Left\ v')$ as needed. Similarly in subcase (b) where, however, in addition we have to use Proposition 1(2) in order to infer $c :: s \notin L(r_1)$ from $s \notin L(r_1 \backslash c)$.

- Case $r = r_1 \cdot r_2$. There are three subcases:

    (a) $v = Left\ (Seq\ v_1\ v_2)$ and *nullable* r_1
    (b) $v = Right\ v_1$ and *nullable* r_1
    (c) $v = Seq\ v_1\ v_2$ and $\neg$ *nullable* r_1

For (a) we know $(s_1, r_1 \backslash c) \rightarrow v_1$ and $(s_2, r_2) \rightarrow v_2$ as well as

$$\nexists s_3 \ s_4.a. \ s_3 \neq [] \ \wedge \ s_3 \ @ \ s_4 = s_2 \ \wedge \ s_1 \ @ \ s_3 \ \in \ L(r_1 \backslash c) \ \wedge \ s_4 \ \in \ L(r_2)$$

From the latter we can infer by Proposition 1(2):

$$\nexists s_3 \ s_4.a. \ s_3 \neq [] \ \wedge \ s_3 \ @ \ s_4 = s_2 \ \wedge \ c :: s_1 \ @ \ s_3 \ \in \ L(r_1) \ \wedge \ s_4 \ \in \ L(r_2)$$

We can use the induction hypothesis for r_1 to obtain $(c :: s_1, r_1) \rightarrow inj \ r_1 \ c$ v_1. Putting this all together allows us to infer $(c :: s_1 \ @ \ s_2, r_1 \cdot r_2) \rightarrow Seq$ $(inj \ r_1 \ c \ v_1) \ v_2$. The case (c) is similar.

For (b) we know $(s, r_2 \backslash c) \rightarrow v_1$ and $s_1 \ @ \ s_2 \notin L((r_1 \backslash c) \cdot r_2)$. From the former we have $(c :: s, r_2) \rightarrow inj \ r_2 \ c \ v_1$ by induction hypothesis for r_2. From the latter we can infer

$$\nexists s_3 \ s_4.a. \ s_3 \neq [] \ \wedge \ s_3 \ @ \ s_4 = c :: s \ \wedge \ s_3 \ \in \ L(r_1) \ \wedge \ s_4 \ \in \ L(r_2)$$

By Lemma 11 we know $([], r_1) \rightarrow mkeps \ r_1$ holds. Putting this all together, we can conclude with $(c :: s, r_1 \cdot r_2) \rightarrow Seq \ (mkeps \ r_1) \ (inj \ r_2 \ c \ v_1)$, as required.

Finally suppose $r = r_1^\star$. This case is very similar to the sequence case, except that we need to also ensure that $|inj \ r_1 \ c \ v_1| \neq []$. This follows from $(c :: s_1, r_1) \rightarrow inj \ r_1 \ c \ v_1$ (which in turn follows from $(s_1, r_1 \backslash c) \rightarrow v_1$ and the induction hypothesis). $\qquad\square$

With Lemma 12 in place, it is completely routine to establish that the Sulzmann and Lu lexer satisfies our specification (returning the null value *None* iff the string is not in the language of the regular expression, and returning a unique POSIX value iff the string *is* in the language):

Theorem 2.
(1) $s \notin L(r)$ if and only if lexer $r \ s = None$
(2) $s \in L(r)$ if and only if $\exists v.$ lexer $r \ s = Some \ v \wedge (s, r) \rightarrow v$

Proof. By induction on s using Lemma 11 and 12. $\qquad\square$

In *(2)* we further know by Theorem 1 that the value returned by the lexer must be unique. A simple corollary of our two theorems is:

Corollary 1.
(1) lexer $r \ s = None$ if and only if $\nexists v.a. \ (s, r) \rightarrow v$
(2) lexer $r \ s = Some \ v$ if and only if $(s, r) \rightarrow v$

This concludes our correctness proof. Note that we have not changed the algorithm of Sulzmann and Lu,[8] but introduced our own specification for what a correct result—a POSIX value—should be. In the next section we show that our specification coincides with another one given by Okui and Suzuki using a different technique.

---

[8] All deviations we introduced are harmless.

6   Ordering of Values according to Okui and Suzuki

While in the previous section we have defined POSIX values directly in terms of a ternary relation (see inference rules in Figure 2), Sulzmann and Lu took a different approach in [13]: they introduced an ordering for values and identified POSIX values as the maximal elements. An extended version of [13] is available at the website of its first author; this includes more details of their proofs, but which are evidently not in final form yet. Unfortunately, we were not able to verify claims that their ordering has properties such as being transitive or having maximal elements.

Okui and Suzuki [10,11] described another ordering of values, which they use to establish the correctness of their automata-based algorithm for POSIX matching. Their ordering resembles some aspects of the one given by Sulzmann and Lu, but overall is quite different. To begin with, Okui and Suzuki identify POSIX values as minimal, rather than maximal, elements in their ordering. A more substantial difference is that the ordering by Okui and Suzuki uses *positions* in order to identify and compare subvalues. Positions are lists of natural numbers. This allows them to quite naturally formalise the Longest Match and Priority rules of the informal POSIX standard. Consider for example the value v

$$v \stackrel{def}{=} Stars\ [Seq\ (Char\ x)\ (Char\ y),\ Char\ z]$$

At position $[0,1]$ of this value is the subvalue $Char\ y$ and at position $[1]$ the subvalue $Char\ z$. At the 'root' position, or empty list $[]$, is the whole value v. Positions such as $[0,1,0]$ or $[2]$ are outside of v. If it exists, the subvalue of v at a position p, written $v\!\downarrow_p$, can be recursively defined by

$$
\begin{aligned}
v\!\downarrow_{[]} &\stackrel{def}{=} v \\
Left\ v\!\downarrow_{0::ps} &\stackrel{def}{=} v\!\downarrow_{ps} \\
Right\ v\!\downarrow_{1::ps} &\stackrel{def}{=} v\!\downarrow_{ps} \\
Seq\ v_1\ v_2\!\downarrow_{0::ps} &\stackrel{def}{=} v_1\!\downarrow_{ps} \\
Seq\ v_1\ v_2\!\downarrow_{1::ps} &\stackrel{def}{=} v_2\!\downarrow_{ps} \\
Stars\ vs\!\downarrow_{n::ps} &\stackrel{def}{=} vs_{[n]}\!\downarrow_{ps}
\end{aligned}
$$

In the last clause we use Isabelle's notation $vs_{[n]}$ for the nth element in a list. The set of positions inside a value v, written $Pos\ v$, is given by

$$
\begin{aligned}
Pos\ (Empty) &\stackrel{def}{=} \{[]\} \\
Pos\ (Char\ c) &\stackrel{def}{=} \{[]\} \\
Pos\ (Left\ v) &\stackrel{def}{=} \{[]\} \cup \{0::ps \mid ps \in Pos\ v\} \\
Pos\ (Right\ v) &\stackrel{def}{=} \{[]\} \cup \{1::ps \mid ps \in Pos\ v\} \\
Pos\ (Seq\ v_1\ v_2) &\stackrel{def}{=} \{[]\} \cup \{0::ps \mid ps \in Pos\ v_1\} \cup \{1::ps \mid ps \in Pos\ v_2\} \\
Pos\ (Stars\ vs) &\stackrel{def}{=} \{[]\} \cup \left(\bigcup n < len\ vs\ \{n::ps \mid ps \in Pos\ vs_{[n]}\}\right)
\end{aligned}
$$

whereby *len* in the last clause stands for the length of a list. Clearly for every position inside a value there exists a subvalue at that position.

To help understanding the ordering of Okui and Suzuki, consider again the earlier value v and compare it with the following w:

$$v \stackrel{def}{=} Stars\,[Seq\,(Char\;x)\,(Char\;y),\,Char\;z]$$
$$w \stackrel{def}{=} Stars\,[Char\;x,\,Char\;y,\,Char\;z]$$

Both values match the string *xyz*, that means if we flatten these values at their respective root position, we obtain *xyz*. However, at position $[0]$, v matches *xy* whereas w matches only the shorter *x*. So according to the Longest Match Rule, we should prefer v, rather than w as POSIX value for string *xyz* (and corresponding regular expression). In order to formalise this idea, Okui and Suzuki introduce a measure for subvalues at position p, called the *norm* of v at position p. We can define this measure in Isabelle as an integer as follows

$$\|v\|_p \stackrel{def}{=} if\;p\,\in\,Pos\;v\;then\;len\;|v{\downarrow}_p|\;else - 1$$

where we take the length of the flattened value at position p, provided the position is inside v; if not, then the norm is -1. The default for outside positions is crucial for the POSIX requirement of preferring a *Left*-value over a *Right*-value (if they can match the same string—see the Priority Rule from the Introduction). For this consider

$$v \stackrel{def}{=} Left\,(Char\;x) \qquad and \qquad w \stackrel{def}{=} Right\,(Char\;x)$$

Both values match *x*. At position $[0]$ the norm of v is 1 (the subvalue matches *x*), but the norm of w is -1 (the position is outside w according to how we defined the 'inside' positions of *Left*- and *Right*-values). Of course at position $[1]$, the norms $\|v\|_{[1]}$ and $\|w\|_{[1]}$ are reversed, but the point is that subvalues will be analysed according to lexicographically ordered positions. According to this ordering, the position $[0]$ takes precedence over $[1]$ and thus also v will be preferred over w. The lexicographic ordering of positions, written $\_ \prec_{lex} \_$, can be conveniently formalised by three inference rules

$$\frac{}{[\,]\;\prec_{lex}\;p::ps} \qquad\qquad \frac{p_1\,<\,p_2}{p_1::ps_1\;\prec_{lex}\;p_2::ps_2} \qquad\qquad \frac{ps_1\;\prec_{lex}\;ps_2}{p::ps_1\;\prec_{lex}\;p::ps_2}$$

With the norm and lexicographic order in place, we can state the key definition of Okui and Suzuki [10]: a value v_1 is *smaller at position p* than v_2, written $v_1 \prec_p v_2$, if and only if (i) the norm at position p is greater in v_1 (that is the string $|v_1{\downarrow}_p|$ is longer than $|v_2{\downarrow}_p|$) and (ii) all subvalues at positions that are inside v_1 or v_2 and that are lexicographically smaller than p, we have the same norm, namely

$$v_1\;\prec_p\;v_2 \stackrel{def}{=} \begin{cases} (i) & \|v_2\|_p\,<\,\|v_1\|_p \quad and \\ (ii) & \forall\,q\in Pos\;v_1\,\cup\,Pos\;v_2.\;q\prec_{lex}\,p \longrightarrow \|v_1\|_q\,=\,\|v_2\|_q \end{cases}$$

The position p in this definition acts as the *first distinct position* of v_1 and v_2, where both values match strings of different length [10]. Since at p the values v_1 and v_2 match different strings, the ordering is irreflexive. Derived from the definition above are the following two orderings:

$$v_1 \prec v_2 \overset{def}{=} \exists\, p.\ v_1 \prec_p v_2$$
$$v_1 \preccurlyeq v_2 \overset{def}{=} v_1 \prec v_2 \vee v_1 = v_2$$

While we encountered a number of obstacles for establishing properties like transitivity for the ordering of Sulzmann and Lu (and which we failed to overcome), it is relatively straightforward to establish this property for the orderings $\_ \prec \_$ and $\_ \preccurlyeq \_$ by Okui and Suzuki.

Lemma 13 (Transitivity). *If* $v_1 \prec v_2$ *and* $v_2 \prec v_3$ *then* $v_1 \prec v_3$.

Proof. From the assumption we obtain two positions p and q, where the values v_1 and v_2 (respectively v_2 and v_3) are 'distinct'. Since $\prec_{lex}$ is trichotomous, we need to consider three cases, namely $p = q$, $p \prec_{lex} q$ and $q \prec_{lex} p$. Let us look at the first case. Clearly $\|v_2\|_p < \|v_1\|_p$ and $\|v_3\|_p < \|v_2\|_p$ imply $\|v_3\|_p < \|v_1\|_p$. It remains to show that for a $p' \in Pos\ v_1 \cup Pos\ v_3$ with $p' \prec_{lex} p$ that $\|v_1\|_{p'} = \|v_3\|_{p'}$ holds. Suppose $p' \in Pos\ v_1$, then we can infer from the first assumption that $\|v_1\|_{p'} = \|v_2\|_{p'}$. But this means that p' must be in $Pos\ v_2$ too (the norm cannot be -1 given $p' \in Pos\ v_1$). Hence we can use the second assumption and infer $\|v_2\|_{p'} = \|v_3\|_{p'}$, which concludes this case with $v_1 \prec v_3$. The reasoning in the other cases is similar. □

The proof for $\preccurlyeq$ is similar and omitted. It is also straightforward to show that $\prec$ and $\preccurlyeq$ are partial orders. Okui and Suzuki furthermore show that they are linear orderings for lexical values [10] of a given regular expression and given string, but we have not formalised this in Isabelle. It is not essential for our results. What we are going to show below is that for a given r and s, the orderings have a unique minimal element on the set $LV\ r\ s$, which is the POSIX value we defined in the previous section. We start with two properties that show how the length of a flattened value relates to the $\prec$-ordering.

Proposition 4.

(1) If $v_1 \prec v_2$ *then* $len\ |v_2| \leq len\ |v_1|$.
(2) If $len\ |v_2| < len\ |v_1|$ *then* $v_1 \prec v_2$.

Both properties follow from the definition of the ordering. Note that *(2)* entails that a value, say v_2, whose underlying string is a strict prefix of another flattened value, say v_1, then v_1 must be smaller than v_2. For our proofs it will be useful to have the following properties—in each case the underlying strings of the compared values are the same:

Proposition 5.

(1) If $|v_1| = |v_2|$ then *Left* $v_1 \prec$ *Right* v_2.

(2) If $|v_1| = |v_2|$ then *Left* $v_1 \prec$ *Left* v_2 iff $v_1 \prec v_2$

(3) If $|v_1| = |v_2|$ then *Right* $v_1 \prec$ *Right* v_2 iff $v_1 \prec v_2$

(4) If $|v_2| = |w_2|$ then *Seq* v $v_2 \prec$ *Seq* v w_2 iff $v_2 \prec w_2$

(5) If $|v_1|$ @ $|v_2| = |w_1|$ @ $|w_2|$ and $v_1 \prec w_1$ then *Seq* v_1 $v_2 \prec$ *Seq* w_1 w_2

(6) If $|vs_1| = |vs_2|$ then *Stars* $(vs$ @ $vs_1) \prec$ *Stars* $(vs$ @ $vs_2)$ iff *Stars* $vs_1 \prec$ *Stars* vs_2

(7) If $|v_1 :: vs_1| = |v_2 :: vs_2|$ and $v_1 \prec v_2$ then *Stars* $(v_1 :: vs_1) \prec$ *Stars* $(v_2 :: vs_2)$

One might prefer that statements *(4)* and *(5)* (respectively *(6)* and *(7)*) are combined into a single *iff*-statement (like the ones for *Left* and *Right*). Unfortunately this cannot be done easily: such a single statement would require an additional assumption about the two values *Seq* v_1 v_2 and *Seq* w_1 w_2 being inhabited by the same regular expression. The complexity of the proofs involved seems to not justify such a 'cleaner' single statement. The statements given are just the properties that allow us to establish our theorems without any difficulty. The proofs for Proposition 5 are routine.

Next we establish how Okui and Suzuki's orderings relate to our definition of POSIX values. Given a *POSIX* value v_1 for r and s, then any other lexical value v_2 in *LV* r s is greater or equal than v_1, namely:

Theorem 3. *If* $(s, r) \to v_1$ *and* $v_2 \in LV$ r s *then* $v_1 \preccurlyeq v_2$.

Proof. By induction on our POSIX rules. By Theorem 1 and the definition of *LV*, it is clear that v_1 and v_2 have the same underlying string s. The three base cases are straightforward: for example for $v_1 = $ *Empty*, we have that $v_2 \in LV$ **1** [] must also be of the form $v_2 = $ *Empty*. Therefore we have $v_1 \preccurlyeq v_2$. The inductive cases for r being of the form $r_1 + r_2$ and $r_1 \cdot r_2$ are as follows:

- Case *P+L* with $(s, r_1 + r_2) \to$ *Left* w_1: In this case the value v_2 is either of the form *Left* w_2 or *Right* w_2. In the latter case we can immediately conclude with $v_1 \preccurlyeq v_2$ since a *Left*-value with the same underlying string s is always smaller than a *Right*-value by Proposition 5*(1)*. In the former case we have $w_2 \in LV$ r_1 s and can use the induction hypothesis to infer $w_1 \preccurlyeq w_2$. Because w_1 and w_2 have the same underlying string s, we can conclude with *Left* $w_1 \preccurlyeq$ *Left* w_2 using Proposition 5*(2)*.

- Case *P+R* with $(s, r_1 + r_2) \to$ *Right* w_1: This case similar to the previous case, except that we additionally know $s \notin L(r_1)$. This is needed when v_2 is of the form *Left* w_2. Since $|v_2| = |w_2| = s$ and $w_2 : r_1$, we can derive a contradiction for $s \notin L(r_1)$ using Proposition 2. So also in this case $v_1 \preccurlyeq v_2$.

- Case *PS* with $(s_1$ @ $s_2, r_1 \cdot r_2) \to$ *Seq* w_1 w_2: We can assume $v_2 = $ *Seq* u_1 u_2 with $u_1 : r_1$ and $u_2 : r_2$. We have s_1 @ $s_2 = |u_1|$ @ $|u_2|$. By the side-condition of the *PS*-rule we know that either $s_1 = |u_1|$ or that $|u_1|$ is a strict prefix of s_1. In the latter case we can infer $w_1 \prec u_1$ by Proposition 4*(2)* and from this $v_1 \preccurlyeq v_2$ by Proposition 5*(5)* (as noted above v_1 and v_2 must have the same underlying string). In the former case we know $u_1 \in LV$ r_1 s_1 and $u_2 \in LV$ r_2 s_2. With this we can use the induction hypotheses to infer $w_1 \preccurlyeq u_1$ and $w_2 \preccurlyeq u_2$. By Proposition 5*(4,5)* we can again infer $v_1 \preccurlyeq v_2$.

The case for $P\star$ is similar to the *PS*-case and omitted.   □

This theorem shows that our *POSIX* value for a regular expression r and string s is in fact a minimal element of the values in *LV r s*. By Proposition 4*(2)* we also know that any value in *LV r s′*, with $s′$ being a strict prefix, cannot be smaller than v_1. The next theorem shows the opposite—namely any minimal element in *LV r s* must be a *POSIX* value. This can be established by induction on r, but the proof can be drastically simplified by using the fact from the previous section about the existence of a *POSIX* value whenever a string $s \in L(r)$.

Theorem 4. *If $v_1 \in LV\ r\ s$ and $\forall\, v_2 \in LV\ r\ s.\ v_2 \nprec v_1$ then $(s,\ r) \to v_1$.*

Proof. If $v_1 \in LV\ r\ s$ then $s \in L(r)$ by Proposition 2. Hence by Theorem 2(2) there exists a *POSIX* value v_P with $(s,\ r) \to v_P$ and by Lemma 10 we also have $v_P \in LV\ r\ s$. By Theorem 3 we therefore have $v_P \preccurlyeq v_1$. If $v_P = v_1$ then we are done. Otherwise we have $v_P \prec v_1$, which however contradicts the second assumption about v_1 being the smallest element in *LV r s*. So we are done in this case too.   □

From this we can also show that if *LV r s* is non-empty (or equivalently $s \in L(r)$) then it has a unique minimal element:

Corollary 2. *If $LV\ r\ s \neq \varnothing$ then $\exists! vmin.\ vmin \in LV\ r\ s \wedge (\forall\, v \in LV\ r\ s.\ vmin \preccurlyeq v)$.*

To sum up, we have shown that the (unique) minimal elements of the ordering by Okui and Suzuki are exactly the *POSIX* values we defined inductively in Section 5. This provides an independent confirmation that our ternary relation formalises the informal POSIX rules.

7   Bitcoded Lexing

Incremental calculation of the value. To simplify the proof we first define the function *flex* which calculates the "iterated" injection function. With this we can rewrite the lexer as

lexer r s = (if nullable (r\s) then Some (flex r id s (mkeps (r\s))) else None)

8   Optimisations

Derivatives as calculated by Brzozowski's method are usually more complex regular expressions than the initial one; the result is that the derivative-based matching and lexing algorithms are often abysmally slow. However, various optimisations are possible, such as the simplifications of $\mathbf{0} + r$, $r + \mathbf{0}$, $\mathbf{1} \cdot r$ and $r \cdot \mathbf{1}$ to r. These simplifications can speed up the algorithms considerably, as noted

in [13]. One of the advantages of having a simple specification and correctness proof is that the latter can be refined to prove the correctness of such simplification steps. While the simplification of regular expressions according to rules like

$$\mathbf{0} + r \Rightarrow r \qquad r + \mathbf{0} \Rightarrow r \qquad \mathbf{1} \cdot r \Rightarrow r \qquad r \cdot \mathbf{1} \Rightarrow r \qquad (2)$$

is well understood, there is an obstacle with the POSIX value calculation algorithm by Sulzmann and Lu: if we build a derivative regular expression and then simplify it, we will calculate a POSIX value for this simplified derivative regular expression, *not* for the original (unsimplified) derivative regular expression. Sulzmann and Lu [13] overcome this obstacle by not just calculating a simplified regular expression, but also calculating a *rectification function* that "repairs" the incorrect value.

The rectification functions can be (slightly clumsily) implemented in Isabelle/HOL as follows using some auxiliary functions:

$$
\begin{aligned}
F_{Right} \; f \; v & \stackrel{\text{def}}{=} Right \; (f \; v) \\
F_{Left} \; f \; v & \stackrel{\text{def}}{=} Left \; (f \; v) \\
F_{Alt} \; f_1 \; f_2 \; (Right \; v) & \stackrel{\text{def}}{=} Right \; (f_2 \; v) \\
F_{Alt} \; f_1 \; f_2 \; (Left \; v) & \stackrel{\text{def}}{=} Left \; (f_1 \; v) \\
F_{Seq1} \; f_1 \; f_2 \; v & \stackrel{\text{def}}{=} Seq \; (f_1 \; ()) \; (f_2 \; v) \\
F_{Seq2} \; f_1 \; f_2 \; v & \stackrel{\text{def}}{=} Seq \; (f_1 \; v) \; (f_2 \; ()) \\
F_{Seq} \; f_1 \; f_2 \; (Seq \; v_1 \; v_2) & \stackrel{\text{def}}{=} Seq \; (f_1 \; v_1) \; (f_2 \; v_2) \\[6pt]
simp_{Alt} \; (\mathbf{0}, \_) \; (r_2, f_2) & \stackrel{\text{def}}{=} (r_2, \; F_{Right} \; f_2) \\
simp_{Alt} \; (r_1, f_1) \; (\mathbf{0}, \_) & \stackrel{\text{def}}{=} (r_1, \; F_{Left} \; f_1) \\
simp_{Alt} \; (r_1, f_1) \; (r_2, f_2) & \stackrel{\text{def}}{=} (r_1 + r_2, \; F_{Alt} \; f_1 \; f_2) \\
simp_{Seq} \; (\mathbf{1}, f_1) \; (r_2, f_2) & \stackrel{\text{def}}{=} (r_2, \; F_{Seq1} \; f_1 \; f_2) \\
simp_{Seq} \; (r_1, f_1) \; (\mathbf{1}, f_2) & \stackrel{\text{def}}{=} (r_1, \; F_{Seq2} \; f_1 \; f_2) \\
simp_{Seq} \; (r_1, f_1) \; (r_2, f_2) & \stackrel{\text{def}}{=} (r_1 \cdot r_2, \; F_{Seq} \; f_1 \; f_2)
\end{aligned}
$$

The functions $simp_{Alt}$ and $simp_{Seq}$ encode the simplification rules in (2) and compose the rectification functions (simplifications can occur deep inside the regular expression). The main simplification function is then

$$
\begin{aligned}
simp \; (r_1 + r_2) & \stackrel{\text{def}}{=} simp_{Alt} \; (simp \; r_1) \; (simp \; r_2) \\
simp \; (r_1 \cdot r_2) & \stackrel{\text{def}}{=} simp_{Seq} \; (simp \; r_1) \; (simp \; r_2) \\
simp \; r & \stackrel{\text{def}}{=} (r, \; id)
\end{aligned}
$$

where *id* stands for the identity function. The function *simp* returns a simplified regular expression and a corresponding rectification function. Note that we do

not simplify under stars: this seems to slow down the algorithm, rather than speed it up. The optimised lexer is then given by the clauses:

$$lexer^+ \; r \; [] \qquad \overset{\text{def}}{=} \; \textit{if nullable r then Some (mkeps r) else None}$$

$$lexer^+ \; r \; (c::s) \; \overset{\text{def}}{=} \; \textit{let } (r_s, f_r) = simp \; (r\backslash c) \textit{ in}$$
$$\qquad \textit{case lexer}^+ \; r_s \; s \textit{ of}$$
$$\qquad \textit{None} \Rightarrow \textit{None}$$
$$\qquad | \; \textit{Some } v \Rightarrow \textit{Some (inj r c (f_r v))}$$

In the second clause we first calculate the derivative $r\backslash c$ and then simpli
    text   *Incremental calculation of the value. To simplify the proof we first define the function @{const flex} which calculates the ''iterated'' injection function. With this we can rewrite the lexer as* \begin{center} @{thm lexer__flex} \end{center} \begin{center} \begin{tabular}{lcl} @{thm (lhs) code.simps(1)} & \$\dn\$ & @{thm (rhs) code.simps(1)}\\ @{thm (lhs) code.simps(2)} & \$\dn\$ & @{thm (rhs) code.simps(2)}\\ @{thm (lhs) code.simps(3)} & \$\dn\$ & @{thm (rhs) code.simps(3)}\\ @{thm (lhs) code.simps(4)} & \$\dn\$ & @{thm (rhs) code.simps(4)}\\ @{thm (lhs) code.simps(5)[of v_1 v_2]} & \$\dn\$ & @{thm (rhs) code.simps(5)[of v_1 v_2]}\\ @{thm (lhs) code.simps(6)} & \$\dn\$ & @{thm (rhs) code.simps(6)}\\ @{thm (lhs) code.simps(7)} & \$\dn\$ & @{thm (rhs) code.simps(7)} \end{tabular} \end{center} \begin{center} \begin{tabular}{lcl} @{term areg} & \$::=\$ & @{term AZERO}\\ & \$\mid\$ & @{term AONE bs}\\ & \$\mid\$ & @{term ACHAR bs c}\\ & \$\mid\$ & @{term AALT bs r1 r2}\\ & \$\mid\$ & @{term ASEQ bs r_1 r_2}\\ & \$\mid\$ & @{term ASTAR bs r} \end{tabular} \end{center} \begin{center} \begin{tabular}{lcl} @{thm (lhs) intern.simps(1)} & \$\dn\$ & @{thm (rhs) intern.simps(1)}\\ @{thm (lhs) intern.simps(2)} & \$\dn\$ & @{thm (rhs) intern.simps(2)}\\ @{thm (lhs) intern.simps(3)} & \$\dn\$ & @{thm (rhs) intern.simps(3)}\\ @{thm (lhs) intern.simps(4)[of r_1 r_2]} & \$\dn\$ & @{thm (rhs) intern.simps(4)[of r_1 r_2]}\\ @{thm (lhs) intern.simps(5)[of r_1 r_2]} & \$\dn\$ & @{thm (rhs) intern.simps(5)[of r_1 r_2]}\\ @{thm (lhs) intern.simps(6)} & \$\dn\$ & @{thm (rhs) intern.simps(6)}\\ \end{tabular} \end{center} \begin{center} \begin{tabular}{lcl} @{thm (lhs) erase.simps(1)} & \$\dn\$ & @{thm (rhs) erase.simps(1)}\\ @{thm (lhs) erase.simps(2)[of bs]} & \$\dn\$ & @{thm (rhs) erase.simps(2)[of bs]}\\ @{thm (lhs) erase.simps(3)[of bs]} & \$\dn\$ & @{thm (rhs) erase.simps(3)[of bs]}\\ @{thm (lhs) erase.simps(4)[of bs r_1 r_2]} & \$\dn\$ & @{thm (rhs) erase.simps(4)[of bs r_1 r_2]}\\ @{thm (lhs) erase.simps(5)[of bs r_1 r_2]} & \$\dn\$ & @{thm (rhs) erase.simps(5)[of bs r_1 r_2]}\\ @{thm (lhs) erase.simps(6)[of bs]} & \$\dn\$ & @{thm (rhs) erase.simps(6)[of bs]}\\ \end{tabular} \end{center} *Some simple facts about erase* \begin{lemma}\mbox{}\\ @{thm erase__bder}\\ @{thm erase__intern} \end{lemma} \begin{center} \begin{tabular}{lcl} @{thm (lhs) bnullable.simps(1)} & \$\dn\$ & @{thm (rhs) bnullable.simps(1)}\\ @{thm (lhs) bnullable.simps(2)} & \$\dn\$ & @{thm (rhs) bnullable.simps(2)}\\ @{thm (lhs) bnullable.simps(3)} & \$\dn\$ & @{thm (rhs) bnullable.simps(3)}\\ @{thm (lhs) bnullable.simps(4)[of bs r_1 r_2]} & \$\dn\$ & @{thm (rhs) bnullable.simps(4)[of bs r_1 r_2]}\\ @{thm (lhs) bnullable.simps(5)[of

bs r_1 r_2]$\}$ & $\$\backslash dn\$$ & @$\{thm$ (rhs) $bnullable.simps(5)[of$ bs r_1 r_2]$\}\backslash\backslash$ @$\{thm$ (lhs) $bnullable.simps(6)\}$ & $\$\backslash dn\$$ & @$\{thm$ (rhs) $bnullable.simps(6)\}\backslash medskip\backslash\backslash$ % $\backslash end\{tabular\}$ % $\backslash end\{center\}$ % $\backslash begin\{center\}$ % $\backslash begin\{tabular\}\{lcl\}$ @$\{thm$ (lhs) $bder.simps(1)\}$ & $\$\backslash dn\$$ & @$\{thm$ (rhs) $bder.simps(1)\}\backslash\backslash$ @$\{thm$ (lhs) $bder.simps(2)\}$ & $\$\backslash dn\$$ & @$\{thm$ (rhs) $bder.simps(2)\}\backslash\backslash$ @$\{thm$ (lhs) $bder.simps(3)\}$ & $\$\backslash dn\$$ & @$\{thm$ (rhs) $bder.simps(3)\}\backslash\backslash$ @$\{thm$ (lhs) $bder.simps(4)[of$ bs r_1 r_2]$\}$ & $\$\backslash dn\$$ & @$\{thm$ (rhs) $bder.simps(4)[of$ bs r_1 r_2]$\}\backslash\backslash$ @$\{thm$ (lhs) $bder.simps(5)[of$ bs r_1 r_2]$\}$ & $\$\backslash dn\$$ & @$\{thm$ (rhs) $bder.simps(5)[of$ bs r_1 r_2]$\}\backslash\backslash$ @$\{thm$ (lhs) $bder.simps(6)\}$ & $\$\backslash dn\$$ & @$\{thm$ (rhs) $bder.simps(6)\}$ $\backslash end\{tabular\}$ $\backslash end\{center\}$ $\backslash begin\{center\}$ $\backslash begin\{tabular\}\{lcl\}$ @$\{thm$ (lhs) $bmkeps.simps(1)\}$ & $\$\backslash dn\$$ & @$\{thm$ (rhs) $bmkeps.simps(1)\}\backslash\backslash$ @$\{thm$ (lhs) $bmkeps.simps(2)[of$ bs r_1 r_2]$\}$ & $\$\backslash dn\$$ & @$\{thm$ (rhs) $bmkeps.simps(2)[of$ bs r_1 r_2]$\}\backslash\backslash$ @$\{thm$ (lhs) $bmkeps.simps(3)[of$ bs r_1 r_2]$\}$ & $\$\backslash dn\$$ & @$\{thm$ (rhs) $bmkeps.simps(3)[of$ bs r_1 r_2]$\}\backslash\backslash$ @$\{thm$ (lhs) $bmkeps.simps(4)\}$ & $\$\backslash dn\$$ & @$\{thm$ (rhs) $bmkeps.simps(4)\}\backslash medskip\backslash\backslash$ $\backslash end\{tabular\}$ $\backslash end\{center\}$ @$\{thm$ [$mode=IfThen$] $bder\_\_retrieve\}$ By $induction$ on ‹r› $\backslash begin\{theorem\}[Main$ $Lemma]\backslash mbox\{\}\backslash\backslash$ @$\{thm$ [$mode=IfThen$] $MAIN\_\_decode\}$ $\backslash end\{theorem\}$ $\backslash noindent$ $Definition$ of the $bitcoded$ $lexer$ @$\{thm$ $blexer\_\_def\}$ $\backslash begin\{theorem\}$ @$\{thm$ $blexer\_\_correctness\}$ $\backslash end\{theorem\}$

section *Optimisations*

text   *Derivatives as calculated by* $\backslash Brz's$ *method are usually more complex regular expressions than the initial one; the result is that the derivative−based matching and lexing algorithms are often abysmally slow. However, various optimisations are possible, such as the simplifications of* @$\{term$ ALT $ZERO$ $r\}$, @$\{term$ ALT r $ZERO\}$, @$\{term$ SEQ ONE $r\}$ *and* @$\{term$ SEQ r $ONE\}$ *to* @$\{term$ $r\}$. *These simplifications can speed up the algorithms considerably, as noted in* $\backslash cite\{Sulzmann2014\}$. *One of the advantages of having a simple specification and correctness proof is that the latter can be refined to prove the correctness of such simplification steps. While the simplification of regular expressions according to rules like* $\backslash begin\{equation\}\backslash label\{Simpl\}$ $\backslash begin\{array\}\{lcllcllcllcl\}$ @$\{term$ ALT $ZERO$ $r\}$ & ‹⇒› & @$\{term$ $r\}$ $\backslash hspace\{8mm\}\%\backslash\backslash$ @$\{term$ ALT r $ZERO\}$ & ‹⇒› & @$\{term$ $r\}$ $\backslash hspace\{8mm\}\%\backslash\backslash$ @$\{term$ SEQ ONE $r\}$ & ‹⇒› & @$\{term$ $r\}$ $\backslash hspace\{8mm\}\%\backslash\backslash$ @$\{term$ SEQ r $ONE\}$ & ‹⇒› & @$\{term$ $r\}$ $\backslash end\{array\}$ $\backslash end\{equation\}$ $\backslash noindent$ *is well understood, there is an obstacle with the POSIX value calculation algorithm by Sulzmann and Lu: if we build a derivative regular expression and then simplify it, we will calculate a POSIX value for this simplified derivative regular expression,* $\backslash emph\{not\}$ *for the original (unsimplified) derivative regular expression. Sulzmann and Lu* $\backslash cite\{Sulzmann2014\}$ *overcome this obstacle by not just calculating a simplified regular expression, but also calculating a* $\backslash emph\{rectification$ $function\}$ *that ''repairs'' the incorrect value. The rectification functions can be (slightly clumsily) implemented in Isabelle/HOL as follows using some auxiliary functions*: $\backslash begin\{center\}$ $\backslash begin\{tabular\}\{lcl\}$ @$\{thm$ (lhs) $F\_\_RIGHT.simps(1)\}$ & $\$\backslash dn\$$ & ‹$Right$ (f v)›$\backslash\backslash$ @$\{thm$ (lhs) $F\_\_LEFT.simps(1)\}$ & $\$\backslash dn\$$ & ‹$Left$ (f v)›$\backslash\backslash$ @$\{thm$ (lhs) $F\_\_ALT.simps(1)\}$ & $\$\backslash dn\$$ & ‹$Right$ (f_2 v)›$\backslash\backslash$

@{thm (lhs) F___ALT.simps(2)} & $\dn$ & ‹Left (f₁ v)›\\ @{thm (lhs) F___SEQ1.simps(1)} & $\dn$ & ‹Seq (f₁ ()) (f₂ v)›\\ @{thm (lhs) F___SEQ2.simps(1)} & $\dn$ & ‹Seq (f₁ v) (f₂ ())›\\ @{thm (lhs) F___SEQ.simps(1)} & $\dn$ & ‹Seq (f₁ v₁) (f₂ v₂)›\medskip\\ %\end{tabular} % %\begin{tabular}{lcl} @{term simp___ALT (ZERO, DUMMY) (r₂, f₂)} & $\dn$ & @{term (r₂, F___RIGHT f₂)}\\ @{term simp___ALT (r₁,f₁) (ZERO, DUMMY)} & $\dn$ & @{term (r₁, F___LEFT f₁)}\\ @{term simp___ALT (r₁, f₁) (r₂, f₂)} & $\dn$ & @{term (ALT r₁ r₂, F___ALT f₁ f₂)}\\ @{term simp___SEQ (ONE, f₁) (r₂, f₂)} & $\dn$ & @{term (r₂, F___SEQ1 f₁ f₂)}\\ @{term simp___SEQ (r₁, f₁) (ONE, f₂)} & $\dn$ & @{term (r₁, F___SEQ2 f₁ f₂)}\\ @{term simp___SEQ (r₁, f₁) (r₂, f₂)} & $\dn$ & @{term (SEQ r₁ r₂, F___SEQ f₁ f₂)}\\ \end{tabular} \end{center} \noindent The functions ‹simp$_{Alt}$› and ‹simp$_{Seq}$› encode the simplification rules in \eqref{Simpl} and compose the rectification functions (simplifications can occur deep inside the regular expression). The main simplification function is then \begin{center} \begin{tabular}{lcl} @{term simp (ALT r₁ r₂)} & $\dn$ & @{term simp___ALT (simp r₁) (simp r₂)}\\ @{term simp (SEQ r₁ r₂)} & $\dn$ & @{term simp___SEQ (simp r₁) (simp r₂)}\\ @{term simp r} & $\dn$ & @{term (r, id)}\\ \end{tabular} \end{center} \noindent where @{term id} stands for the identity function. The function @{const simp} returns a simplified regular expression and a corresponding rectification function. Note that we do not simplify under stars: this seems to slow down the algorithm, rather than speed it up. The optimised lexer is then given by the clauses: \begin{center} \begin{tabular}{lcl} @{thm (lhs) slexer.simps(1)} & $\dn$ & @{thm (rhs) slexer.simps(1)}\\ @{thm (lhs) slexer.simps(2)} & $\dn$ & ‹let (r$_s$, f$_r$) = simp (r ›$\backslash$‹ c) in›\\ & & ‹case› @{term slexer r$_s$ s} ‹of›\\ & & @{term None} ‹⇒› @{term None}\\ & & $\|$ @{term Some v} ‹⇒› ‹Some (inj r c (f$_r$ v))› \end{tabular} \end{center} \noindent In the second clause we first calculate the derivative @{term der c r} and then simplify the result. This gives us a simplified derivative ‹r$_s$› and a rectification function ‹f$_r$›. The lexer is then recursively called with the simplified derivative, but before we inject the character @{term c} into the value @{term v}, we need to rectify @{term v} (that is construct @{term f$_r$ v}). Before we can establish the correctness of @{term slexer}, we need to show that simplification preserves the language and simplification preserves our POSIX relation once the value is rectified (recall @{const simp} generates a (regular expression, rectification function) pair): \begin{lemma}\mbox{}\smallskip\\\label{slexeraux} \begin{tabular}{ll} (1) & @{thm L___fst___simp[symmetric]}\\ (2) & @{thm[mode=IfThen] Posix___simp} \end{tabular} \end{lemma} \begin{proof} Both are by induction on ‹r›. There is no interesting case for the first statement. For the second statement, of interest are the @{term r = ALT r₁ r₂} and @{term r = SEQ r₁ r₂} cases. In each case we have to analyse four subcases whether @{term fst (simp r₁)} and @{term fst (simp r₂)} equals @{const ZERO} (respectively @{const ONE}). For example for @{term r = ALT r₁ r₂}, consider the subcase @{term fst (simp r₁) = ZERO} and @{term fst (simp r₂) ≠ ZERO}. By assumption we know @{term s

∈ fst (simp (ALT r₁ r₂)) → v}. From this we can infer @{term s ∈ fst (simp r₂) → v} and by IH also (∗) @{term s ∈ r₂ → (snd (simp r₂) v)}. Given @{term fst (simp r₁) = ZERO} we know @{term L (fst (simp r₁)) = {}}. By the first statement @{term L r₁} is the empty set, meaning (∗∗) @{term s ∉ L r₁}. Taking (∗) and (∗∗) together gives by the \mbox{‹P+R›}−rule @{term s ∈ ALT r₁ r₂ → Right (snd (simp r₂) v)}. In turn this gives @{term s ∈ ALT r₁ r₂ → snd (simp (ALT r₁ r₂)) v} as we need to show. The other cases are similar.\qed \end{proof} \noindent We can now prove relatively straightfor- wardly that the optimised lexer produces the expected result: \begin{theorem} @{thm slexer__correctness} \end{theorem} \begin{proof} By induction on @{term s} generalising over @{term r}. The case @{term []} is trivial. For the cons−case suppose the string is of the form @{term c # s}. By induction hypothesis we know @{term slexer r s = lexer r s} holds for all @{term r} (in particular for @{term r} being the derivative @{term der c r}). Let @{term rₛ} be the simplified derivative regular expression, that is @{term fst (simp (der c r))}, and @{term fᵣ} be the rectification function, that is @{term snd (simp (der c r))}. We distinguish the cases whether (∗) @{term s ∈ L (der c r)} or not. In the first case we have by Theorem~\ref{lexercorrect}(2) a value @{term v} so that @{term lexer (der c r) s = Some v} and @{term s ∈ der c r → v} hold. By Lemma~\ref{slexeraux}(1) we can also infer from~(∗) that @{term s ∈ L rₛ} holds. Hence we know by Theorem~\ref{lexercorrect}(2) that there exists a @{term v'} with @{term lexer rₛ s = Some v'} and @{term s ∈ rₛ → v'}. From the latter we know by Lemma~\ref{slexeraux}(2) that @{term s ∈ der c r → (fᵣ v')} holds. By the uniqueness of the POSIX relation (Theorem~\ref{posixdeterm}) we can infer that @{term v} is equal to @{term fᵣ v'}−−−that is the rectification function applied to @{term v'} produces the original @{term v}. Now the case follows by the definitions of @{const lexer} and @{const slexer}. In the second case where @{term s ∉ L (der c r)} we have that @{term lexer (der c r) s = None} by Theorem~\ref{lexercorrect}(1). We also know by Lemma~\ref{slexeraux}(1) that @{term s ∉ L rₛ}. Hence @{term lexer rₛ s = None} by Theorem~\ref{lexercorrect}(1) and by IH then also @{term slexer rₛ s = None}. With this we can conclude in this case too.\qed \end{proof} fy the result. This gives us a simplified derivative r_s and a rec- tification function f_r. The lexer is then recursively called with the simplified derivative, but before we inject the character c into the value v, we need to rec- tify v (that is construct $f_r\ v$). Before we can establish the correctness of $lexer^+$, we need to show that simplification preserves the language and simplification preserves our POSIX relation once the value is rectified (recall *simp* generates a (regular expression, rectification function) pair):

Lemma 14.
(1) $L(fst\ (simp\ r)) = L(r)$
(2) If $(s,\ fst\ (simp\ r)) → v$ then $(s,\ r) → snd\ (simp\ r)\ v$.

Proof. Both are by induction on r. There is no interesting case for the first statement. For the second statement, of interest are the $r = r_1 + r_2$ and $r = r_1 \cdot r_2$ cases. In each case we have to analyse four subcases whether *fst (simp*

r_1) and fst ($simp\ r_2$) equals $\mathbf{0}$ (respectively $\mathbf{1}$). For example for $r = r_1 + r_2$, consider the subcase fst ($simp\ r_1$) $= \mathbf{0}$ and fst ($simp\ r_2$) $\neq \mathbf{0}$. By assumption we know (s, fst ($simp\ (r_1 + r_2)$)) $\rightarrow v$. From this we can infer (s, fst ($simp\ r_2$)) $\rightarrow v$ and by IH also (*) (s, r_2) $\rightarrow snd$ ($simp\ r_2$) v. Given fst ($simp\ r_1$) $= \mathbf{0}$ we know $L(fst\ (simp\ r_1)) = \varnothing$. By the first statement $L(r_1)$ is the empty set, meaning (**) $s \notin L(r_1)$. Taking (*) and (**) together gives by the $P{+}R$-rule (s, $r_1 + r_2$) $\rightarrow Right$ (snd ($simp\ r_2$) v). In turn this gives (s, $r_1 + r_2$) $\rightarrow snd$ ($simp\ (r_1 + r_2)$) v as we need to show. The other cases are similar.   □

We can now prove relatively straightforwardly that the optimised lexer produces the expected result:

Theorem 5. $lexer^+\ r\ s = lexer\ r\ s$

Proof. By induction on s generalising over r. The case $[]$ is trivial. For the cons-case suppose the string is of the form $c{::}s$. By induction hypothesis we know $lexer^+\ r\ s = lexer\ r\ s$ holds for all r (in particular for r being the derivative $r\backslash c$). Let r_s be the simplified derivative regular expression, that is fst ($simp\ (r\backslash c)$), and f_r be the rectification function, that is snd ($simp\ (r\backslash c)$). We distinguish the cases whether (*) $s \in L(r\backslash c)$ or not. In the first case we have by Theorem 2(2) a value v so that $lexer\ (r\backslash c)\ s = Some\ v$ and (s, $r\backslash c$) $\rightarrow v$ hold. By Lemma 14(1) we can also infer from (*) that $s \in L(r_s)$ holds. Hence we know by Theorem 2(2) that there exists a v' with $lexer\ r_s\ s = Some\ v'$ and (s, r_s) $\rightarrow v'$. From the latter we know by Lemma 14(2) that (s, $r\backslash c$) $\rightarrow f_r\ v'$ holds. By the uniqueness of the POSIX relation (Theorem 1) we can infer that v is equal to $f_r\ v'$—that is the rectification function applied to v' produces the original v. Now the case follows by the definitions of $lexer$ and $lexer^+$.

    In the second case where $s \notin L(r\backslash c)$ we have that $lexer\ (r\backslash c)\ s = None$ by Theorem 2(1). We also know by Lemma 14(1) that $s \notin L(r_s)$. Hence $lexer\ r_s\ s = None$ by Theorem 2(1) and by IH then also $lexer^+\ r_s\ s = None$. With this we can conclude in this case too.   □

9   HERE

Lemma 15. *If* $v : (r^{\downarrow})\backslash c$ *then* $retrieve\ (r\backslash\!\backslash c)\ v = retrieve\ r\ (inj\ (r^{\downarrow})\ c\ v)$.

Proof. By induction on the definition of $r^{\downarrow}$. The cases for rule 1) and 2) are straightforward as $\mathbf{0}\backslash c$ and $\mathbf{1}\backslash c$ are both equal to $\mathbf{0}$. This means $v : \mathbf{0}$ cannot hold. Similarly in case of rule 3) where r is of the form $ACHAR\ d$ with $c = d$. Then by assumption we know $v : \mathbf{1}$, which implies $v = Empty$. The equation follows by simplification of left- and right-hand side. In case $c \neq d$ we have again $v : \mathbf{0}$, which cannot hold.

    For rule 4a) we have again $v : \mathbf{0}$. The property holds by IH for rule 4b). The induction hypothesis is

$$retrieve\ (r\backslash\!\backslash c)\ v = retrieve\ r\ (inj\ (r^{\downarrow})\ c\ v)$$

which is what left- and right-hand side simplify to. The slightly more interesting case is for 4c). By assumption we have $v : ((r_1{}^{\downarrow})\backslash c) + (((AALTs\ bs\ (r_2 :: rs))^{\downarrow})\backslash c)$. This means we have either (*) $v1 : (r_1{}^{\downarrow})\backslash c$ with $v = Left\ v1$ or (**) $v2 : ((AALTs\ bs\ (r_2 :: rs))^{\downarrow})\backslash c$ with $v = Right\ v2$. The former case is straightforward by simplification. The second case is ...TBD.

Rule 5) TBD.

Finally for rule 6) the reasoning is as follows: By assumption we have $v : ((r^{\downarrow})\backslash c) \cdot (r^{\downarrow})^{\star}$. This means we also have $v = Seq\ v1\ v2$, $v1 : (r^{\downarrow})\backslash c$ and $v2 = Stars\ vs$. We want to prove

$$retrieve\ (ASEQ\ bs\ (fuse\ [Z]\ (r\backslash\!\backslash c))\ (ASTAR\ []\ r))\ v \tag{3}$$

$$= retrieve\ (ASTAR\ bs\ r)\ (inj\ ((r^{\downarrow})^{\star})\ c\ v) \tag{4}$$

The right-hand side *inj*-expression is equal to $Stars\ (inj\ (r^{\downarrow})\ c\ v1 :: vs)$, which means the *retrieve*-expression simplifies to

$$bs\ @\ [Z]\ @\ retrieve\ r\ (inj\ (r^{\downarrow})\ c\ v1)\ @\ retrieve\ (ASTAR\ []\ r)\ (Stars\ vs)$$

The left-hand side (3) above simplifies to

$$bs\ @\ retrieve\ (fuse\ [Z]\ (r\backslash\!\backslash c))\ v1\ @\ retrieve\ (ASTAR\ []\ r)\ (Stars\ vs)$$

We can move out the *fuse* $[Z]$ and then use the IH to show that left-hand side and right-hand side are equal. This completes the proof.

References

1. The Open Group Base Specification Issue 6 IEEE Std 1003.1 2004 Edition, 2004. http://pubs.opengroup.org/onlinepubs/009695399/basedefs/xbd_chap09.html.
2. F. Ausaf, R. Dyckhoff, and C. Urban. POSIX Lexing with Derivatives of Regular Expressions (Proof Pearl). In *Proc. of the 7th International Conference on Interactive Theorem Proving (ITP)*, volume 9807 of *LNCS*, pages 69–86, 2016.
3. J. A. Brzozowski. Derivatives of Regular Expressions. *Journal of the ACM*, 11(4):481–494, 1964.
4. T. Coquand and V. Siles. A Decision Procedure for Regular Expression Equivalence in Type Theory. In *Proc. of the 1st International Conference on Certified Programs and Proofs (CPP)*, volume 7086 of *LNCS*, pages 119–134, 2011.
5. A. Frisch and L. Cardelli. Greedy Regular Expression Matching. In *Proc. of the 31st International Conference on Automata, Languages and Programming (ICALP)*, volume 3142 of *LNCS*, pages 618–629, 2004.
6. N. B. B. Grathwohl, F. Henglein, and U. T. Rasmussen. A Crash-Course in Regular Expression Parsing and Regular Expressions as Types. Technical report, University of Copenhagen, 2014.
7. H. Hosoya, J. Vouillon, and B. C. Pierce. Regular Expression Types for XML. *ACM Transactions on Programming Languages and Systems (TOPLAS)*, 27(1):46–90, 2005.
8. A. Krauss and T. Nipkow. Proof Pearl: Regular Expression Equivalence and Relation Algebra. *Journal of Automated Reasoning*, 49:95–106, 2012.

9. C. Kuklewicz. Regex Posix. https://wiki.haskell.org/Regex_Posix.
10. S. Okui and T. Suzuki. Disambiguation in Regular Expression Matching via Position Automata with Augmented Transitions. In *Proc. of the 15th International Conference on Implementation and Application of Automata (CIAA)*, volume 6482 of *LNCS*, pages 231–240, 2010.
11. S. Okui and T. Suzuki. Disambiguation in Regular Expression Matching via Position Automata with Augmented Transitions. Technical report, University of Aizu, 2013.
12. S. Owens and K. Slind. Adapting Functional Programs to Higher Order Logic. *Higher-Order and Symbolic Computation*, 21(4):377–409, 2008.
13. M. Sulzmann and K. Lu. POSIX Regular Expression Parsing with Derivatives. In *Proc. of the 12th International Conference on Functional and Logic Programming (FLOPS)*, volume 8475 of *LNCS*, pages 203–220, 2014.
14. S. Vansummeren. Type Inference for Unique Pattern Matching. *ACM Transactions on Programming Languages and Systems*, 28(3):389–428, 2006.