POSIX Lexing with Derivatives of Regular Expressions (Proof Pearl)

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. Sulzmann and Lu have made available on-line what they call a "rigorous proof" of the correctness of their algorithm w.r.t. their specification; regrettably, it appears to us to have unfillable gaps. In the first part of 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 also prove the correctness of an optimised version of the POSIX matching algorithm. Our definitions and proof are much simpler than those by Sulzmann and Lu and can be easily formalised in Isabelle/HOL. In the second part we analyse the correctness argument by Sulzmann and Lu and explain why it seems hard to turn it into a proof rigorous enough to be accepted by a system such as Isabelle/HOL.

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

1 Introduction

Brzozowski [1] 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

completely formalised correctness proof of this matcher in for example HOL4 has been mentioned in [5]. Another one in Isabelle/HOL is part of the work in [3].

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 [6] extended this matcher to allow generation not just of a YES/NO answer but of an actual matching, called a [lexical] *value*. They give a simple algorithm to calculate a value that appears to be the value associated with POSIX matching [4,7]. 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 [6] 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 [2] on a GREEDY regular expression matching algorithm. Beginning with our observations that, without evidence that it is transitive, it cannot be called an "order relation", and that the relation is called a "total order" despite being evidently not total[4], we identify problems with this approach (of which some of the proofs are not published in [6]); perhaps more importantly, we give 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 computes such a value and that such a value is unique. 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 Kuklewitz [4] who found that nearly all POSIX matching implementations are "buggy" [6, Page 203].

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 [2] and the other is POSIX matching [4,6,7]. 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

---

[4] The relation $\geq_r$ defined in [6] is a relation on the values for the regular expression r; but it only holds between v and v' in cases where v and v' have the same flattening (underlying string). So a counterexample to totality is given by taking two values v and v' for r that have different flattenings (see Section 3). A different relation $\geq_{r,s}$ on the set of values for r with flattening s is definable by the same approach, and is indeed total; but that is not what Proposition 1 of [6] does.

identifiers, respectively. There are two underlying (informal) rules behind tokenising a string in a POSIX 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 regular expression that can match determines the token.

Consider for example r_{key} 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.

Contributions: (NOT DONE YET) We have implemented in Isabelle/HOL the derivative-based regular expression matching algorithm as described by Sulzmann and Lu [6]. We have proved the correctness of this algorithm according to our specification of what a POSIX value is. Sulzmann and Lu sketch in [6] an informal correctness proof: but to us it contains unfillable gaps.

informal correctness proof given in [6] is in final form[5] and to us contains unfillable gaps.

Our specification of a POSIX value consists of a simple inductive definition that given a string and a regular expression uniquely determines this value. Derivatives as calculated by Brzozowski's method are usually more complex regular expressions than the initial one; various optimisations are possible, such as the simplifications of $0 + r$, $r + 0$, $1 \cdot r$ and $r \cdot 1$ to r. One of the advantages of having a simple specification and correctness proof is that the latter can be refined to allow for such optimisations and simple correctness proof.

An extended version of [6] is available at the website of its first author; this includes some "proofs", claimed in [6] to be "rigorous". Since these are evidently not in final form, we make no comment thereon, preferring to give general reasons for our belief that the approach of [6] is problematic rather than to discuss details of unpublished work.

2   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 $[c]$. 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:

5

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

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 clauses:

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

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:

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

For semantic derivatives we have the following equations (for example mechanically proved in [3]):

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

Brzozowski's derivatives of regular expressions [1] 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:

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

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

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

$$r\backslash [] \quad \overset{\text{def}}{=} \ r$$
$$r\backslash (c :: s) \overset{\text{def}}{=} \ (r\backslash c)\backslash s$$

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 [6] is to append another phase to this algorithm in order to calculate a [lexical] value. We will explain the details next.

3   POSIX Regular Expression Matching

The clever idea in [6] is to introduce values for encoding *how* a regular expression matches a string and then define a function on values that mirrors (but inverts) the construction of the derivative on regular expressions. *Values* are defined as the inductive datatype

$$v := () \mid Char\ c \mid Left\ v \mid Right\ v \mid Seq\ v_1\ v_2 \mid Stars\ vs$$

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 [2], and Sulzmann and Lu [6] for POSIX matching). The string underlying a value can be calculated by the *flat* function, written $|\_|$ and defined as:

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

Sulzmann and Lu also define inductively an inhabitation relation that associates values to regular expressions:

$$\overline{() : \mathbf{1}} \qquad \overline{Char\ c : c}$$

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

$$\frac{v_1 : r_1 \qquad v_2 : r_2}{Seq\ v_1\ v_2 : r_1 \cdot r_2}$$

$$\overline{Stars\ [] : r^\star} \qquad \frac{v : r \qquad Stars\ vs : r^\star}{Stars\ (v :: vs) : r^\star}$$

Note that no values are associated with the regular expression **0**, and that the only value associated with the regular expression **1** is (), pronounced (if one must) as *Void*. It is routine to establish how values "inhabiting" a regular expression correspond to the language of a regular expression, namely

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

In general there is more than one value associated with a regular expression. In case of POSIX matching the problem is to calculate the unique 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 *der/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

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

$$\Big\downarrow \textit{mkeps}$$

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

Fig. 1. The two phases of the algorithm by Sulzmann & Lu [6], matching the string $[a, b, c]$. The first phase (the arrows from left to right) is Brzozowski's matcher building succesive 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 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.

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 rules in case there are several ways). This functions is defined by the clauses:

$$\textit{mkeps} \, (\mathbf{1}) \qquad\quad \overset{\text{def}}{=} \; ()$$
$$\textit{mkeps} \, (r_1 \cdot r_2) \quad \overset{\text{def}}{=} \; Seq \, (\textit{mkeps } r_1) \, (\textit{mkeps } r_2)$$
$$\textit{mkeps} \, (r_1 + r_2) \overset{\text{def}}{=} \; \textit{if nullable } r_1 \textit{ then Left } (\textit{mkeps } r_1) \textit{ else Right } (\textit{mkeps } r_2)$$
$$\textit{mkeps} \, (r^\star) \qquad\; \overset{\text{def}}{=} \; Stars \, []$$

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 an error is raised instead. 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.

The most interesting idea from Sulzmann and Lu [6] 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 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 corresponding to the input regular expression. The *inj* function is by recursion on the regular expressions and by analysing the shape of values (corresponding to the derivative regular expressions).

$$
\begin{array}{lll}
(1) & inj\ d\ c\ () & \stackrel{def}{=}\ Char\ d \\
(2) & inj\ (r_1 + r_2)\ c\ (Left\ v_1) & \stackrel{def}{=}\ Left\ (inj\ r_1\ c\ v_1) \\
(3) & inj\ (r_1 + r_2)\ c\ (Right\ v_2) & \stackrel{def}{=}\ Right\ (inj\ r_2\ c\ v_2) \\
(4) & inj\ (r_1 \cdot r_2)\ c\ (Seq\ v_1\ v_2) & \stackrel{def}{=}\ Seq\ (inj\ r_1\ c\ v_1)\ v_2 \\
(5) & inj\ (r_1 \cdot r_2)\ c\ (Left\ (Seq\ v_1\ v_2)) & \stackrel{def}{=}\ Seq\ (inj\ r_1\ c\ v_1)\ v_2 \\
(6) & inj\ (r_1 \cdot r_2)\ c\ (Right\ v_2) & \stackrel{def}{=}\ Seq\ (mkeps\ r_1)\ (inj\ r_2\ c\ v_2) \\
(7) & inj\ (r^\star)\ c\ (Seq\ v\ (Stars\ vs)) & \stackrel{def}{=}\ Stars\ (inj\ r\ c\ v :: vs)
\end{array}
$$

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 *der*-function for sequence regular expressions:

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

Consider first the else-branch where the derivative is $(r_1\backslash c) \cdot r_2$. The corresponding value must therefore be 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 $r \cdot r^\star$. 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 is to build this constraint explicitly into our function definition.[6]

The idea of *inj* to "inject back" a character into a value can be made precise by the first part of the following lemma; the second part shows that the underlying string of an *mkeps*-value is the empty string.

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

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

---

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

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

If the regular expression does not match, *None* is returned, indicating an error is raised. 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 a 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 the longest match and priority rule; as correctly argued in [6], this needs formal specification. Sulzmann and Lu define a *dominance* relation[7] between values and argue that there is a maximum value, as given by the derivative-based algorithm. In contrast, we shall next introduce a simple inductive definition to specify 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 in [7]. The relation we define is ternary and written as $(s, r) \rightarrow v$, relating strings, regular expressions and values.

$$\overline{([], \mathbf{1}) \rightarrow ()} \qquad \overline{([c], c) \rightarrow Char\ c}$$

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

$$\frac{(s_1, r_1) \rightarrow v_1 \qquad (s_2, r_2) \rightarrow v_2}{\nexists s_3\ s_4.\ s_3 \neq [] \wedge s_3\ @\ s_4 = s_2 \wedge s_1\ @\ s_3 \in L(r_1) \wedge s_4 \in L(r_2)}{(s_1\ @\ s_2, r_1 \cdot r_2) \rightarrow Seq\ v_1\ v_2}$$

$$\overline{([], r^\star) \rightarrow Stars\ []}$$

$$\frac{(s_1, r) \rightarrow v \qquad (s_2, r^\star) \rightarrow Stars\ vs \qquad |v| \neq []}{\nexists s_3\ s_4.\ s_3 \neq [] \wedge s_3\ @\ s_4 = s_2 \wedge s_1\ @\ s_3 \in L(r) \wedge s_4 \in L(r^\star)}{(s_1\ @\ s_2, r^\star) \rightarrow Stars\ (v::vs)}$$

We claim that this relation captures the idea behind the two informal POSIX rules shown in the Introduction: Consider the second line where the POSIX values for an alternative regular expression 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 the rule on right). Interesting is also the rule for sequence regular expressions (third line). 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.

Theorem 1. *If* $(s, r) \rightarrow v_1$ *and* $(s, r) \rightarrow v_2$ *then* $v_1 = v_2$.

---

[7] Sulzmann and Lu call it an ordering relation, but without giving evidence that it is transitive.

Lemma 2. *If nullable r then* $([], r) \to mkeps\ r.$

Lemma 3. *If* $(s, r\backslash c) \to v$ *then* $(c :: s, r) \to inj\ r\ c\ v.$

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

4   The Argument by Sulzmmann and Lu

5   Conclusion

Nipkow lexer from 2000

We have also introduced a slightly restricted version of this relation where the last rule is restricted so that $|v| \neq []$. This relation for *non-problematic* is written $\models v : r.$

Our version of Sulzmann's ordering relation

$$\frac{v_1 \succ_{r_1} v_1{}' \quad v_1 \neq v_1{}'}{(Seq\ v_1\ v_2) \succ (r_1 \cdot r_2) (Seq\ v_1{}'\ v_2{}')} \qquad \frac{v_2 \succ_{r_2} v_2{}'}{(Seq\ v_1\ v_2) \succ (r_1 \cdot r_2) (Seq\ v_1\ v_2{}')}$$

$$\frac{len\ (|v_1|) \leq len\ (|v_2|)}{(Left\ v_2) \succ (r_1 + r_2) (Right\ v_1)} \qquad \frac{len\ (|v_2|) < len\ (|v_1|)}{(Right\ v_1) \succ (r_1 + r_2) (Left\ v_2)}$$

$$\frac{v_2 \succ_{r_2} v_2{}'}{(Right\ v_2) \succ (r_1 + r_2) (Right\ v_2{}')} \qquad \frac{v_1 \succ_{r_1} v_1{}'}{(Left\ v_1) \succ (r_1 + r_2) (Left\ v_1{}')}$$

$$\frac{}{(()) \succ (\mathbf{1}) (())} \qquad \frac{}{(Char\ c) \succ (c) (Char\ c)}$$

$$\frac{|Stars\ (v :: vs)| = []}{(Stars\ []) \succ (r^\star) (Stars\ (v :: vs))} \qquad \frac{|Stars\ (v :: vs)| \neq []}{(Stars\ (v :: vs)) \succ (r^\star) (Stars\ [])}$$

$$\frac{v_1 \succ_r v_2 \quad v_1 \neq v_2}{(Stars\ (v_1 :: vs_1)) \succ (r^\star) (Stars\ (v_2 :: vs_2))}$$

$$\frac{(Stars\ vs_1) \succ (r^\star) (Stars\ vs_2)}{(Stars\ (v :: vs_1)) \succ (r^\star) (Stars\ (v :: vs_2))} \qquad \frac{}{(Stars\ []) \succ (r^\star) (Stars\ [])}$$

A prefix of a string s

$$s_1 \sqsubseteq s_2 \stackrel{def}{=} \exists s3.\ s_1 \mathbin{@} s3 = s_2$$

Values and non-problematic values

$$Values\ r\ s \stackrel{def}{=} \{v \mid v : r \wedge (|v|) \sqsubseteq s\}$$

The point is that for a given s and r there are only finitely many non-problematic values.

Some lemmas we have proved:

$L(r) = \{|v| \mid v : r\}$
$L(r) = \{|v| \mid \models v : r\}$
If nullable r then mkeps $r : r$.
If nullable r then $|mkeps\ r| = [\,]$.
If $v : r \backslash c$ then inj $r\ c\ v : r$.
If $v : r \backslash c$ then $|inj\ r\ c\ v| = c :: |v|$.
If nullable r then $([\,], r) \rightarrow mkeps\ r$.
If $(s, r) \rightarrow v$ then $|v| = s$.
If $(s, r) \rightarrow v$ then $\models v : r$.
If $(s, r) \rightarrow v_1$ and $(s, r) \rightarrow v_2$ then $v_1 = v_2$.

This is the main theorem that lets us prove that the algorithm is correct according to $(s, r) \rightarrow v$:

If $(s, r \backslash c) \rightarrow v$ then $(c :: s, r) \rightarrow inj\ r\ c\ v$.

Proof The proof is by induction on the definition of *der*. Other inductions would go through as well. The interesting case is for $r_1 \cdot r_2$. First we analyse the case where *nullable r_1*. We have by induction hypothesis

$$(IH1) \quad \forall s\ v.\ \text{if}\ (s, r_1 \backslash c) \rightarrow v\ \text{then}\ (c :: s, r_1) \rightarrow inj\ r_1\ c\ v$$
$$(IH2) \quad \forall s\ v.\ \text{if}\ (s, r_2 \backslash c) \rightarrow v\ \text{then}\ (c :: s, r_2) \rightarrow inj\ r_2\ c\ v$$

and have

$$(s, (r_1 \backslash c) \cdot r_2 + (r_2 \backslash c)) \rightarrow v$$

There are two cases what v can be: (1) *Left v'* and (2) *Right v'*.

(1) We know $(s, (r_1 \backslash c) \cdot r_2) \rightarrow v'$ holds, from which we can infer that there are s_1, s_2, v_1, v_2 with

$$(s_1, r_1 \backslash c) \rightarrow v_1 \qquad \text{and} \qquad (s_2, r_2) \rightarrow v_2$$

and also

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

and have to prove

$$(c :: s_1 @ s_2, r_1 \cdot r_2) \rightarrow Seq\ (inj\ r_1\ c\ v_1)\ v_2$$

The two requirements $(c :: s_1, r_1) \rightarrow inj\ r_1\ c\ v_1$ and $(s_2, r_2) \rightarrow v_2$ can be proved by the induction hypothese (IH1) and the fact above.
This leaves to prove

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

which holds because $c :: s_1 @ s_3 \in L(r_1)$ implies $s_1 @ s_3 \in L(r_1 \backslash c)$

(2) This case is similar.

The final case is that $\neg$ *nullable* r_1 holds. This case again similar to the cases above.

References

1. J. A. Brzozowski. Derivatives of Regular Expressions. *Journal of the ACM*, 11(4):481–494, 1964.
2. 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.
3. A. Krauss and T. Nipkow. Proof Pearl: Regular Expression Equivalence and Relation Algebra. *Journal of Automated Reasoning*, 49:95–106, 2012.
4. C. Kuklewicz. Regex Posix. https://wiki.haskell.org/Regex_Posix.
5. S. Owens and K. Slind. Adapting Functional Programs to Higher Order Logic. *Higher-Order and Symbolic Computation*, 21(4):377–409, 2008.
6. 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.
7. S. Vansummeren. Type Inference for Unique Pattern Matching. *ACM Transactions on Programming Languages and Systems*, 28(3):389–428, 2006.

6   Roy's Rules

$$Void \lhd \epsilon \qquad Char\ c \lhd Lit\ c$$

$$\frac{v_1 \lhd r_1}{Left\ v_1 \lhd r_1 + r_2} \qquad \frac{v_2 \lhd r_2 \qquad |v_2| \notin L(r_1)}{Right\ v_2 \lhd r_1 + r_2}$$

$$\frac{v_1 \lhd r_1 \qquad v_2 \lhd r_2 \qquad s \in L(r_1 \backslash |v_1|) \wedge |v_2| \backslash s\ \epsilon\ L(r_2) \Rightarrow s = []}{(v_1, v_2) \lhd r_1 \cdot r_2}$$

$$\frac{v \lhd r \qquad vs \lhd r^* \qquad |v| \neq []}{(v :: vs) \lhd r^*} \qquad [] \lhd r^*$$