

POSIX Lexing with Bitcoded Derivatives

Chengsong Tan ✉

King's College London

Christian Urban ✉

King's College London

Abstract

Sulzmann and Lu describe a lexing algorithm that calculates Brzozowski derivatives using bitcodes annotated to regular expressions. 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. This information is needed in the context of lexing in order to extract and to classify tokens. The purpose of the bitcodes is to generate POSIX values incrementally while derivatives are calculated. They also help with designing an “aggressive” simplification function that keeps the size of derivatives finite. Without simplification the size of some derivatives can grow arbitrarily big resulting in an extremely slow lexing algorithm. In this paper we describe a variant of Sulzmann and Lu’s algorithm: Our variant is a recursive functional program, whereas Sulzmann and Lu’s version involves a fixpoint construction. We (i) prove in Isabelle/HOL that our algorithm is correct and generates unique POSIX values; we also (ii) establish a finite bound for the size of the derivatives.

2012 ACM Subject Classification Design and analysis of algorithms; Formal languages and automata theory

Keywords and phrases POSIX matching and lexing, derivatives of regular expressions, Isabelle/HOL

Digital Object Identifier 10.4230/LIPIcs...

1 Introduction

In the last fifteen or so years, Brzozowski’s derivatives of regular expressions have sparked quite a bit of interest in the functional programming and theorem prover communities. The beauty of Brzozowski’s derivatives [4] 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. Derivatives of a regular expression, written $r \setminus c$, give 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*). We are aware of a mechanised correctness proof of Brzozowski’s derivative-based matcher in HOL4 by Owens and Slind [10]. Another one in Isabelle/HOL is part of the work by Krauss and Nipkow [7]. And another one in Coq is given by Coquand and Siles [5]. Also Ribeiro and Du Bois give one in Agda [11].

However, there are two difficulties with derivative-based matchers: First, Brzozowski’s original matcher only generates a yes/no answer for whether a regular expression matches a string or not. This is too little information in the context of lexing where separate tokens must be identified and also classified (for example as keywords or identifiers). Sulzmann and Lu [12] overcome this difficulty by cleverly extending Brzozowski’s matching algorithm. Their extended version generates additional information on *how* a regular expression matches a string following the POSIX rules for regular expression matching. They achieve this by adding a second “phase” to Brzozowski’s algorithm involving an injection function. In our own earlier work we provided the formal specification of what POSIX matching means and proved in Isabelle/HOL the correctness of Sulzmann and Lu’s extended algorithm accordingly [3].

The second difficulty is that Brzozowski’s derivatives can grow to arbitrarily big sizes. For example if we start with the regular expression $(a + aa)^*$ and take successive derivatives according to the character a , we end up with a sequence of ever-growing derivatives like



© :

licensed under Creative Commons License CC-BY 4.0

Leibniz International Proceedings in Informatics



LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

$$\begin{aligned}
(a + aa)^* &\xrightarrow{-\lambda^a} (\mathbf{1} + \mathbf{1}a) \cdot (a + aa)^* \\
&\xrightarrow{-\lambda^a} (\mathbf{0} + \mathbf{0}a + \mathbf{1}) \cdot (a + aa)^* + (\mathbf{1} + \mathbf{1}a) \cdot (a + aa)^* \\
&\xrightarrow{-\lambda^a} (\mathbf{0} + \mathbf{0}a + \mathbf{0}) \cdot (a + aa)^* + (\mathbf{1} + \mathbf{1}a) \cdot (a + aa)^* + \\
&\quad (\mathbf{0} + \mathbf{0}a + \mathbf{1}) \cdot (a + aa)^* + (\mathbf{1} + \mathbf{1}a) \cdot (a + aa)^* \\
&\xrightarrow{-\lambda^a} \dots \quad (\text{regular expressions of sizes } 98, 169, 283, 468, 767, \dots)
\end{aligned}$$

where after around 35 steps we run out of memory on a typical computer (we shall define shortly the precise details of our regular expressions and the derivative operation). Clearly, the notation involving **0**s and **1**s already suggests simplification rules that can be applied to regular regular expressions, for example $\mathbf{0}r \Rightarrow \mathbf{0}$, $\mathbf{1}r \Rightarrow r$, $\mathbf{0} + r \Rightarrow r$ and $r + r \Rightarrow r$. While such simple-minded simplifications have been proved in our earlier work to preserve the correctness of Sulzmann and Lu’s algorithm [3], they unfortunately do *not* help with limiting the growth of the derivatives shown above: the growth is slowed, but the derivatives can still grow rather quickly beyond any finite bound.

Sulzmann and Lu overcome this “growth problem” in a second algorithm [12] where they introduce bitcoded regular expressions. In this version, POSIX values are represented as bitsequences and such sequences are incrementally generated when derivatives are calculated. The compact representation of bitsequences and regular expressions allows them to define a more “aggressive” simplification method that keeps the size of the derivatives finite no matter what the length of the string is. They make some informal claims about the correctness and linear behaviour of this version, but do not provide any supporting proof arguments, not even “pencil-and-paper” arguments. They write about their bitcoded *incremental parsing method* (that is the algorithm to be formalised in this paper):

“Correctness Claim: We further claim that the incremental parsing method [...] in combination with the simplification steps [...] yields POSIX parse trees. We have tested this claim extensively [...] but yet have to work out all proof details.” [12, Page 14]

Contributions: We have implemented in Isabelle/HOL the derivative-based lexing algorithm of Sulzmann and Lu [12] where regular expressions are annotated with bitsequences. We define the crucial simplification function as a recursive function, without the need of a fix-point operation. One objective of the simplification function is to remove duplicates of regular expressions. For this Sulzmann and Lu use in their paper the standard *nub* function from Haskell’s list library, but this function does not achieve the intended objective with bitcoded regular expressions. The reason is that in the bitcoded setting, each copy generally has a different bitcode annotation—so *nub* would never “fire”. Inspired by Scala’s library for lists, we shall instead use a *distinctBy* function that finds duplicates under an erasing function which deletes bitcodes. We shall also introduce our own argument and definitions for establishing the correctness of the bitcoded algorithm when simplifications are included.

In this paper, we shall first briefly introduce the basic notions of regular expressions and describe the basic definitions of POSIX lexing from our earlier work [3]. This serves as a reference point for what correctness means in our Isabelle/HOL proofs. We shall then prove the correctness for the bitcoded algorithm without simplification, and after that extend the proof to include simplification.

2 Background

In our Isabelle/HOL formalisation strings are lists of characters with the empty string being represented by the empty list, written `[]`, and list-cons being written as `_::__`; string concatenation is `_@_`. We often use the usual bracket notation for lists also for strings; for example a string consisting of just a

single character c is written $[c]$. Our regular expressions are defined as usual as the elements of the following inductive datatype:

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

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 constructors $+$ and \cdot represent alternatives and sequences, respectively. We sometimes omit the \cdot in a sequence regular expression for brevity. The *language* of a regular expression, written L , is defined as usual and we omit giving the definition here (see for example [3]).

Central to Brzowski's regular expression matcher are two functions called *nullable* and *derivative*. The latter is written $r \setminus c$ for the derivative of the regular expression r w.r.t. the character c . Both functions are defined by recursion over regular expressions.

$$\begin{array}{ll} \mathbf{0} \setminus c & \stackrel{\text{def}}{=} \mathbf{0} & \text{nullable}(\mathbf{0}) & \stackrel{\text{def}}{=} \text{False} \\ \mathbf{1} \setminus c & \stackrel{\text{def}}{=} \mathbf{0} & \text{nullable}(\mathbf{1}) & \stackrel{\text{def}}{=} \text{True} \\ d \setminus c & \stackrel{\text{def}}{=} \text{if } c = d \text{ then } \mathbf{1} \text{ else } \mathbf{0} & \text{nullable}(c) & \stackrel{\text{def}}{=} \text{False} \\ (r_1 + r_2) \setminus c & \stackrel{\text{def}}{=} (r_1 \setminus c) + (r_2 \setminus c) & \text{nullable}(r_1 + r_2) & \stackrel{\text{def}}{=} \text{nullable } r_1 \vee \text{nullable } r_2 \\ (r_1 \cdot r_2) \setminus c & \stackrel{\text{def}}{=} \text{if nullable } r_1 & \text{nullable}(r_1 \cdot r_2) & \stackrel{\text{def}}{=} \text{nullable } r_1 \wedge \text{nullable } r_2 \\ & \text{then } (r_1 \setminus c) \cdot r_2 + (r_2 \setminus c) & \text{nullable}(r^*) & \stackrel{\text{def}}{=} \text{True} \\ & \text{else } (r_1 \setminus c) \cdot r_2 & & \\ (r^*) \setminus c & \stackrel{\text{def}}{=} (r \setminus c) \cdot r^* & & \end{array}$$

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

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

Using *nullable* and the derivative operation, we can define the following simple regular expression matcher:

$$\text{match } s \ r \stackrel{\text{def}}{=} \text{nullable}(r \setminus s)$$

This is essentially Brzowski's algorithm from 1964. Its main virtue is that the algorithm can be easily implemented as a functional program (either in a functional programming language or in a theorem prover). The correctness proof for *match* amounts to establishing the property

► **Proposition 1.** *match* $s \ r$ if and only if $s \in L(r)$

It is a fun exercise to formally prove this property in a theorem prover.

The novel idea of Sulzmann and Lu is to extend this algorithm for lexing, where it is important to find out which part of the string is matched by which part of the regular expression. For this Sulzmann and Lu presented two lexing algorithms in their paper [12]. The first algorithm consists of two phases: first a matching phase (which is Brzowski's algorithm) and then a value construction phase. The values encode *how* a regular expression matches a string. *Values* are defined as the inductive datatype

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

where we use vs to stand for a list of values. The string underlying a value can be calculated by a *flat* function, written $|_|$. It traverses a value and collects the characters contained in it. Sulzmann and Lu also define inductively an inhabitation relation that associates values to regular expressions:

$$\begin{array}{c}
 \frac{}{(\mathbf{0}, \mathbf{1}) \rightarrow \text{Empty}} P\mathbf{1} \qquad \frac{}{([c], c) \rightarrow \text{Char } c} Pc \\
 \frac{(s, r_1) \rightarrow v}{(s, r_1 + r_2) \rightarrow \text{Left } v} P+L \qquad \frac{(s, r_2) \rightarrow v \quad s \notin L r_1}{(s, r_1 + r_2) \rightarrow \text{Right } v} P+R \\
 \frac{(s_1, r_1) \rightarrow v_1 \quad (s_2, r_2) \rightarrow v_2 \quad \nexists s_3 s_4. s_3 \neq \mathbf{0} \wedge s_3 @ s_4 = s_2 \wedge s_1 @ s_3 \in L r_1 \wedge s_4 \in L r_2}{(s_1 @ s_2, r_1 \cdot r_2) \rightarrow \text{Seq } v_1 v_2} PS \\
 \frac{}{(\mathbf{0}, r^*) \rightarrow \text{Stars } \mathbf{0}} P\mathbf{0} \qquad \frac{(s_1, r) \rightarrow v \quad (s_2, r^*) \rightarrow \text{Stars } vs \quad |v| \neq \mathbf{0} \quad \nexists s_3 s_4. s_3 \neq \mathbf{0} \wedge s_3 @ s_4 = s_2 \wedge s_1 @ s_3 \in L r \wedge s_4 \in L(r^*)}{(s_1 @ s_2, r^*) \rightarrow \text{Stars } (v :: vs)} P\star
 \end{array}$$

■ **Figure 1** The inductive definition of POSIX values taken from our earlier paper [3]. The ternary relation, written $(s, r) \rightarrow v$, formalises the notion of given a string s and a regular expression r what is the unique value v that satisfies the informal POSIX constraints for regular expression matching.

$$\begin{array}{c}
 \frac{}{\vdash \text{Empty} : \mathbf{1}} \qquad \frac{}{\vdash \text{Char } c : c} \\
 \frac{\vdash v_1 : r_1}{\vdash \text{Left } v_1 : r_1 + r_2} \qquad \frac{\vdash v_2 : r_2}{\vdash \text{Right } v_2 : r_1 + r_2} \\
 \frac{\vdash v_1 : r_1 \quad \vdash v_2 : r_2}{\vdash \text{Seq } v_1 v_2 : r_1 \cdot r_2} \qquad \frac{\forall v \in vs. \vdash v : r \wedge |v| \neq \mathbf{0}}{\vdash \text{Stars } vs : r^*}
 \end{array}$$


Note that no values are associated with the regular expression $\mathbf{0}$, since it cannot match any string. It is routine to establish how values “inhabiting” a regular expression correspond to the language of a regular expression, namely

► **Proposition 2.** $L r = \{|v| \mid \vdash v : r\}$

In general there is more than one value inhabited by a regular expression (meaning regular expressions can typically match more than one string). But even when fixing a string from the language of the regular expression, there are generally more than one way of how the regular expression can match this string. POSIX lexing is about identifying the unique value for a given regular expression and a string that satisfies the informal POSIX rules (see [1, 8, 9, 12, 13]).¹ Sometimes these informal rules are called *maximal much rule* and *rule priority*. One contribution of our earlier paper is to give a convenient specification for what POSIX values are (the inductive rules are shown in Figure 1).

The clever idea by Sulzmann and Lu [12] in their first algorithm is to define an injection function on values that mirrors (but inverts) the construction of the derivative on regular expressions. Essentially it injects back a character into a value. For this they define two functions called *mkeps* and *inj*:

¹ POSIX lexing acquired its name from the fact that the corresponding rules were described as part of the POSIX specification for Unix-like operating systems [1].



■ **Figure 2** The two phases of the first algorithm by Sulzmann & Lu [12], 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 value v_1 is the result of the algorithm representing the POSIX value for this string and regular expression.

$$\begin{aligned}
 \mathit{mkeps} \mathbf{1} & \stackrel{\text{def}}{=} \mathit{Empty} \\
 \mathit{mkeps} (r_1 \cdot r_2) & \stackrel{\text{def}}{=} \mathit{Seq} (\mathit{mkeps} r_1) (\mathit{mkeps} r_2) \\
 \mathit{mkeps} (r_1 + r_2) & \stackrel{\text{def}}{=} \mathit{if nullable } r_1 \mathit{ then Left } (\mathit{mkeps} r_1) \mathit{ else Right } (\mathit{mkeps} r_2) \\
 \mathit{mkeps} (r^*) & \stackrel{\text{def}}{=} \mathit{Stars} [] \\
 \mathit{inj} d c (\mathit{Empty}) & \stackrel{\text{def}}{=} \mathit{Char} d \\
 \mathit{inj} (r_1 + r_2) c (\mathit{Left} v_1) & \stackrel{\text{def}}{=} \mathit{Left} (\mathit{inj} r_1 c v_1) \\
 \mathit{inj} (r_1 + r_2) c (\mathit{Right} v_2) & \stackrel{\text{def}}{=} \mathit{Right} (\mathit{inj} r_2 c v_2) \\
 \mathit{inj} (r_1 \cdot r_2) c (\mathit{Seq} v_1 v_2) & \stackrel{\text{def}}{=} \mathit{Seq} (\mathit{inj} r_1 c v_1) v_2 \\
 \mathit{inj} (r_1 \cdot r_2) c (\mathit{Left} (\mathit{Seq} v_1 v_2)) & \stackrel{\text{def}}{=} \mathit{Seq} (\mathit{inj} r_1 c v_1) v_2 \\
 \mathit{inj} (r_1 \cdot r_2) c (\mathit{Right} v_2) & \stackrel{\text{def}}{=} \mathit{Seq} (\mathit{mkeps} r_1) (\mathit{inj} r_2 c v_2) \\
 \mathit{inj} (r^*) c (\mathit{Seq} v (\mathit{Stars} vs)) & \stackrel{\text{def}}{=} \mathit{Stars} (\mathit{inj} r c v :: vs)
 \end{aligned}$$

The function *mkeps* is run when the last derivative is nullable, that is the string to be matched is in the language of the regular expression. It generates a value for how the last derivative can match the empty string. The injection function then calculates the corresponding value for each intermediate derivative until a value for the original regular expression is generated. Graphically the algorithm by Sulzmann and Lu can be illustrated by the picture in Figure 2 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. The picture above shows the steps required when a regular expression, say r_1 , matches the string $[a, b, c]$. The first lexing algorithm by Sulzmann and Lu can be defined as:

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

We have shown in our earlier paper [3] that this algorithm is correct, that is it generates POSIX values. The central property we established relates the derivative operation to the injection function.

► **Proposition 3.** *If $(s, r \setminus c) \rightarrow v$ then $(c :: s, r) \rightarrow \mathit{inj} r c v$.*

With this in place we were able to prove:

► **Proposition 4.**

XX:6 POSIX Lexing with Bitcoded Derivatives

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

In fact we have shown that in the success case the generated POSIX value v is unique and in the failure case that there is no POSIX value v that satisfies $(s, r) \rightarrow v$. While the algorithm is correct, it is excruciatingly slow in cases where the derivatives grow arbitrarily (recall the example from the Introduction). However it can be used as a convenient reference point for the correctness proof of the second algorithm by Sulzmann and Lu, which we shall describe next.

3 Bitcoded Regular Expressions and Derivatives

In the second part of their paper [12], Sulzmann and Lu describe another algorithm that also generates POSIX values but dispenses with the second phase where characters are injected “back” into values. For this they annotate bitcodes to regular expressions, which we define in Isabelle/HOL as the datatype

$$\begin{aligned} breg & ::= ZERO \mid ONE\ bs \\ & \mid CHAR\ bs\ c \\ & \mid ALTs\ bs\ rs \\ & \mid SEQ\ bs\ r_1\ r_2 \\ & \mid STAR\ bs\ r \end{aligned}$$

where bs stands for bitsequences; r , r_1 and r_2 for bitcoded regular expressions; and rs for lists of bitcoded regular expressions. The binary alternative $ALT\ bs\ r_1\ r_2$ is just an abbreviation for $ALTs\ bs\ [r_1, r_2]$. For bitsequences we use lists made up of the constants Z and S . The idea with bitcoded regular expressions is to incrementally generate the value information (for example *Left* and *Right*) as bitsequences. For this Sulzmann and Lu define a coding function for how values can be coded into bitsequences.

$$\begin{array}{ll} code\ (Empty) & \stackrel{\text{def}}{=} [] \\ code\ (Char\ c) & \stackrel{\text{def}}{=} [] \\ code\ (Left\ v) & \stackrel{\text{def}}{=} Z :: code\ v \\ code\ (Right\ v) & \stackrel{\text{def}}{=} S :: code\ v \end{array} \qquad \begin{array}{ll} code\ (Seq\ v_1\ v_2) & \stackrel{\text{def}}{=} code\ v_1 @ code\ v_2 \\ code\ (Stars\ []) & \stackrel{\text{def}}{=} [S] \\ code\ (Stars\ (v :: vs)) & \stackrel{\text{def}}{=} Z :: code\ v @ code\ (Stars\ vs) \end{array}$$

As can be seen, this coding is “lossy” in the sense that we do not record explicitly character values and also not sequence values (for them we just append two bitsequences). However, the different alternatives for *Left*, respectively *Right*, are recorded as Z and S followed by some bitsequence. Similarly, we use Z to indicate if there is still a value coming in the list of *Stars*, whereas S indicates the end of the list. The lossiness makes the process of decoding a bit more involved, but the point is that if we have a regular expression *and* a bitsequence of a corresponding value, then we can always decode the value accurately. The decoding can be defined by using two functions called *decode'* and *decode*:

$$\begin{aligned}
\text{decode}' \text{ bs } (\mathbf{1}) &\stackrel{\text{def}}{=} (\text{Empty}, \text{bs}) \\
\text{decode}' \text{ bs } (c) &\stackrel{\text{def}}{=} (\text{Char } c, \text{bs}) \\
\text{decode}' (Z::\text{bs}) (r_1 + r_2) &\stackrel{\text{def}}{=} \text{let } (v, \text{bs}_1) = \text{decode}' \text{ bs } r_1 \text{ in } (\text{Left } v, \text{bs}_1) \\
\text{decode}' (S::\text{bs}) (r_1 + r_2) &\stackrel{\text{def}}{=} \text{let } (v, \text{bs}_1) = \text{decode}' \text{ bs } r_2 \text{ in } (\text{Right } v, \text{bs}_1) \\
\text{decode}' \text{ bs } (r_1 \cdot r_2) &\stackrel{\text{def}}{=} \text{let } (v_1, \text{bs}_1) = \text{decode}' \text{ bs } r_1 \text{ in} \\
&\quad \text{let } (v_2, \text{bs}_2) = \text{decode}' \text{ bs } r_2 \text{ in } (\text{Seq } v_1 \ v_2, \text{bs}_2) \\
\text{decode}' (Z::\text{bs}) (r^*) &\stackrel{\text{def}}{=} (\text{Stars } [], \text{bs}) \\
\text{decode}' (S::\text{bs}) (r^*) &\stackrel{\text{def}}{=} \text{let } (v, \text{bs}_1) = \text{decode}' \text{ bs } r \text{ in} \\
&\quad \text{let } (\text{Stars } \text{vs}, \text{bs}_2) = \text{decode}' \text{ bs } r^* \text{ in } (\text{Stars } v::\text{vs}, \text{bs}_2) \\
\\
\text{decode } \text{ bs } r &\stackrel{\text{def}}{=} \text{let } (v, \text{bs}') = \text{decode}' \text{ bs } r \text{ in} \\
&\quad \text{if } \text{bs}' = [] \text{ then } \text{Some } v \text{ else } \text{None}
\end{aligned}$$

The function *decode* checks whether all of the bitsequence is consumed and returns the corresponding value as *Some* v ; otherwise it fails with *None*. We can establish that for a value v inhabited by a regular expression r , the decoding of its bitsequence never fails.

► **Lemma 5.** *If $\vdash v : r$ then $\text{decode}(\text{code } v) r = \text{Some } v$.*

Proof. This follows from the property that $\text{decode}'((\text{code } v) @ \text{bs}) r = (v, \text{bs})$ holds for any bitsequence bs and $\vdash v : r$. This property can be easily proved by induction on $\vdash v : r$. ◀

Sulzmann and Lu define the function *internalise* in order to transform (standard) regular expressions into annotated regular expressions. We write this operation as r^\uparrow . This internalisation uses the following *fuse* function.

$$\begin{aligned}
\text{fuse } \text{bs } (\text{ZERO}) &\stackrel{\text{def}}{=} \text{ZERO} \\
\text{fuse } \text{bs } (\text{ONE } \text{bs}') &\stackrel{\text{def}}{=} \text{ONE } (\text{bs } @ \text{bs}') \\
\text{fuse } \text{bs } (\text{CHAR } \text{bs}' \ c) &\stackrel{\text{def}}{=} \text{CHAR } (\text{bs } @ \text{bs}') \ c \\
\text{fuse } \text{bs } (\text{ALTS } \text{bs}' \ r \ s) &\stackrel{\text{def}}{=} \text{ALTS } (\text{bs } @ \text{bs}') \ r \ s \\
\text{fuse } \text{bs } (\text{SEQ } \text{bs}' \ r_1 \ r_2) &\stackrel{\text{def}}{=} \text{SEQ } (\text{bs } @ \text{bs}') \ r_1 \ r_2 \\
\text{fuse } \text{bs } (\text{STAR } \text{bs}' \ r) &\stackrel{\text{def}}{=} \text{STAR } (\text{bs } @ \text{bs}') \ r
\end{aligned}$$

A regular expression can then be *internalised* into a bitcoded regular expression as follows:

$$\begin{aligned}
(\mathbf{0})^\uparrow &\stackrel{\text{def}}{=} \text{ZERO} \\
(\mathbf{1})^\uparrow &\stackrel{\text{def}}{=} \text{ONE } [] \\
(c)^\uparrow &\stackrel{\text{def}}{=} \text{CHAR } [] \ c \\
(r_1 + r_2)^\uparrow &\stackrel{\text{def}}{=} \text{ALT } [] \ (\text{fuse } [Z] \ r_1^\uparrow) \ (\text{fuse } [S] \ r_2^\uparrow) \\
(r_1 \cdot r_2)^\uparrow &\stackrel{\text{def}}{=} \text{SEQ } [] \ r_1^\uparrow \ r_2^\uparrow \\
(r^*)^\uparrow &\stackrel{\text{def}}{=} \text{STAR } [] \ r^\uparrow
\end{aligned}$$

There is also an *erase*-function, written r^\downarrow , which transforms a bitcoded regular expression into a (standard) regular expression by just erasing the annotated bitsequences. We omit the straightforward definition. For defining the algorithm, we also need the functions *bnullable* and *bmkeys*(s), which are the “lifted” versions of *nullable* and *mkeys* acting on bitcoded regular expressions.

| | | | |
|------------------------------|---|---------------------------|--|
| $bnullable (ZERO)$ | $\stackrel{\text{def}}{=} False$ | $bmkeps (ONE bs)$ | $\stackrel{\text{def}}{=} bs$ |
| $bnullable (ONE bs)$ | $\stackrel{\text{def}}{=} True$ | $bmkeps (ALTs bs rs)$ | $\stackrel{\text{def}}{=} bs @ bmkeps rs$ |
| $bnullable (CHAR bs c)$ | $\stackrel{\text{def}}{=} False$ | $bmkeps (SEQ bs r_1 r_2)$ | $\stackrel{\text{def}}{=} bs @ bmkeps r_1 @ bmkeps r_2$ |
| $bnullable (ALTs bs rs)$ | $\stackrel{\text{def}}{=} \exists r \in rs. bnullable r$ | $bmkeps (STAR bs r)$ | $\stackrel{\text{def}}{=} bs @ [S]$ |
| $bnullable (SEQ bs r_1 r_2)$ | $\stackrel{\text{def}}{=} bnullable r_1 \wedge bnullable r_2$ | $bmkeps (r :: rs)$ | $\stackrel{\text{def}}{=} \text{if } bnullable r \text{ then } bmkeps r \text{ else } bmkeps rs$ |
| $bnullable (STAR bs r)$ | $\stackrel{\text{def}}{=} True$ | | |

The key function in the bitcoded algorithm is the derivative of a bitcoded regular expression. This derivative function calculates the derivative but at the same time also the incremental part of the bitsequences that contribute to constructing a POSIX value.

| | |
|--------------------------------|---|
| $(ZERO) \setminus c$ | $\stackrel{\text{def}}{=} ZERO$ |
| $(ONE bs) \setminus c$ | $\stackrel{\text{def}}{=} ZERO$ |
| $(CHAR bs d) \setminus c$ | $\stackrel{\text{def}}{=} \text{if } c = d \text{ then } ONE bs \text{ else } ZERO$ |
| $(ALTs bs rs) \setminus c$ | $\stackrel{\text{def}}{=} ALTs bs (map (_ \setminus c) rs)$ |
| $(SEQ bs r_1 r_2) \setminus c$ | $\stackrel{\text{def}}{=} \text{if } bnullable r_1 \text{ then } ALT bs (SEQ [] (r_1 \setminus c) r_2) \text{ (fuse (bmkeps } r_1) (r_2 \setminus c)) \text{ else } SEQ bs (r_1 \setminus c) r_2$ |
| $(STAR bs r) \setminus c$ | $\stackrel{\text{def}}{=} SEQ bs (fuse [Z](r \setminus c)) (STAR [] r)$ |

This function can also be extended to strings, written $r \setminus s$, just like the standard derivative. We omit the details. Finally we can define Sulzmann and Lu’s bitcoded lexer, which we call *blexer*:

$$blexer r s \stackrel{\text{def}}{=} \text{let } r_{der} = (r^\dagger) \setminus s \text{ in } \text{if } bnullable(r_{der}) \text{ then } decode (bmkeps r_{der}) r \text{ else } None$$

This bitcoded lexer first internalises the regular expression r and then builds the bitcoded derivative according to s . If the derivative is (b)nullable the string is in the language of r and it extracts the bitsequence using the *bmkeps* function. Finally it decodes the bitsequence into a value. If the derivative is *not* nullable, then *None* is returned. We can show that this way of calculating a value generates the same result as *lexer*.

Before we can proceed we need to define a helper function, called *retrieve*, which Sulzmann and Lu introduced for the correctness proof.

| | |
|---|--|
| $retrieve (ONE bs) (Empty)$ | $\stackrel{\text{def}}{=} bs$ |
| $retrieve (CHAR bs c) (Char d)$ | $\stackrel{\text{def}}{=} bs$ |
| $retrieve (ALTs bs [r]) v$ | $\stackrel{\text{def}}{=} bs @ retrieve r v$ |
| $retrieve (ALTs bs (r :: rs)) (Left v)$ | $\stackrel{\text{def}}{=} bs @ retrieve r v$ |
| $retrieve (ALTs bs (r :: rs)) (Right v)$ | $\stackrel{\text{def}}{=} bs @ retrieve (ALTs [] rs) v$ |
| $retrieve (SEQ bs r_1 r_2) (Seq v_1 v_2)$ | $\stackrel{\text{def}}{=} bs @ retrieve r_1 v_1 @ retrieve r_2 v_2$ |
| $retrieve (STAR bs r) (Stars [])$ | $\stackrel{\text{def}}{=} bs @ [S]$ |
| $retrieve (STAR bs r) (Stars (v :: vs))$ | $\stackrel{\text{def}}{=} bs @ [Z] @ retrieve r v @ retrieve (STAR [] r) (Stars vs)$ |

The idea behind this function is to retrieve a possibly partial bitsequence from a bitcoded regular expression, where the retrieval is guided by a value. For example if the value is *Left* then we descend into the left-hand side of an alternative in order to assemble the bitcode. Similarly for *Right*. The

property we can show is that for a given v and r with $\vdash v : r$, the retrieved bitsequence from the internalised regular expression is equal to the bitcoded version of v .

► **Lemma 6.** *If $\vdash v : r$ then $\text{code } v = \text{retrieve } (r^\dagger) v$.*

We also need some auxiliary facts about how the bitcoded operations relate to the “standard” operations on regular expressions. For example if we build a bitcoded derivative and erase the result, this is the same as if we first erase the bitcoded regular expression and then perform the “standard” derivative operation.

► **Lemma 7.**

- (1) $(r \setminus s)^\dagger = (r^\dagger) \setminus s$
- (2) $\text{bnulable}(r) \text{ iff } \text{nulable}(r^\dagger)$
- (3) $\text{bmkeps}(r) = \text{retrieve } r (\text{mkeps } (r^\dagger)) \text{ provided } \text{nulable}(r^\dagger)$.

Proof. All properties are by induction on annotated regular expressions. There are no interesting cases. ◀

The only difficulty left for the correctness proof is that the bitcoded algorithm has only a “forward phase” where POSIX values are generated incrementally. We can achieve the same effect with *lexer* (which has two phases) by stacking up injection functions during the forward phase. An auxiliary function, called *flex*, allows us to recast the rules of *lexer* in terms of a single phase and stacked up injection functions.

$$\begin{aligned} \text{flex } r \ f \ [] & \stackrel{\text{def}}{=} f \\ \text{flex } r \ f \ (c :: s) & \stackrel{\text{def}}{=} \text{flex } (r \setminus c) (\lambda v. f (\text{inj } r \ c \ v)) s \end{aligned}$$

The point of this function is that when reaching the end of the string, we just need to apply the stacked up injection functions to the value generated by *mkeps*. Using this function we can recast the success case in *lexer* as follows:

► **Proposition 8.** *If $\text{lexer } r \ s = \text{Some } v$ then $v = \text{flex } r \ \text{id } s (\text{mkeps}(r \setminus s))$.*

Note we did not redefine *lexer*, we just established that the value generated by *lexer* can also be obtained by a different method. While this different method is not efficient (we essentially need to traverse the string s twice, once for building the derivative $r \setminus s$ and another time for stacking up injection functions using *flex*), it helps us with proving that incrementally building up values in *blexer* generates the same result.

This brings us to our main lemma in this section: if we calculate a derivative, say $r \setminus s$, and have a value, say v , inhabited by this derivative, then we can produce the result *lexer* generates by applying this value to the stacked-up injection functions that *flex* assembles. The lemma establishes that this is the same value as if we build the annotated derivative $r^\dagger \setminus s$ and then retrieve the corresponding bitcoded version, followed by a decoding step.

► **Lemma 9 (Main Lemma).** *If $\vdash v : r \setminus s$ then*

$$\text{Some } (\text{flex } r \ \text{id } s \ v) = \text{decode}(\text{retrieve } (r^\dagger \setminus s) \ v) \ r$$

Proof. This can be proved by induction on s and generalising over v . The interesting point is that we need to prove this in the reverse direction for s . This means instead of cases $[]$ and $c :: s$, we have cases $[]$ and $s @ [c]$ where we unravel the string from the back.²

² Isabelle/HOL provides an induction principle for this way of performing the induction.

XX:10 POSIX Lexing with Bitcoded Derivatives

The case for \square is routine using Lemmas 5 and 6. In the case $s @ [c]$, we can infer from the assumption that $\vdash v : (r \setminus s) \setminus c$ holds. Hence by Prop. 3 we know that $(*) \vdash inj (r \setminus s) c v : r \setminus s$ holds too. By definition of *flex* we can unfold the left-hand side to be

$$Some (flex r id (s @ [c]) v) = Some (flex r id s (inj (r \setminus s) c v))$$

By induction hypothesis and $(*)$ we can rewrite the right-hand side to

$$decode (retrieve (r^\uparrow \setminus s) (inj (r \setminus s) c v)) r$$

which is equal to $decode (retrieve (r^\uparrow \setminus (s @ [c])) v) r$ as required. The last rewrite step is possible because we generalised over v in our induction. \blacktriangleleft

With this lemma in place, we can prove the correctness of *blexer*—it indeed produces the same result as *lexer*.

► **Theorem 10.** $lexer r s = blexer r s$

Proof. We can first expand both sides using Prop. 8 and the definition of *blexer*. This gives us two *if*-statements, which we need to show to be equal. By Lemma 7(2) we know the *if*-tests coincide:

$$bnullable(r^\uparrow \setminus s) \text{ iff } nullable(r \setminus s)$$

For the *if*-branch suppose $r_d \stackrel{\text{def}}{=} r^\uparrow \setminus s$ and $d \stackrel{\text{def}}{=} r \setminus s$. We have $(*) nullable d$. We can then show by Lemma 7(3) that

$$decode(bmkeys r_d) r = decode(retrieve r_d (mkeys d)) r$$

where the right-hand side is equal to $Some (flex r id s (mkeys d))$ by Lemma 9 (we know $\vdash mkeys d : d$ by $(*)$). This shows the *if*-branches return the same value. In the *else*-branches both *lexer* and *blexer* return *None*. Therefore we can conclude the proof. \blacktriangleleft

This establishes that the bitcoded algorithm by Sulzmann and Lu *without* simplification produces correct results. This was only conjectured by Sulzmann and Lu in their paper [12]. The next step is to add simplifications.

4 Simplification

Derivatives as calculated by Brzozowski’s method are usually more complex regular expressions than the initial one; the result is that derivative-based matching and lexing algorithms are often abysmally slow if the “growth problem” is not addressed. As Sulzmann and Lu wrote, various optimisations are possible, such as the simplifications $\mathbf{0} r \Rightarrow \mathbf{0}$, $\mathbf{1} r \Rightarrow r$, $\mathbf{0} + r \Rightarrow r$ and $r + r \Rightarrow r$. While these simplifications can considerably speed up the two algorithms in many cases, they do not solve fundamentally the growth problem with derivatives. To see this let us return to the example from the Introduction that shows the derivatives for $(a + aa)^*$. If we delete in the 3rd step all $\mathbf{0}s$ and $\mathbf{1}s$ according to the simplification rules shown above we obtain

$$(a + aa)^* \xrightarrow{-\setminus[a,a,a]} \underbrace{(\mathbf{1} + a) \cdot (a + aa)^*}_r + ((a + aa)^* + \underbrace{(\mathbf{1} + a) \cdot (a + aa)^*}_r) \quad (1)$$

This is a simpler derivative, but unfortunately we cannot make any further simplifications. This is a problem because the outermost alternatives contains two copies of the same regular expression

(underlined with r). These copies will spawn new copies in later derivative steps and they in turn even more copies. This destroys any hope of taming the size of the derivatives. But the second copy of r in (1) will never contribute to a value, because POSIX lexing will always prefer matching a string with the first copy. So it could be safely removed without affecting the correctness of the algorithm. The dilemma with the simple-minded simplification rules above is that the rule $r + r \Rightarrow r$ will never be applicable because as can be seen in this example the regular expressions are not next to each other but separated by another regular expression.

But here is where Sulzmann and Lu’s representation of generalised alternatives in the bitcoded algorithm shines: in $ALTs\ bs\ rs$ we can define a more aggressive simplification by recursively simplifying all regular expressions in rs and then analyse the resulting list and remove any duplicates. Another advantage with the bitsequences in bitcoded regular expressions is that they can be easily modified such that simplification does not interfere with the value constructions. For example we can “flatten”, or de-nest, $ALTs$ as follows

$$ALTs\ bs_1\ (ALTs\ bs_2\ rs_2 :: rs_1) \xrightarrow{bsimp} ALTs\ bs_1\ (map\ (fuse\ bs_2)\ rs_2 :: rs_1)$$

where we just need to fuse the bitsequence that has accumulated in bs_2 to the alternatives in rs_2 . As we shall show below this will ensure that the correct value corresponding to the original (unsimplified) regular expression can still be extracted.

However there is one problem with the definition for the more aggressive simplification rules described by Sulzmann and Lu. Recasting their definition with our syntax they define the step of removing duplicates as

$$bsimp\ (ALTs\ bs\ rs) \stackrel{\text{def}}{=} ALTs\ bs\ (nub\ (map\ bsimp\ rs))$$

where they first recursively simplify the regular expressions in rs (using map) and then use Haskell’s nub -function to remove potential duplicates. While this makes sense when considering the example shown in (1), nub is the inappropriate function in the case of bitcoded regular expressions. The reason is that in general the elements in rs will have a different annotated bitsequence and in this way nub will never find a duplicate to be removed. One correct way to handle this situation is to first *erase* the regular expressions when comparing potential duplicates. This is inspired by Scala’s list functions of the form $distinctBy\ rs\ f\ acc$ where a function is applied first before two elements are compared. We define this function in Isabelle/HOL as

$$\begin{aligned} distinctBy\ []\ f\ acc &\stackrel{\text{def}}{=} [] \\ distinctBy\ (x :: xs)\ f\ acc &\stackrel{\text{def}}{=} \text{if } fx \in acc \text{ then } distinctBy\ xs\ f\ acc \text{ else } x :: distinctBy\ xs\ f\ (\{fx\} \cup acc) \end{aligned}$$

where we scan the list from left to right (because we have to remove later copies). In $distinctBy, f$ is a function and acc is an accumulator for regular expressions—essentially a set of regular expressions that we have already seen while scanning the list. Therefore we delete an element, say x , from the list provided fx is already in the accumulator; otherwise we keep x and scan the rest of the list but add fx as another “seen” element to acc . We will use $distinctBy$ where f is the erase function, $_ \downarrow$, that deletes bitsequences from bitcoded regular expressions. This is clearly a computationally more expensive operation than nub , but is needed in order to make the removal of unnecessary copies to work properly.

Our simplification function depends on three helper functions, one is called fts and analyses lists of regular expressions coming from alternatives. It is defined as follows:

$$\begin{aligned} fts\ [] &\stackrel{\text{def}}{=} [] \\ fts\ (ZERO :: rs) &\stackrel{\text{def}}{=} fts\ rs \\ fts\ (ALTs\ bs'\ rs' :: rs) &\stackrel{\text{def}}{=} map\ (fuse\ bs')\ rs' @ fts\ rs \end{aligned}$$

XX:12 POSIX Lexing with Bitcoded Derivatives

The second clause of *fts* removes all instances of *ZERO* in alternatives and the third “spills” out nested alternatives (but retaining the bitsequence *bs*’ accumulated in the inner alternative). There are some corner cases to be considered when the resulting list inside an alternative is empty or a singleton list. We take care of those cases in the *bsimpALTs* function; similarly we define a helper function that simplifies sequences according to the usual rules about *ZERO*s and *ONE*s:

$$\begin{array}{lll}
 \text{bsimpALTs } bs \ [] \stackrel{\text{def}}{=} \text{ZERO} & \text{bsimpSEQ } bs \ _ \text{ZERO} \stackrel{\text{def}}{=} \text{ZERO} & \\
 \text{bsimpALTs } bs \ [r] \stackrel{\text{def}}{=} \text{fuse } bs \ r & \text{bsimpSEQ } bs \ \text{ZERO } _ \stackrel{\text{def}}{=} \text{ZERO} & \\
 \text{bsimpALTs } bs \ rs \stackrel{\text{def}}{=} \text{ALTs } bs \ rs & \text{bsimpSEQ } bs_1 \ (\text{ONE } bs_2) \ r_2 \stackrel{\text{def}}{=} \text{fuse } (bs_1 \ @ \ bs_2) \ r_2 & \\
 & \text{bsimpSEQ } bs \ r_1 \ r_2 \stackrel{\text{def}}{=} \text{SEQ } bs \ r_1 \ r_2 &
 \end{array}$$

With this in place we can define our simplification function as

$$\begin{array}{ll}
 \text{bsimp } (\text{SEQ } bs \ r_1 \ r_2) \stackrel{\text{def}}{=} \text{bsimpSEQ } bs \ (\text{bsimp } r_1) \ (\text{bsimp } r_2) & \\
 \text{bsimp } (\text{ALTs } bs \ rs) \stackrel{\text{def}}{=} \text{bsimpALT } bs \ (\text{distinctBy } (\text{fts } (\text{map } \text{bsimp } rs)) \ \text{erase } \emptyset) & \\
 \text{bsimp } r \stackrel{\text{def}}{=} r &
 \end{array}$$

As far as we can see, our recursive function *bsimp* simplifies regular expressions as intended by Sulzmann and Lu. There is no point in applying the *bsimp* function repeatedly (like the simplification in their paper which needs to be applied until a fixpoint is reached) because we can show that *bsimp* is idempotent, that is

► **Proposition 11.** $\text{bsimp } (\text{bsimp } r) = \text{bsimp } r$

This can be proved by induction on *r* but requires a detailed analysis that the de-nesting of alternatives always results in a flat list of regular expressions. We omit the details since it does not concern the correctness proof.

Next we can include simplification after each derivative step leading to the following notion of bitcoded derivatives:

$$r \setminus_{\text{bsimp}} [] \stackrel{\text{def}}{=} r \quad r \setminus_{\text{bsimp}} (c :: s) \stackrel{\text{def}}{=} \text{bsimp } (r \setminus c) \setminus_{\text{bsimp}} s$$

and use it in the improved lexing algorithm defined as

$$\text{blexer}^+ r \ s \stackrel{\text{def}}{=} \text{let } r_{\text{der}} = (r^\dagger) \setminus_{\text{bsimp}} s \text{ in} \\
 \text{if } \text{bnulable}(r_{\text{der}}) \text{ then } \text{decode } (\text{bmkeys } r_{\text{der}}) \ r \ \text{else } \text{None}$$

The remaining task is to show that *blexer* and *blexer*⁺ generate the same answers.

When we first attempted this proof we encountered a problem with the idea in Sulzmann and Lu’s paper where the argument seems to be to appeal again to the *retrieve*-function defined for the unsimplified version of the algorithm. But this does not work, because desirable properties such as

$$\text{retrieve } r \ v = \text{retrieve } (\text{bsimp } r) \ v$$

do not hold under simplification—this property essentially purports that we can retrieve the same value from a simplified version of the regular expression. To start with *retrieve* depends on the fact that the value *v* correspond to the structure of the regular expressions—but the whole point of simplification is to “destroy” this structure by making the regular expression simpler. To see this consider the regular expression $r = r' + 0$ and a corresponding value $v = \text{Left } v'$. If we annotate bitcodes to *r*, then we can use *retrieve* and *v* in order to extract a corresponding bitsequence. The reason that this works is that *r* is an alternative regular expression and *v* a corresponding value. However, if we simplify *r*, then

v does not correspond to the shape of the regular expression anymore. So unless one can somehow synchronise the change in the simplified regular expressions with the original POSIX value, there is no hope of appealing to *retrieve* in the correctness argument for *blexer*⁺.

We found it more helpful to introduce the rewriting systems shown in Figure 3. The idea is to generate simplified regular expressions in small steps (unlike the *bsimp*-function which does the same in a big step), and show that each of the small steps preserves the bitcodes that lead to the final POSIX value. The rewrite system is organised such that \rightsquigarrow is for bitcoded regular expressions and \rightsquigarrow^s for lists of bitcoded regular expressions. The former essentially implements the simplifications of *bsimpSEQ* and *fts*; while the latter implements the simplifications in *bsimpALTs*. We can show that any bitcoded regular expression reduces in zero or more steps to the simplified regular expression generated by *bsimp*:

► **Lemma 12.** $r \rightsquigarrow^* \text{bsimp } r$

Proof. By induction on r . For this we can use the properties $rs \rightsquigarrow^s \text{fts } rs$ and $rs \rightsquigarrow^s \text{distinctBy } rs \text{ erase } \emptyset$. The latter uses repeated applications of the *LD* rule which allows the removal of duplicates that can recognise the same strings. ◀

We can show that this rewrite system preserves *bnullable*, that is simplification, essentially, does not affect nullability:

► **Lemma 13.** *If $r_1 \rightsquigarrow r_2$ then $\text{bnullable } r_1 = \text{bnullable } r_2$.*

Proof. Straightforward mutual induction on the definition of \rightsquigarrow and \rightsquigarrow^s . The only interesting case is the rule *LD* where the property holds since by the side-conditions of that rule the empty string will be in both $L(rs_a @ [r_1] @ rs_b @ [r_2] @ rs_c)$ and $L(rs_a @ [r_1] @ rs_b @ rs_c)$. ◀

From this, we can show that *bmkeys* will produce the same bitsequence as long as one of the bitcoded regular expressions in \rightsquigarrow is nullable (this lemma establishes the missing fact we were not able to establish using *retrieve*, as suggested in the paper by Sulzmann and Lu).

► **Lemma 14.** *If $r_1 \rightsquigarrow r_2$ and $\text{bnullable } r_1$ then $\text{bmkeys } r_1 = \text{bmkeys } r_2$.*

Proof. By straightforward mutual induction on the definition of \rightsquigarrow and \rightsquigarrow^s . Again the only interesting case is the rule *LD* where we need to ensure that

$$\text{bmkeys}(rs_a @ [r_1] @ rs_b @ [r_2] @ rs_c) = \text{bmkeys}(rs_a @ [r_1] @ rs_b @ rs_c)$$

holds. This is indeed the case because according to the POSIX rules the generated bitsequence is determined by the first alternative that can match the string (in this case being nullable). ◀

Crucial is also the fact that derivative steps and simplification steps can be interleaved, which is shown by the fact that \rightsquigarrow is preserved under derivatives.

► **Lemma 15.** *If $r_1 \rightsquigarrow r_2$ then $r_1 \setminus c \rightsquigarrow^* r_2 \setminus c$.*

Proof. By straightforward mutual induction on the definition of \rightsquigarrow and \rightsquigarrow^s . The case for *LD* holds because $L((r_2 \setminus c)^\downarrow) \subseteq L((r_1 \setminus c)^\downarrow)$ if and only if $L(r_2^\downarrow) \subseteq L(r_1^\downarrow)$. ◀

Using this fact together with Lemma 12 allows us to prove the central lemma that the unsimplified derivative (with a string s) reduces to the simplified derivative (with the same string).

► **Lemma 16.** $r \setminus s \rightsquigarrow^* r \setminus \text{bsimp } s$

Proof. By reverse induction on s generalising over r . ◀

$$\begin{array}{c}
\frac{}{(SEQ\ bs\ ZERO\ r_2) \rightsquigarrow (ZERO)} S0_l \quad \frac{}{(SEQ\ bs\ r_1\ ZERO) \rightsquigarrow (ZERO)} S0_r \\
\frac{}{(SEQ\ bs_1\ (ONE\ bs_2)\ r) \rightsquigarrow fuse\ (bs_1\ @\ bs_2)\ r} S1 \\
\frac{r_1 \rightsquigarrow r_2}{(SEQ\ bs\ r_1\ r_3) \rightsquigarrow (SEQ\ bs\ r_2\ r_3)} SL \quad \frac{r_3 \rightsquigarrow r_4}{(SEQ\ bs\ r_1\ r_3) \rightsquigarrow (SEQ\ bs\ r_1\ r_4)} SR \\
\frac{}{(ALTs\ bs\ []) \rightsquigarrow (ZERO)} A0 \quad \frac{}{(ALTs\ bs\ [r]) \rightsquigarrow fuse\ bs\ r} A1 \\
\frac{rs_1 \rightsquigarrow^s rs_2}{(ALTs\ bs\ rs_1) \rightsquigarrow (ALTs\ bs\ rs_2)} AL \\
\frac{rs_1 \rightsquigarrow^s rs_2}{r :: rs_1 \rightsquigarrow^s r :: rs_2} LH \quad \frac{r_1 \rightsquigarrow r_2}{r_1 :: rs \rightsquigarrow^s r_2 :: rs} LT \\
\frac{}{ZERO :: rs \rightsquigarrow^s rs} L0 \quad \frac{}{ALTs\ bs\ rs_1 :: rs_2 \rightsquigarrow^s (map\ (fuse\ bs)\ rs_1\ @\ rs_2)} LS \\
\frac{L(r_2^\downarrow) \subseteq L(r_1^\downarrow)}{(rs_1\ @\ [r_1]\ @\ rs_2\ @\ [r_2]\ @\ rs_3) \rightsquigarrow^s (rs_1\ @\ [r_1]\ @\ rs_2\ @\ rs_3)} LD
\end{array}$$

■ **Figure 3** The rewrite rules that generate simplified regular expressions in small steps: $r_1 \rightsquigarrow r_2$ is for bitcoded regular expressions and $rs_1 \rightsquigarrow^s rs_2$ for lists of bitcoded regular expressions. Interesting is the *LD* rule that allows copies of regular expressions to be removed provided a regular expression earlier in the list can match the same strings.

With these lemmas in place we can finally establish that $blexer^+$ and $blexer$ generate the same value, and using Theorem 10 from the previous section that this value is indeed the POSIX value.

► **Theorem 17.** $blexer\ r\ s = blexer^+\ r\ s$

Proof. By unfolding the definitions and using Lemmas 16 and 14. ◀

This completes the correctness proof for the second POSIX lexing algorithm by Sulzmann and Lu. The interesting point of this algorithm is that the sizes of derivatives do not grow arbitrarily, which we shall show next.

5 Finiteness of Derivatives

In this section let us sketch our argument for why the size of the simplified derivatives with the aggressive simplification function is finite. Suppose we have a size function for bitcoded regular expressions, written $|r|$, which counts the number of nodes if we regard r as a tree (we omit the precise definition). For this we show that for every r there exists a bound N such that

$$\forall s. |r \setminus_{bsimp} s| < N$$

We prove this by induction on r . The base cases for *ZERO*, *ONE bs* and *CHAR bs c* are straightforward. The interesting case is for sequences of the form *SEQ bs r₁ r₂*. In this case our induction hypotheses state $\forall s. |r_1 \setminus_{bsimp} s| < N_1$ and $\forall s. |r_2 \setminus_{bsimp} s| < N_2$. We can reason as follows

$$\begin{aligned}
& |(SEQ\ bs\ r_1\ r_2)\backslash_{bsimp\ s}| \\
= & |bsimp(ALTs\ bs\ ((r_1\backslash_{bsimp\ s}) \cdot r_2) :: [r_2\backslash_{bsimp\ s'} \mid s' \in Suffix(s)])| \quad (1) \\
\leq & |distinctBy\ (flts\ ((r_1\backslash_{bsimp\ s}) \cdot r_2) :: [r_2\backslash_{bsimp\ s'} \mid s' \in Suffix(s)])| + 1 \quad (2) \\
\leq & |(r_1\backslash_{bsimp\ s}) \cdot r_2| + |distinctBy\ (flts\ [r_2\backslash_{bsimp\ s'} \mid s' \in Suffix(s)])| + 1 \quad (3) \\
\leq & N_1 + |r_2| + 2 + |distinctBy\ (flts\ [r_2\backslash_{bsimp\ s'} \mid s' \in Suffix(s)])| \quad (4) \\
\leq & N_1 + |r_2| + 2 + l_{N_2} * N_2 \quad (5)
\end{aligned}$$

where in (1) the $Suffix(s')$ are the suffixes where $r_1\backslash_{bsimp\ s''}$ is nullable for $s = s'' @ s'$. In (3) we know that $|(r_1\backslash_{bsimp\ s}) \cdot r_2|$ is bounded by $N_1 + |r_2|$. In (5) we know the list comprehension contains only regular expressions of size smaller than N_2 . The list length after $distinctBy$ is bounded by a number, which we call l_{N_2} . It stands for the number of distinct regular expressions with a maximum size N_2 (there can only be finitely many of them). We reason similarly in the *Star*-case.

Clearly we give in this finiteness argument (Step (5)) a very loose bound that is far from the actual bound we can expect. We can do better than this, but this does not improve the finiteness property we are proving. If we are interested in a polynomial bound, one would hope to obtain a similar tight bound as for partial derivatives introduced by Antimirov [2]. After all the idea with $distinctBy$ is to maintain a “set” of alternatives (like the sets in partial derivatives). Unfortunately to obtain the exact same bound would mean we need to introduce simplifications such as

$$(r_1 + r_2) \cdot r_3 \longrightarrow (r_1 \cdot r_3) + (r_2 \cdot r_3)$$

which exist for partial derivatives. However, if we introduce them in our setting we would lose the POSIX property of our calculated values. We leave better bounds for future work.

6 Conclusion

We set out in this work to prove in Isabelle/HOL the correctness of the second POSIX lexing algorithm by Sulzmann and Lu [12]. This follows earlier work where we established the correctness of the first algorithm [3]. In the earlier work we needed to introduce our own specification about what POSIX values are, because the informal definition given by Sulzmann and Lu did not stand up to a formal proof. Also for the second algorithm we needed to introduce our own definitions and proof ideas in order to establish the correctness. Our interest in the second algorithm lies in the fact that by using bitcoded regular expressions and an aggressive simplification method there is a chance that the derivatives can be kept universally small (we established in this paper that they can be kept finite for any string). This is important if one is after an efficient POSIX lexing algorithm.

Having proved the correctness of the POSIX lexing algorithm, which lessons have we learned? Well, we feel this is a very good example where formal proofs give further insight into the matter at hand. For example it is very hard to see a problem with *nub* vs *distinctBy* with only experimental data—one would still see the correct result but find that simplification does not simplify in well-chosen, but not obscure, examples. We found that from an implementation point-of-view it is really important to have the formal proofs of the corresponding properties at hand.

We have also developed a healthy suspicion when experimental data is used to back up efficiency claims. For example Sulzmann and Lu write about their equivalent of *blexer*⁺ “...we can incrementally compute bitcoded parse trees in linear time in the size of the input” [12, Page 14]. Given the growth of the derivatives in some cases even after aggressive simplification, this is a hard to believe fact. A similar claim about a theoretical runtime of $O(n^2)$ is made for the Verbatim lexer, which calculates tokens according to POSIX rules [6]. For this it uses Brzozowski’s derivatives like in our work. They write: “The results of our empirical tests [...] confirm that Verbatim has $O(n^2)$ time complexity.” [6, Section VII]. While their correctness proof for Verbatim is formalised in Coq, the claim about

the runtime complexity is only supported by some empirical evidence obtained by using the code extraction facilities of Coq. In the context of our observation with the “growth problem” of derivatives, we tried out their extracted OCaml code with the example $(a + aa)^*$ as a single lexing rule, and it took for us around 5 minutes to tokenise a string of 40 a ’s and that increased to approximately 19 minutes when the string is 50 a ’s long. Given that derivatives are not simplified in the Verbatim lexer, such numbers are not surprising. Clearly our result of having finite derivatives might sound rather weak in this context but we think such efficiency claims really require further scrutiny.

Our Isabelle/HOL code is available under <https://github.com/urbanchr/posix>.

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 V. Antimirov. Partial Derivatives of Regular Expressions and Finite Automata Constructions. *Theoretical Computer Science*, 155:291–319, 1995.
- 3 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.
- 4 J. A. Brzozowski. Derivatives of Regular Expressions. *Journal of the ACM*, 11(4):481–494, 1964.
- 5 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.
- 6 D. Egolf, S. Lasser, and K. Fisher. Verbatim: A Verified Lexer Generator. In *2021 IEEE Security and Privacy Workshops (SPW)*, pages 92–100, 2021.
- 7 A. Krauss and T. Nipkow. Proof Pearl: Regular Expression Equivalence and Relation Algebra. *Journal of Automated Reasoning*, 49:95–106, 2012.
- 8 C. Kuklewicz. Regex Posix. https://wiki.haskell.org/Regex_Posix.
- 9 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.
- 10 S. Owens and K. Slind. Adapting Functional Programs to Higher Order Logic. *Higher-Order and Symbolic Computation*, 21(4):377–409, 2008.
- 11 R. Ribeiro and A. Du Bois. Certified Bit-Coded Regular Expression Parsing. In *Proc. of the 21st Brazilian Symposium on Programming Languages*, New York, NY, USA, 2017. Association for Computing Machinery.
- 12 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.
- 13 S. Vansummeren. Type Inference for Unique Pattern Matching. *ACM Transactions on Programming Languages and Systems*, 28(3):389–428, 2006.