# Optimally Streaming Greedy Regular Expression Parsing^ 

Niels Bjørn Bugge Grathwohl, Fritz Henglein, and Ulrik Terp Rasmussen<br>Department of Computer Science, University of Copenhagen (DIKU)


#### Abstract

We study the problem of streaming regular expression parsing: Given a regular expression and an input stream of symbols, how to output a serialized syntax tree representation as an output stream during input stream processing. We show that optimally streaming regular expression parsing, outputting bits of the output as early as is semantically possible for any regular expression of size $m$ and any input string of length $n$, can be performed in time $O\left(2^{m \log m}+m n\right)$ on a unit-cost random-access machine. This is for the wide-spread greedy disambiguation strategy for choosing parse trees of grammatically ambiguous regular expressions. In particular, for a fixed regular expression, the algorithm's run-time scales linearly with the input string length. The exponential is due to the need for preprocessing the regular expression to analyze state coverage of its associated NFA, a PSPACE-hard problem, and tabulating all reachable ordered sets of NFA-states. Previous regular expression parsing algorithms operate in multiple phases, always requiring processing or storing the whole input string before outputting the first bit of output, not only for those regular expressions and input prefixes where reading to the end of the input is strictly necessary.


## 1 Introduction

In programming, regular expressions are often used to extract information from an input, which requires an intensional interpretation of regular expressions as denoting parse trees, and not just their ordinary language-theoretic interpretation as denoting strings.

This is a nontrivial change of perspective. We need to deal with grammatical ambiguity - which parse tree to return, not just that it has one - and memory requirements become a critical factor: Deciding whether a string belongs to the language denoted by $(\mathrm{ab})^{\star}+(\mathrm{a}+\mathrm{b})^{\star}$ can be done in constant space, but outputting the first bit, whether the string matches the first alternative or only the second, may require buffering the whole input string. This is an instructive case of deliberate grammatical ambiguity to be resolved by the prefer-the-leftalternative policy of greedy disambiguation: Try to match the left alternative;

[^0]if that fails, return a match according to the right alternative as a fallback. Straight-forward application of automata-theoretic techniques does not help: $(\mathrm{ab})^{\star}+(\mathrm{a}+\mathrm{b})^{\star}$ denotes the same language as $(\mathrm{a}+\mathrm{b})^{\star}$, which is unambiguous and corresponds to a small DFA, but is also useless: it doesn't represent any more when a string consists of a sequence of $a b$-groups.

Previous parsing algorithms $93 / 5 / 10136$ require at least one full pass over the input string before outputting any output bits representing the parse tree. This is the case even for regular expressions requiring only bounded lookahead such as one-unambiguous regular expressions [1].

In this paper we study the problem of optimally streaming parsing. Consider $(\mathrm{ab})^{\star}+(\mathrm{a}+\mathrm{b})^{\star}$, which is ambiguous and in general requires unbounded input buffering, and consider the particular input string ab ...abaababababab.... An optimally streaming parsing algorithm needs to buffer the prefix $\mathrm{ab} .$. ab in some form because the complete parse might match either of the two alternatives in the regular expression, but once encountering aa, only the right alternative is possible. At this point it outputs this information and the output representation for the buffered string as parsed by the second alternative. After this, it outputs a bit for each input symbol read, with no internal buffering: input symbols are discarded before reading the next symbol. Optimality means that output bits representing the eventual parse tree must be produced earliest possible: as soon as they are semantically determined by the input processed so far under the assumption that the parse will succeed.

Outline. In Section 2 we recall the type interpretation of regular expressions, where a regular expression denotes parse trees, along with the bit-coding of parse trees.

In Section 3 we introduce a class of Thompson-style augmented nondeterministic finite automata (aNFAs). Paths in such an aNFA naturally represent complete parse trees, and paths to intermediate states represent partial parse trees for prefixes of an input string.

We recall the greedy disambiguation strategy in Section 4, which specifies a deterministic mapping of accepted strings to NFA-paths.

Section 5 contains a definition of what it means to be an optimally streaming implementation of a parsing function.

We define what it means for a set of aNFA-states to cover another state in Section 6, which constitutes the computationally hardest part needed in our algorithm.

Section 7 contains the main results. We present path trees as a way of organizing partial parse trees, and based on these we present our algorithm for an optimally streaming parsing function and analyze its asymptotic run-time complexity.

Finally, in Section 8, the algorithm is demonstrated by illustrative examples alluding to its expressive power and practical utility.

## 2 Preliminaries

In the following section, we recall definitions of regular expressions and their interpretation as types [10].

Definition 1 (Regular expression). A regular expression ( $R E$ ) over a finite alphabet $\Sigma$ is an expression $E$ generated by the grammar

$$
E::=\mathbf{0}|\mathbf{1}| a\left|E_{1} E_{2}\right| E_{1}+E_{2} \mid E_{1}^{\star}
$$

where $a \in \Sigma$.
Concatenation (juxtaposition) and alternation ( + ) associates to the right; parentheses may be inserted to override associativity. Kleene star ( $\star$ ) binds tightest, followed by concatenation and alternation.

The standard interpretation of regular expressions is as descriptions of regular languages.

Definition 2 (Language interpretation). Every $R E E$ denotes a language $\mathcal{L} \llbracket E \rrbracket \subseteq \Sigma^{\star}$ given as follows:

$$
\begin{aligned}
& \mathcal{L} \llbracket \mathbf{0} \rrbracket=\emptyset \quad \mathcal{L} \llbracket E_{1} E_{2} \rrbracket=\mathcal{L} \llbracket E_{1} \rrbracket \mathcal{L} \llbracket E_{2} \rrbracket \quad \mathcal{L} \llbracket a \rrbracket=\{a\} \\
& \mathcal{L} \llbracket \mathbf{1} \rrbracket=\{\epsilon\} \quad \mathcal{L} \llbracket E_{1}+E_{2} \rrbracket=\mathcal{L} \llbracket E_{1} \rrbracket \cup \mathcal{L} \llbracket E_{2} \rrbracket \quad \mathcal{L} \llbracket E_{1}^{\star} \rrbracket=\bigcup_{n \geq 0} \mathcal{L} \llbracket E_{1} \rrbracket^{n}
\end{aligned}
$$

where we have $A_{1} A_{2}=\left\{w_{1} w_{2} \mid w_{1} \in A_{1}, w_{2} \in A_{2}\right\}$, and $A^{0}=\{\epsilon\}$ and $A^{n+1}=$ $A A^{n}$.

Proviso: Henceforth we shall restrict ourselves to REs $E$ such that $\mathcal{L} \llbracket E \rrbracket \neq \emptyset$.
For regular expression parsing, we consider an alternative interpretation of regular expressions as types.

Definition 3 (Type interpretation). Let the syntax of values be given by

$$
v::=()\left|\operatorname{inl} v_{1}\right| \operatorname{inr} v_{1}\left|\left\langle v_{1}, v_{2}\right\rangle\right|\left[v_{1}, v_{2}, \ldots, v_{n}\right]
$$

Every $R E E$ can be seen as a type describing a set $\mathcal{T} \llbracket E \rrbracket$ of well-typed values:

$$
\begin{aligned}
& \mathcal{T} \llbracket \mathbf{0} \rrbracket=\emptyset \quad \mathcal{T} \llbracket E_{1} E_{2} \rrbracket=\left\{\left\langle v_{1}, v_{2}\right\rangle \mid v_{1} \in \mathcal{T} \llbracket E_{1} \rrbracket, v_{2} \in \mathcal{T} \llbracket E_{2} \rrbracket\right\} \\
& \mathcal{T} \llbracket \mathbf{1} \rrbracket=\{()\} \quad \mathcal{T} \llbracket E_{1}+E_{2} \rrbracket=\left\{\text { inl } v \mid v \in \mathcal{T} \llbracket E_{1} \rrbracket\right\} \cup\left\{\text { inr } v \mid v \in \mathcal{T} \llbracket E_{2} \rrbracket\right\} \\
& \mathcal{T} \llbracket a \rrbracket=\{a\} \quad \mathcal{T} \llbracket E_{1}^{\star} \rrbracket=\left\{\left[v_{1}, \ldots, v_{n}\right] \mid n \geq 0 \wedge \forall 1 \leq i \leq n . v_{i} \in \mathcal{T} \llbracket E_{1} \rrbracket\right\}
\end{aligned}
$$

We write $|v|$ for the flattening of a value, defined as the word obtained by doing an in-order traversal of $v$ and writing down all the symbols in the order they are visited. We write $\mathcal{T}_{w} \llbracket E \rrbracket$ for the restricted set $\{v \in \mathcal{T} \llbracket E \rrbracket||v|=w\}$. Regular expression parsing is a generalization of the acceptance problem of determining whether a word $w$ belongs to the language of some RE $E$, where additionally we produce a parse tree from $\mathcal{T}_{w} \llbracket E \rrbracket$. We say that an RE $E$ is ambiguous iff there exists a $w$ such that $\left|\mathcal{T}_{w} \llbracket E \rrbracket\right|>1$.

Any well-typed value can be serialized into a sequence of bits.

Definition 4 (Bit-coding). Given a value $v \in \mathcal{T} \llbracket E \rrbracket$, we denote its bit-code by $\ulcorner v\urcorner \subseteq\{0,1\}^{\star}$, defined as follows:

$$
\left.\left.\begin{array}{rlrl}
\ulcorner()\urcorner & =\epsilon & \ulcorner a\urcorner & =\epsilon \\
\left\ulcorner\left\langle v_{1}, v_{2}\right\rangle\right\urcorner & =\left\ulcorner v_{1}\right\urcorner\left\ulcorner v_{2}\right\urcorner & \left\ulcorner\left[v_{1}, \ldots, v_{n}\right]\right\urcorner & =0\left\ulcorner v_{1}\right\urcorner \ldots 0\left\ulcorner v_{n}\right\urcorner 1
\end{array} r \text { inl } v\right\urcorner=0\ulcorner v\urcorner\right)
$$

We write $\mathcal{B} \llbracket E \rrbracket$ for the set $\{\ulcorner v\urcorner \mid v \in \mathcal{T} \llbracket E \rrbracket\}$ and $\mathcal{B}_{w} \llbracket E \rrbracket$ for the set restricted to bit-codes for values with a flattening $w$. Note that for any $R E E$, bit-coding is an isomorphism when seen as a function $\ulcorner\cdot\urcorner_{E}: \mathcal{T} \llbracket E \rrbracket \rightarrow \mathcal{B} \llbracket E \rrbracket$.

## 3 Augmented Automata

In this section we recall from [6] the construction of finite automata from regular expressions. Our construction is similar to that of Thompson [15], but augmented with extra annotations on non-deterministic $\epsilon$-transitions. The resulting automata can be seen as non-deterministic transducers which for each accepted input string in the language of the underlying regular expression outputs the bit-codes for the corresponding parse trees.

Definition 5 (Augmented non-deterministic finite automaton). An augmented non-deterministic finite automaton (aNFA) is a tuple (State, $\delta, q^{\text {in }}, q^{\text {fin }}$ ), where State is a finite set of states, $q^{\text {in }}, q^{\text {fin }} \in$ State are initial and final states, respectively, and $\delta \subseteq$ State $\times \Gamma \times$ State is a labeled transition relation with labels $\Gamma=\Sigma \uplus\{0,1, \epsilon\}$.

Transition labels are divided into the disjoint sets $\Sigma$ (symbol labels); $\{0,1\}$ (bit-labels); and $\{\epsilon\}$ ( $\epsilon$-labels). $\Sigma$-transitions can be seen as input actions, and bit-transitions as output actions.

Definition 6 (aNFA construction). Let $E$ be an $R E$ and define an aNFA $M_{E}=\left(\operatorname{State}_{E}, \delta_{E}, q_{E}^{\mathrm{in}}, q_{E}^{\mathrm{fin}}\right)$ by induction on $E$. We give the definition diagrammatically by cases:


In the above, the notation (q)- $q_{1}$ means that $q_{1}, q_{2}$ are initial and final states, respectively, in some (sub-)automaton $M$.

See Figure 1 for an example.


Fig. 1: Example automaton for the $\operatorname{RE}(a+b)^{\star} b$

Definition 7 (Path). A path in an aNFA is a finite non-empty sequence $\alpha \in$ State $^{\star}$ of the form $\alpha=p_{0} p_{1} \ldots p_{n-1}$ such that for each $i<n$, we have $\left(p_{i}, \gamma_{i}, p_{i+1}\right) \in \delta_{E}$ for some $\gamma_{i}$. As a shorthand for this fact we might write $p_{0} \stackrel{\alpha}{\rightsquigarrow} p_{n-1}$ (note that a single state is a path to itself).

Each path $\alpha$ is associated with a (possibly empty) sequence of labels lab $(\alpha)$ : we let read $(\alpha)$ and write $(\alpha)$ refer to the corresponding subsequences of $\operatorname{lab}(\alpha)$ filtered by $\Sigma$ and $\{0,1\}$, respectively. An automaton accepts a word $w$ iff $q^{\text {in }} \stackrel{\alpha}{\sim} q^{\text {fin }}$ for some $\alpha$ where $\operatorname{read}(\alpha)=w$. There is a one-to-one correspondence between bit-codes and accepting paths:

Proposition 1. For any $R E E$ with aNFA $M_{E}$, we have for each $w \in \mathcal{L} \llbracket E \rrbracket$ that

$$
\left\{\text { write }(\alpha) \mid q^{\text {in }} \stackrel{\alpha}{\rightsquigarrow} q^{\text {fin }}, \operatorname{read}(\alpha)=w\right\}=\mathcal{B}_{w} \llbracket E \rrbracket
$$

Determinization. Given a state set $Q$, define its closure as the set closure $(Q)=$ $\left\{q^{\prime} \mid q \in Q \wedge \exists \alpha \cdot \operatorname{read}(\alpha)=\epsilon \wedge q \stackrel{\alpha}{\rightsquigarrow} q^{\prime}\right\}$. For any aNFA $M=\left(\right.$ State, $\left.\delta, q^{\text {in }}, q^{\text {fin }}\right)$, let $D(M)=\left(\right.$ DState $\left._{M}, I_{M}, F_{M}, \Delta_{M}\right)$ be the deterministic automaton obtained by applying the standard subset sum construction: Here, $I_{M}=\operatorname{closure}\left(\left\{q^{\text {in }}\right\}\right)$ is the initial state, and DState ${ }_{M} \subseteq \mathbf{2}^{\text {State }}$ is the set of states, defined to be the smallest set containing $I_{M}$ and closed under the transition function $\Delta_{M}(Q, a)=$ closure $\left(\left\{q^{\prime} \mid\left(q, a, q^{\prime}\right) \in \delta, q \in Q\right\}\right)$. The set of final states $F_{M}$ is the set $\{Q \in$ DState $\left._{M} \mid q^{\text {fin }} \in Q\right\}$.

## 4 Disambiguation

A regular expression parsing algorithm has to produce a parse tree for an input word whenever the word is in the language for the underlying RE. In the case of ambiguous REs, the algorithm has to choose one of several candidates. We do not want the choice to be arbitrary, but rather a parse tree which is uniquely identified by a disambiguation policy. Since there is a one-to-one correspondence between words in the language of an RE $E$ and accepting paths in $M_{E}$, a disambiguation policy can be seen as a deterministic choice between aNFA paths recognizing the same string.

We will focus on greedy disambiguation, which corresponds to choosing the first result that would have been found by a backtracking regular expression
parsing algorithm such as the one found in the Perl programming language [16]. The greedy strategy has successfully been implemented in previous work [56, and is simpler to define and implement than other strategies such as POSIX [8]4] whose known parsing algorithms are technically more complicated [11|1314].

Greedy disambiguation can be seen as picking the accepting path with the lexicographically least bitcode. A well-known problem with backtracking parsing is non-termination in the case of regular expressions with nullable subexpressions under Kleene star, which means that the lexicographically least path is not always well-defined. This problem can easily be solved by not considering paths with non-productive loops, as in [5].

## 5 Optimal Streaming

In this section we specify what it means to be an optimally streaming implementation of a function from sequences to sequences.

We write $w \sqsubseteq w^{\prime \prime}$ if $w$ is a prefix of $w^{\prime \prime}$, that is $w w^{\prime}=w^{\prime \prime}$ for some $w^{\prime}$. Note that $\sqsubseteq$ is a partial order with greatest lower bounds for nonempty sets: $\Pi L=w$ if $w \sqsubseteq w^{\prime \prime}$ for all $w^{\prime \prime} \in L$ and $\forall w^{\prime} .\left(\forall w^{\prime \prime} \in S . w^{\prime} \sqsubseteq w^{\prime \prime}\right) \Rightarrow w^{\prime} \sqsubseteq w . \sqcap L$ is the longest common prefix of all words in $L$.

Definition 8 (Completions). The set of completions $C_{E}(w)$ of $w$ in $E$ is the set of all words in $\mathcal{L} \llbracket E \rrbracket$ that have $w$ as a prefix:

$$
C_{E}(w)=\left\{w^{\prime \prime} \mid w \sqsubseteq w^{\prime \prime} \wedge w^{\prime \prime} \in \mathcal{L} \llbracket E \rrbracket\right\} .
$$

Note that $C_{E}(w)$ may be empty.
Definition 9 (Extension). For nonempty $C_{E}(w)$ the unique extension $\hat{w}_{E}$ of $w$ under $E$ is the longest extension of $w$ with a suffix such that all successful extensions of $w$ to an element of $\mathcal{L} \llbracket E \rrbracket$ are also extensions of $\hat{w}$ :

$$
\hat{w}_{E}=\prod C_{E}(w)
$$

Word $w$ is extended under $E$ if $w=\hat{w}$; otherwise it is unextended.
Extension is a closure operation: $\hat{\hat{w}}=\hat{w}$; in particular, extensions are extended.
Definition 10 (Reduction). For empty $C_{E}(w)$ the unique reduction $\bar{w}_{E}$ of $w$ under $E$ is the longest prefix $w^{\prime}$ of $w$ such that $C_{E}\left(w^{\prime}\right) \neq \emptyset$.

Given parse function $\mathrm{P}_{E}(\cdot): \mathcal{L} \llbracket E \rrbracket \rightarrow \mathcal{B} \llbracket E \rrbracket$ for complete input strings, we can now define what it means for an implementation of it to be optimally streaming:

Definition 11 (Optimally streaming). The optimally streaming function corresponding to $\mathrm{P}_{E}(\cdot)$ is

$$
O_{E}(w)= \begin{cases}\prod\left\{\mathrm{P}_{E}\left(w^{\prime \prime}\right) \mid w^{\prime \prime} \in C_{E}(w)\right\} & \text { if } C_{E}(w) \neq \emptyset \\ \left(\prod O_{E}(\bar{w})\right) \sharp & \text { if } C_{E}(w)=\emptyset\end{cases}
$$

The first condition expresses that after seeing prefix $w$ the function must output all bits that are a common prefix of all bit-coded parse trees of words in $\mathcal{L} \llbracket E \rrbracket$ that $w$ can be extended to. The second condition expresses that as soon as it is clear that a prefix has no extension to an element of $\mathcal{L} \llbracket E \rrbracket$, an indicator $\sharp$ of failure must be emitted, with no further output after that. In this sense $O_{E}$ is optimally streaming: It produces output bits at the semantically earliest possible time during input processing.

It is easy to check that $O_{E}$ is a streaming function:

$$
w \sqsubseteq w^{\prime} \Rightarrow O_{E}(w) \sqsubseteq O_{E}\left(w^{\prime}\right)
$$

The definition has the, at first glance, surprising consequence that $O_{E}$ may output bits for parts of the input it has not even read yet:

Proposition 2. $O_{E}(w)=O_{E}(\hat{w})$
E.g. for $E=(\mathrm{a}+\mathrm{a})(\mathrm{a}+\mathrm{a})$ we have $O_{E}(\epsilon)=00$; that is, $O_{E}$ outputs 00 off the bat, before reading any input symbols, in anticipation of aa being the only possible successful extension. Assume the input is ab. After reading a it does not output anything, and after reading b it outputs $\sharp$ to indicate a failed parse, the total output being 00\#.

## 6 Coverage

Our algorithm is based on simulating aNFAs in lock-step, maintaining a set of partial paths reading the prefix $w$ of the input that has been consumed so far. In order to be optimally streaming, we have to identify partial paths which are guaranteed not be a prefixes of a greedy parse for a word in $C_{E}(w)$.

In this section, we define a coverage relation which our parsing algorithm relies on in order to detect the aforementioned situation. In the following, fix an RE $E$ and its aNFA $M_{E}=\left(\operatorname{State}_{E}, \delta_{E}, q_{E}^{\text {in }}, q_{E}^{\text {fin }}\right)$.

Definition 12 (Coverage). Let $p \in$ State $_{E}$ be a state and $Q \subseteq$ State $_{E}$ a state set. We say that $Q$ covers $p$, written $Q \sqsupseteq p$, iff

$$
\begin{equation*}
\left\{\operatorname{read}(\alpha) \mid q \stackrel{\alpha}{\rightsquigarrow} q^{\text {fin }}, q \in Q\right\} \supseteq\left\{\operatorname{read}(\beta) \mid p \stackrel{\beta}{\rightsquigarrow} q^{\text {fin }}\right\} \tag{1}
\end{equation*}
$$

Coverage can be seen as a slight generalization of language inclusion. That is, if $Q \sqsupseteq p$, then every word suffix read by a path from $p$ to the final state can also be read by a path from one of the states in $Q$ to the final state.

Let $\overline{M_{e}}$ refer to the automaton obtained by reversing the direction of all transitions and swapping the initial and final states. It can easily be verified that if (1) holds for some $Q, p$, then the following property also holds in the reverse automaton $\overline{M_{E}}$ :

$$
\begin{equation*}
\left\{\operatorname{read}(\alpha) \mid q^{\text {in }} \stackrel{\alpha}{\rightsquigarrow} q, q \in Q\right\} \supseteq\left\{\operatorname{read}(\beta) \mid q^{\text {in }} \stackrel{\alpha}{\rightsquigarrow} p\right\} \tag{2}
\end{equation*}
$$

If we consider $D\left(\overline{M_{E}}\right)$, the deterministic automaton generated from $\overline{M_{E}}$, then we see that $(2)$ is satisfied iff

$$
\begin{equation*}
\forall S \in \text { DState }_{\overline{M_{E}}} \cdot p \in S \Rightarrow Q \cap S \neq \emptyset \tag{3}
\end{equation*}
$$

This is true since a DFA state $S$ is reachable by reading a word $w$ in $D\left(\overline{M_{E}}\right)$ iff every $q \in S$ is reachable by reading $w$ in $\overline{M_{E}}$. Since a DFA accepts the same language as the underlying aNFA, this implies that condition (2) must hold iff $Q$ has a non-empty intersection with all DFA states containing $p$.

The equivalence of (1) and (3) gives us a method to decide $\sqsupseteq$ in an aNFA $M$, provided that we have computed $D(\bar{M})$ beforehand. Checking (3) for a particular $Q$ and $p$ can be done by intersecting all states of DState ${\overline{M_{E}}}$ with $Q$, using time $O\left(|Q|\left|\mathrm{DState} \overline{\bar{M}_{E}}\right|\right)=O\left(|Q| 2^{O(m)}\right)$, where $m$ is the size of the RE $E$.

The exponential cost appears to be unavoidable - the problem of deciding coverage is inherently hard to compute:

Proposition 3. The problem of deciding coverage, that is the set $\{(E, Q, p) \mid$ $Q \subseteq$ State $\left._{E} \wedge Q \sqsupseteq p\right\}$, is PSPACE-hard.

Proof. We can reduce regular expression equivalence to coverage: Given regular expressions $E$ and $F$, produce an aNFA $M_{E+F}$ for $E+F$ and observe that $M_{E}$ and $M_{F}$ are subautomata. Now observe that there is a path $q_{E+F}^{\text {in }} \underset{\sim}{\alpha} q_{E}^{\mathrm{fin}}$ (respectively $\left.q_{E+F}^{\text {in }} \stackrel{\beta}{\rightsquigarrow} q_{F}^{\text {fin }}\right)$ in $M_{E+F}$ iff there is a path $q_{E}^{\text {in }} \stackrel{\alpha^{\prime}}{\rightsquigarrow} q_{E}^{\text {fin }}$ with $\operatorname{read}(\alpha)=\operatorname{read}\left(\alpha^{\prime}\right)$ in $M_{E}$ (respectively $q_{F}^{\text {in }} \stackrel{\beta^{\prime}}{\rightsquigarrow} q_{F}^{\text {fin }}$ with $\operatorname{read}(\beta)=\operatorname{read}\left(\beta^{\prime}\right)$ in $M_{F}$ ). Hence, we have $\left\{q_{F}^{\text {in }}\right\} \sqsupseteq q_{E}^{\text {in }}$ in $M_{E+F}$ iff $\mathcal{L} \llbracket E \rrbracket \subseteq \mathcal{L} \llbracket F \rrbracket$. Since regular expression containment is PSPACE-complete [12] this shows that coverage is PSPACE-hard.

Even after having computed a determinized automaton, the decision version of the coverage problem is still NP-complete, which we show by reduction to and from Min-Cover, a well-known NP-complete problem. Let State-Cover refer to the problem of deciding membership for the language $\{(M, D(M), p, k) \mid$ $\exists Q .|Q|=k \wedge p \notin Q \wedge Q \sqsupseteq p$ in $M\}$. Recall that Min-Cover is the problem of deciding membership for the language $\{(X, \mathcal{F}, k)|\exists \mathcal{C} \subseteq \mathcal{F} .|\mathcal{C}|=k \wedge X=\bigcup \mathcal{C}\}$.

Proposition 4. State-Cover is NP-complete.
Proof. State-Cover $\Rightarrow$ Min-Cover: Let $(M, D(M), p, k)$ be given. Define $X=$ $\left\{S \in\right.$ DState $\left._{M} \mid p \in S\right\}$ and $\mathcal{F}=\left\{R_{q} \mid q \in \bigcup X\right\}$ where $R_{q}=\{S \in X \mid q \in S\}$. Then any $k$-sized set cover $\mathcal{C}=\left\{R_{q_{1}}, \ldots, R_{q_{k}}\right\}$ gives a state cover $Q=\left\{q_{1}, \ldots, q_{k}\right\}$ and vice-versa.

Min-Cover $\Rightarrow$ State-Cover: Let $(X, \mathcal{F}, k)$ be given, where $|X|=m$ and $|\mathcal{F}|=n$. Construct an aNFA $M_{X, \mathcal{F}}$ over the alphabet $\Sigma=X \uplus\{\$\}$. Define its states to be the set $\left\{q^{\text {in }}, q^{\text {fin }}, p\right\} \cup\left\{F_{1}, \ldots, F_{n}\right\}$, and for each $F_{i}$, add transitions $F_{i} \xrightarrow{\$} q^{\text {fin }}$ and $q^{\text {in }} \xrightarrow{x_{i j}} F_{i}$ for each $x_{i j} \in F_{i}$. Finally add transitions $p \xrightarrow{\$} q^{\text {fin }}$ and $q^{\text {in }} \xrightarrow{x} p$ for each $x \in X$.

Observe that $D\left(M_{X, \mathcal{F}}\right)$ will have states $\left\{\left\{q^{\text {in }}\right\},\left\{q^{\text {fin }}\right\}\right\} \cup\left\{S_{x} \mid x \in X\right\}$ where $S_{x}=\{F \in \mathcal{F} \mid x \in F\} \cup\{p\}$, and $\Delta\left(\left\{q^{\text {in }}\right\}, x\right)=S_{x}$. Also, the time to compute
$D\left(M_{X, \mathcal{F}}\right)$ is bounded by $O(|X||\mathcal{F}|)$. Then any $k$-sized state cover $Q=\left\{F_{1}, \ldots, F_{k}\right\}$ is also a set cover.

## 7 Algorithm

Our parsing algorithm produces a bit-coded parse tree from an input string $w$ for a given RE $E$. We will simulate $M_{E}$ in lock-step, reading a symbol from $w$ in each step. The simulation maintains a set of all partial paths that read the prefix of $w$ that has been consumed so far; there are always only finitely many paths to consider, since we restrict ourselves to paths without non-productive loops. When a path reaches a non-deterministic choice, it will "fork" into two paths with the same prefix. Thus, the path set can be represented as a tree of states, where the root is the initial state, the edges are transitions between states, and the leaves are the reachable states.

Definition 13 (Path trees). A path tree is a rooted, ordered, binary tree with internal nodes of outdegrees 1 or 2 . Nodes are labeled by aNFA-states and edges by $\Gamma=\Sigma \cup\{0,1\} \cup\{\epsilon\}$. Binary nodes have a pair of 0 - and 1-labeled edges (in this order only), respectively.

We use the following notation:

- $\operatorname{root}(T)$ is the root node of path tree $T$.
- path $(n, c)$ is the path from $n$ to $c$, where $c$ is a descendant of $n$.
$-\operatorname{init}(T)$ is the path from the root to the first binary node reachable or to the unique leaf of $T$ if it has no binary node.
- leaves $(T)$ is the ordered list of leaf nodes.
- Trempty ${ }^{\text {is the empty tree. }}$

As a notational convenience, the tree with a root node labeled $q$ and no children is written $q\langle\cdot\rangle$, where $q$ is an aNFA-state. Similarly, a tree with a root labeled $q$ with children $l$ and $r$ is written $q\langle 0: l, 1: r\rangle$, where $q$ is an aNFA-state and $l$ and $r$ are path trees and the edges from $q$ to $l$ and $r$ are labeled 0 and 1 , respectively. Unary nodes are labelled by $\Sigma \cup\{\epsilon\}$ and are written $q\langle\ell: c\rangle$, denoting a tree rooted at $q$ with only one $\ell$-labelled child $c$.

In the following we shall use $T_{w}$ to refer to a path tree created after processing input word $w$ and $T$ to refer to path trees in general, where the input string giving rise to the tree is irrelevant.

Definition 14 (Path tree invariant). Let $T_{w}$ be a path tree and $w$ a word. Define $I\left(T_{w}\right)$ as the proposition that all of the following hold:
(i) The leaves $\left(T_{w}\right)$ have pairwise distinct node labels; all labels are symbol sources, that is states with a single symbol transition, or the accept state.
(ii) All paths from the root to a leaf read w: $\forall n \in \operatorname{leaves}\left(T_{w}\right) . \operatorname{read}\left(\operatorname{path}\left(\operatorname{root}\left(T_{w}\right), n\right)\right)=w$.
(iii) For each leaf $n \in \operatorname{leaves}\left(T_{w}\right)$ there exists $w^{\prime \prime} \in C_{E}(w)$ such that the bit-coded parse of $w^{\prime \prime}$ starts with write $\left(\operatorname{path}\left(\operatorname{root}\left(T_{w}\right), n\right)\right)$.

```
Algorithm 1 Optimally streaming greedy regular expression parsing algorithm.
Require: An aNFA \(M\), a coverage relation \(\sqsupseteq\), and an input stream \(S\).
Ensure: The greedy leftmost parse tree, emitted in an optimally-streaming fashion.
    function \(\operatorname{Stream-Parse}(M, \sqsupseteq, S)\)
        \(w \leftarrow \epsilon\)
        \(\left(T_{\epsilon}, \ldots\right) \leftarrow \operatorname{CLOSURE}\left(M, \emptyset, q^{\text {in }}\right) \triangleright\) Initialize path tree as the output of CLOSURE
        while \(S\) has another input symbol \(a\) do
            if \(C_{E}(w a)=\emptyset\) then
                return write(init \(\left.\left(T_{w}\right)\right)\) followed by \(\sharp\) and exit.
            \(T_{w a} \leftarrow\) Establish-Invariant \(\left(T_{w}, a, \sqsupseteq\right)\)
            Output new bits on the path to the first binary node in \(T_{w a}\), if any.
            \(w \leftarrow w a\)
        if \(q^{\text {fin }} \in \operatorname{leaves}\left(T_{w}\right)\) then
            return write \(\left(\operatorname{path}\left(\operatorname{root}\left(T_{w}\right), q^{\text {fin }}\right)\right)\)
        else
            return write \(\left(\operatorname{init}\left(T_{w}\right)\right)\) followed by \(\sharp\)
```

```
Algorithm 2 Establishing invariant \(I\left(T_{w a}\right)\)
Require: A path tree \(T_{w}\) satisfying invariant \(I\left(T_{w}\right)\), a character \(a\), and a coverage
    relation \(\sqsupseteq\).
Ensure: A path tree \(T_{w a}\) satisfying invariant \(I\left(T_{w a}\right)\).
    function Establish-Invariant \(\left(T_{w}, a, \sqsupseteq\right)\)
        Remove leaves from \(T_{w}\) that do not have a transition on \(a\).
        Extend \(T_{w}\) to \(T_{w a}\) by following all \(a\)-transitions.
        for each leaf \(n\) in \(T_{w a}\) do
            \(\left(T^{\prime}, \ldots\right) \leftarrow \operatorname{CLOSURE}(M, \emptyset, n)\).
            Replace the leaf \(n\) with the tree \(T^{\prime}\) in \(T_{w a}\).
        return \(\operatorname{PRUNE}\left(T_{w a}, \sqsupseteq\right)\)
```

(iv) For each $w^{\prime \prime} \in C_{E}(w)$ there exists $n \in$ leaves $\left(T_{w}\right)$ such that the bit-coded parse of $w^{\prime \prime}$ starts with write $\left(\operatorname{path}\left(\operatorname{root}\left(T_{w}\right), n\right)\right)$.

The path tree invariant is maintained by Algorithm 2 ; line 2 establishes parti line 3 establishes part iif and lines $4-7$ establishes part iii and iv

Theorem 1 (Optimal streaming property). Assume extended $w, C_{E}(w) \neq$ $\emptyset$. Consider the path tree $T_{w}$ after reading $w$ upon entry into the while-loop of the algorithm in Algorithm 1. Then write $\left(\operatorname{init}\left(T_{w}\right)\right)=O_{E}(w)$.

In other words, the initial path from the root of $T_{w}$ to the first binary node in $T_{w}$ is the longest common prefix of all paths accepting an extension of $w$. Operationally, whenever that path gets longer by pruning branches, we output the bits on the extension.

Proof. Assume $w$ extended, that is $w=\hat{w}$; assume $C_{E}(w) \neq \emptyset$, that is there exists $w^{\prime \prime}$ such that $w \sqsubseteq w^{\prime \prime}$ and $w^{\prime \prime} \in \mathcal{L} \llbracket E \rrbracket$.

```
Algorithm 3 Pruning algorithm.
Require: A path tree \(T\) and a covering relation \(\sqsupseteq\).
Ensure: A pruned path tree \(T^{\prime}\) where all leaves are alive.
    function \(\operatorname{PRUNE}(T, \sqsupseteq)\)
        for each \(l\) in reverse(leaves \((T))\) do
            \(S \leftarrow\{n \mid n\) comes before \(l\) in leaves \((T)\}\)
            if \(S \sqsupseteq l\) then
                    \(p \leftarrow \operatorname{parent}(l)\)
                    Delete \(l\) from \(T\)
                    \(T \leftarrow \operatorname{CuT}(T, p)\)
        return \(T\)
    function \(\operatorname{CuT}(T, n) \quad\) Cuts a chain of 1-ary nodes.
        if \(\mid\) children \((n) \mid=0\) then
            \(p \leftarrow \operatorname{parent}(n)\)
            \(T^{\prime} \leftarrow T\) with \(n\) removed
            return \(\operatorname{CuT}\left(T^{\prime}, p\right)\)
        else
            return \(T\)
```

Claim: $\mid$ leaves $\left(T_{w}\right) \mid \geq 2$ or the unique node in leaves $\left(T_{w}\right)$ is labeled by the accept state. Proof of claim: Assume otherwise, that is $\mid$ leaves $\left(T_{w}\right) \mid=1$, but its node is not the accept state. By i of $I\left(T_{w}\right)$, this means the node must have a symbol transition on some symbol $a$. In this case, all accepting paths $C_{E}(w a)=$ $C_{E}(w)$ and thus $\hat{w}=\hat{w a}$; in particular $\hat{w} \neq w$, which, however, is a contradiction to the assumption that $w$ is extended.

This means we have two cases. The case $\mid$ leaves $\left(T_{w}\right) \mid=1$ with the sole node being labeled by the accept state is easy: It spells a single path from initial to accept state. By ii and iii of $I\left(T_{w}\right)$ we have that that path is correct for $w$. By iv and since the accept state has no outgoing transitions, we have $C_{E}(w)=\{w\}$, and the theorem follows for this case.

Let us consider the case $\mid$ leaves $\left(T_{w}\right) \mid \geq 2$ then. Recall that $C_{E}(w) \neq \emptyset$ by assumption. By iv of $I\left(T_{w}\right)$ the accepting path of every $w^{\prime \prime} \in C_{E}(w)$ starts with path $\left(\operatorname{root}\left(T_{w}\right), n\right)$ for some $n \in$ leaves $\left(T_{w}\right)$, and by iii each path from the root to a leaf is the start of some accept path. Since $\mid$ leaves $\left(T_{w}\right) \mid \geq 2$ we know that there exists a binary node in $T_{w}$. Consider the first on the path from the root to a leaf. It has both 0 - and 1-labeled out-edges. Thus the longest common prefix of $\left\{\operatorname{write}(p) \mid n \in \operatorname{leaves}\left(T_{w}\right), p \in \operatorname{path}\left(\operatorname{root}\left(T_{w}\right), n\right)\right\}$ is write $\left(\operatorname{init}\left(T_{w}\right)\right)$, the bits on the initial path from the root of $T_{w}$ to its first binary node.

The algorithm, as given, is only optimally streaming for extended prefixes. It can be made to work for all prefixes by enclosing it in an outer loop that for each prefix $w$ computes $\hat{w}$ and calls the given algorithm with $\hat{w}$. The outer loop then checks that subsequent symbols match until $\hat{w}$ is reached. By Proposition 2 the resulting algorithm gives the right result for all input prefixes, not only extended ones.

```
Algorithm 4 -closure with path tree construction.
Require: An aNFA \(M\), a set of visited states \(V\), and a state \(q\)
Ensure: A path tree \(T\) and a set of visited states \(V^{\prime}\)
    function Closure \((M, V, q)\)
        if \(q \xrightarrow{0} q_{l}\) and \(q \xrightarrow{1} q_{r}\) then
            \(\left(T^{l}, V_{l}\right) \leftarrow \operatorname{CLOSURE}\left(M, V \cup\{q\}, q_{l}\right) \quad \triangleright\) Try left option first.
            \(\left(T^{r}, V_{l r}\right) \leftarrow \operatorname{Closure}\left(M, V_{l}, q_{r}\right) \quad \triangleright\) Use \(V_{l}\) to skip already-visited nodes.
            return \(\left(q\left\langle T^{l}: T^{r}\right\rangle, V_{l r}\right)\)
        if \(q \xrightarrow{\epsilon} p\) then
            if \(p \in V\) then \(\triangleright\) Stop loops.
                return ( \(\mathrm{Tr}_{\text {empty }}, V\) )
            else
                \(\left(T^{\prime}, V^{\prime}\right) \leftarrow \operatorname{Closure}(M, V \cup\{q\}, p)\)
                return \(\left(q\left\langle\epsilon: T^{\prime}\right\rangle, V^{\prime}\right)\)
        else \(\quad \triangleright q\) is a symbol source or the final state.
            return \((q\langle\cdot\rangle, V)\)
```

Theorem 2. The optimally streaming algorithm can be implemented to run in time $O\left(2^{m \log m}+m n\right)$, where $m=|E|$ and $n=|w|$.

Proof (Sketch). As shown in Section 6, we can decide coverage in time $O\left(m 2^{O(m)}\right)$. The set of ordered lists leaves $(T)$ for any $T$ reachable from the initial state can be precomputed and covered states marked in it. (This requires unit-cost random access since there are $O\left(2^{m \log m}\right)$ such lists.) The $\epsilon$-closure can be computed in time $O(m)$ for each input symbol, and pruning can be amortized over $\epsilon$-closure computation by charging each edge removed to its addition to a tree path.

For fixed regular expression $E$ this is linear time in $n$ and thus asymptotically optimal. An exponential in $m$ as an additive preprocessing cost appears practically unavoidable since we require the coverage relation, which is inherently hard to compute (Proposition 3).

## 8 Example

Consider the RE (aaa +aa$)^{\star}$. A simplified version of its symmetric position automaton is shown in Figure 2 . The following two observations are requirements for an earliest parse of this expression:

- After one a has been read, the algorithm must output a 0 to indicate that one iteration of the Kleene star has been made, but:
- five consecutive as determine that the leftmost possibility in the Kleene star choice was taken, meaning that the first three as are consumed in that branch.

The first point can be seen by noting that any parse of a non-zero number of as must follow a path through the Kleene star. This guarantees that if a
successful parse is eventually performed, it must be the case that at least one iteration was made.

The second point can be seen by considering the situation where only four input as have been read: It is not known whether these are the only four or more input symbols in the stream. In the former case, the correct (and only) parse is two iterations with the right alternative, but in the latter case, the first three symbols are consumed in the left branch instead.

These observations correspond intuitively to what "earliest" parsing is; as soon as it is impossible that an iteration was not made, a bit indicating this fact is emitted, and as soon as the first three symbols must have been parsed in the left alternative, this fact is output. Furthermore, a 0-bit is emitted to indicate that (at least) another iteration is performed.

Figure 2 shows the evolution of the path tree during execution with the RE $(a a a+a a)^{\star}$ on the input aaaaa.

By similar reasoning as above, after five as it is safe to commit to the left alternative after every third a. Hence, for the inputs aaaaa(aaa) ${ }^{n}$, aaaaa(aaa) ${ }^{n}$ a, and aaaaa(aaa) ${ }^{n}$ aa the "commit points" are placed as follows (• indicate end-ofinput):


Complex coverage. The previous example does not exhibit any non-trivial coverage, i.e., situations where a state $n$ is covered by $k>1$ other states. One can construct an expression that contains non-trivial coverage relations by observing that if each symbol source $s$ in the aNFA is associated with the RE representing the language recognized from $s$, coverage can be expressed as a set of (in)equations in Kleene algebra. Thus, the coverage $\left\{n_{0}, n_{1}\right\} \sqsupseteq n$ becomes $R E\left(n_{0}\right)+R E\left(n_{1}\right) \geq R E(n)$ in KA, where $R E(\cdot)$ is the function that yields the RE from a symbol source in an aNFA.

Any expression of the form $x_{1} z y_{1}+x_{2} z y_{2}+x_{3} z\left(y_{1}+y_{2}\right)$ satisfies the property that two subterms cover a third. If the coverage is to play a role in the algorithm, however, the languages denoted by $x_{1}$ and $x_{2}$ must not subsume that of $x_{3}$, otherwise the part starting with $x_{3}$ would never play a role due to greedy leftmost disambiguation.

Choose $x_{1}=x_{2}=(\mathrm{aa})^{\star}, x_{3}=\mathrm{a}^{\star}, y_{1}=\mathrm{a}$, and $y_{2}=\mathrm{b}$. Figure 3 shows the expression

$$
(\mathrm{aa})^{\star} \mathrm{za}+\mathrm{aa}^{\star} \mathrm{zb}+\mathrm{a}^{\star} z \mathrm{a}+\mathrm{b}=(\mathrm{aa})^{\star}(\mathrm{za}+\mathrm{zb})+a^{\star} \mathrm{z}(\mathrm{a}+\mathrm{b})
$$

The earliest point where any bits can be output is when the $z$ is reached. Then it becomes known whether there was an even or odd number of as. Due to the coverage $\{8,13\} \sqsupseteq 20$ state 20 is pruned away on the input aazb, thereby causing the path tree to have a large trunk that can be output.


Fig. 2: Example run of the algorithm on the regular expression $E=(\mathrm{aaa}+\mathrm{aa})^{\star}$ and the input string aaaaa. The dashed edges represent the partial parse trees that can be emitted: thus, after one a we can emit a 0 , and after five as we can emit 00 because the bottom "leg" of the tree has been removed in the pruning step. The automaton for $E$ and its associated minimal covering relation are shown in the inset.


Fig. 3: Example run of the algorithm on $E=(a \mathrm{a})^{\star}(\mathrm{za}+\mathrm{zb})+\mathrm{a}^{\star} \mathrm{z}(\mathrm{a}+\mathrm{b})$. Note that state 20 is covered by the combination of states 8 and 13 . The earliest time the algorithm can do a commit is when a $\mathbf{z}$ is encountered, which decides whether there is an even or odd number of as. The topmost figure shows the evolution of the path tree on the input aaazb. There is a long "trunk" from state 1 to state 21 after reading $z$, as the rest of the branches have been pruned (not shown). The desired output, corresponding to taking the rightmost option in the sum, can be read off the labels on the edges. Likewise in the second figure, we see that if the $\mathbf{z}$ comes after an even number of as, a binary-node-free path from 1 to 7 emerges. Due to the cover $\{8,13\} \sqsupseteq 20$, the branch starting from 20 is not expanded further, even though there could be a z-transition on it. This is indicated with $\downarrow$. Overall, the resulting parse tree corresponds to the leftmost option in the sum.

CSV files. The expression $\left((\mathrm{a}+\mathrm{b})^{\star}\left(;(\mathrm{a}+\mathrm{b})^{\star}\right)^{\star} \mathrm{n}\right)^{\star}$ defines the format of a simple semicolon-delimited data format, with data consisting of words over $\{\mathrm{a}, \mathrm{b}\}$ and rows separated by the newline character, n. Our algorithm emits the partial parse trees after each letter has been parsed, as illustrated on the example input below:

Due to the star-height of three, many widespread implementations would not be able to meaningfully handle this expression using only the RE engine. Capturing groups under Kleene stars return either the first or last match, but not a list of matches-and certainly not a list of lists of matches! Hence, if using an implementation like Perl's [16], one is forced to rewrite the expression by removing the iteration in the outer Kleene star and reintroduce it as a looping construct in Perl.

## 9 Related and Future Work

Parsing regular expressions is not new [6|5|3|10|14], and streaming parsing of XML documents has been investigated for more than a decade in the context of XQuERY and XPath-see, e.g., [2|7|17]. However, streaming regular expression parsing appears to be new.

In earlier work [6] we described a compact "lean log" format for storing intermediate information required for two-phase regular expression parsing. The algorithm presented here may degenerate to two passes, but requires often just one pass in the sense being effectively streaming, using only $O(m)$ work space, independent of $n$. The preprocessing of the regular expression and the intermediate data structure durig input string processing are more complex, however. It may be possible to merge the two approaches using a tree of lean log frames with associated counters, observing that edges in the path tree that are not labeled 0 or 1 are redundant. This is future work.

Acknowledgements. This work is supported by the Danish Independent Research Council under Project "Kleene Meets Church: Regular Expressions and Types". We would like to thank the anonymous referees for their helpful criticisms.

## References

1. Brüggemann-Klein, A., Wood, D.: One-unambiguous regular languages. Information and computation $140(2)$, 229-253 (1998), http://dx.doi.org/10.1006/inco 1997.2688
2. Debarbieux, D., Gauwin, O., Niehren, J., Sebastian, T., Zergaoui, M.: Early nested word automata for XPath query answering on XML streams. In: Konstantinidis, S. (ed.) Implementation and Application of Automata, Lecture Notes in Computer Science, vol. 7982, pp. 292-305. Springer Berlin Heidelberg (2013)
3. Dubé, D., Feeley, M.: Efficiently building a parse tree from a regular expression. Acta Informatica $37(2), 121-144$ (2000), http://dx.doi.org/10.1007/s002360000037
4. Fowler, G.: An interpretation of the POSIX regex standard. http://www2.research att.com/~astopen/testregex/re-interpretation.html (January 2003)
5. Frisch, A., Cardelli, L.: Greedy regular expression matching. In: Proc. 31st International Colloquium on Automata, Languages and Programming (ICALP). Lecture Notes in Computer Science (LNCS), vol. 3142, pp. 618-629. Springer, Turku, Finland (July 2004), http://dx.doi.org/10.1007/978-3-540-27836-8_53
6. Grathwohl, N.B.B., Henglein, F., Nielsen, L., Rasmussen, U.T.: Two-pass greedy regular expression parsing. In: Proc. 18th International Conference on Implementation and Application of Automata (CIAA), Lecture Notes in Computer Science (LNCS), vol. 7982, pp. 60-71. Springer (July 2013), http://dx.doi.org/10.1007/ 978-3-642-39274-0_7
7. Gupta, A.K., Suciu, D.: Stream processing of XPath queries with predicates. In: Proc. 2003 ACM SIGMOD International Conference on Management of Data. pp. 419-430. SIGMOD '03, ACM, New York, NY, USA (2003), http://dx.doi.org/ 10.1145/872757.872809
8. IEEE Computer Society: Standard for Information Technology - Portable Operating System Interface (POSIX), Base Specifications, Issue 7. IEEE (2008), http://dx doi.org/10.1109/IEEESTD.2008.4694976, IEEE Std 1003.1
9. Kearns, S.: Extending regular expressions with context operators and parse extraction. Software - Practice and Experience 21(8), 787-804 (1991), http://dx.doi org/10.1002/spe. 4380210803
10. Nielsen, L., Henglein, F.: Bit-coded regular expression parsing. In: Proc. 5th Int'l Conf. on Language and Automata Theory and Applications (LATA). pp. 402413. Lecture Notes in Computer Science (LNCS), Springer (May 2011), http: //dx.doi.org/10.1007/978-3-642-21254-3_32
11. Okui, S., Suzuki, T.: Disambiguation in regular expression matching via position automata with augmented transitions. In: Domaratzki, M., Salomaa, K. (eds.) Implementation and Application of Automata, Lecture Notes in Computer Science, vol. 6482, pp. 231-240. Springer Berlin Heidelberg (2011), http://dx.doi.org/10 1007/978-3-642-18098-9_25
12. Stockmeyer, L.J., Meyer, A.R.: Word problems requiring exponential time (preliminary report). In: Proc. Fifth Annual ACM Symposium on Theory of Computing. pp. 1-9. ACM (1973)
13. Sulzmann, M., Lu, K.Z.M.: Regular expression sub-matching using partial derivatives. In: Proc. 14th symposium on Principles and practice of declarative programming. pp. 79-90. PPDP '12, ACM, New York, NY, USA (2012), http: //dx.doi.org/10.1145/2370776.2370788
14. Sulzmann, M., Lu, K.Z.M.: POSIX regular expression parsing with derivatives. In: Proc. 12th International Symposium on Functional and Logic Programming. FLOPS '14, Kanazawa, Japan (June 2014), http://dx.doi.org/10.1007/ 978-3-319-07151-0_13
15. Thompson, K.: Programming techniques: Regular expression search algorithm. Commun. ACM 11(6), 419-422 (1968), http://dx.doi.org/10.1145/363347 363387
16. Wall, L., Christiansen, T., Orwant, J.: Programming Perl. O'Reilly Media, Incorporated (2000)
17. Wu, X., Theodoratos, D.: A survey on XML streaming evaluation techniques. The VLDB Journal 22(2), 177-202 (Apr 2013), http://dx.doi.org/10.1007/ s00778-012-0281-y

[^0]:    * This work has been partially supported by The Danish Council for Independent Research under Project 11-106278, "Kleene Meets Church: Regular Expressions and Types". The order of authors is insignificant.

