

From Finite Automata to Regular Expressions and Back—A Summary on Descriptive Complexity

Hermann Gruber

knowledgepark AG, Leonrodstr. 68,
80636 München, Germany

hermann.gruber@knowledgepark-ag.de

Markus Holzer

Institut für Informatik, Universität Giessen,
Arndtstr. 2, 35392 Giessen, Germany

holzer@informatik.uni-giessen.de

The equivalence of finite automata and regular expressions dates back to the seminal paper of Kleene on events in nerve nets and finite automata from 1956. In the present paper we tour a fragment of the literature and summarize results on upper and lower bounds on the conversion of finite automata to regular expressions and *vice versa*. We also briefly recall the known bounds for the removal of spontaneous transitions (ε -transitions) on non- ε -free nondeterministic devices. Moreover, we report on recent results on the average case descriptive complexity bounds for the conversion of regular expressions to finite automata and brand new developments on the state elimination algorithm that converts finite automata to regular expressions.

1 Introduction

There is a vast literature documenting the importance of the notion of finite automata and regular expressions as an enormously valuable concept in theoretical computer science and applications. It is well known that these two formalisms are equivalent, and in almost all monographs on automata and formal languages one finds appropriate constructions for the conversion of finite automata to equivalent regular expressions and back. Regular expressions, introduced by Kleene [68], are well suited for human users and therefore are often used as interfaces to specify certain patterns or languages. For example, in the widely available programming environment UNIX, regular(-like) expressions can be found in legion of software tools like, e.g., `awk`, `ed`, `emacs`, `egrep`, `lex`, `sed`, `vi`, etc., to mention a few of them. On the other hand, automata [94] immediately translate to efficient data structures, and are very well suited for programming tasks. This naturally raises the interest in conversions among these two different notions. Our tour on the subject covers some (recent) results in the fields of descriptive and computational complexity. During the last decade descriptive aspects on finite automata and regular expressions formed an extremely vivid area of research. For recent surveys on descriptive complexity issues of finite automata and regular expressions we refer to, for example, [39, 56, 57, 58, 59, 60, 103]. This was not only triggered by appropriate conferences and workshops on that subject, but also by the availability of mathematical tools and the influence of empirical studies. For obvious reasons, this survey lacks completeness, as finite automata and regular expressions fall short of exhausting the large number of related problems considered in the literature. We give a view of what constitutes, in our opinion, the most interesting recent links to the problem area under consideration.

Before we start our tour some definitions are in order. First of all, our nomenclature of finite automata is as follows: a *nondeterministic finite automaton with ε -transitions* (ε -NFA) is a quintuple $A = (Q, \Sigma, \delta, q_0, F)$, where Q is the finite set of *states*, Σ is the finite set of *input symbols*, $q_0 \in Q$ is the *initial state*, $F \subseteq Q$ is the set of *accepting states*, and $\delta : Q \times (\Sigma \cup \{\varepsilon\}) \rightarrow 2^Q$ is the *transition function*. If a finite automaton has no ε -transitions, i.e., the transition function is restricted to $\delta : Q \times \Sigma \rightarrow 2^Q$, then we simply speak of a *nondeterministic finite automaton* (NFA). Moreover, a nondeterministic finite

automaton is *deterministic* (DFA) if and only if $|\delta(q, a)| = 1$, for all states $q \in Q$ and letters $a \in \Sigma$. The *language accepted* by the finite automaton A is defined as $L(A) = \{w \in \Sigma^* \mid \delta(q_0, w) \cap F \neq \emptyset\}$, where the transition function is recursively extended to $\delta : Q \times \Sigma^* \rightarrow 2^Q$. Second, we turn to the definition of regular expressions: the *regular expressions* over an alphabet Σ and the languages they describe are defined inductively in the usual way:¹ \emptyset , ε , and every letter a with $a \in \Sigma$ is a regular expression, and when s and t are regular expressions, then $(s+t)$, $(s \cdot t)$, and $(s)^*$ are also regular expressions. The language defined by a regular expression r , denoted by $L(r)$, is defined as follows: $L(\emptyset) = \emptyset$, $L(\varepsilon) = \{\varepsilon\}$, $L(a) = \{a\}$, $L(s+t) = L(s) \cup L(t)$, $L(s \cdot t) = L(s) \cdot L(t)$, and $L(s^*) = L(s)^*$. For further details on finite automata and regular expressions we refer to, e.g., [61].

We start our tour on the subject with the question on the appropriate measure for finite automata and regular expressions. We discuss this topic in detail in Section 2. There we also concentrate on two specific measures: on star height for regular expressions and cycle rank for the automaton side. By Eggen’s theorem [26] both measures are related to each other. Recent developments, in particular on the conversion from finite automata to regular expressions, utilize this connection to prove upper and lower bounds. Then in Section 3 we take a closer on the conversion from regular expressions to equivalent finite automata. We recall the most prominent conversion algorithms such as Thompson’s construction and its optimized version the follow automaton, the position or Glushkov automaton, and conversion by computations of the (partial-)derivatives. We summarize the known relations on these devices, which were mostly found during the last decade. Significant differences on these constructions are pointed out and the presented developments on lower bound and upper bound results enlighten the efficiency of these algorithms. Some of the bounds are sensitive to the size of the alphabet. Besides worst case descriptive complexity results on the synthesis problem of finite automata from regular expressions, we also list some recent results on the average case complexity of the transformation of regular expressions to finite automata. Finally, in Section 4 we consider the converse transformation. Again, we summarize some of the few conversion techniques, but then stick in more detail to the so-called state elimination technique. The reason for that is, that in [97], it was shown that almost all conversion methods can be recast as variants of the state elimination technique. Here, the ordering in which the states are eliminated can largely affect the size of the regular expression corresponding to the given finite automaton. We survey some heuristics that have been proposed for this goal. For appropriate choices of the ordering, nontrivial upper bounds on regular expression size can be proved. By looking at the transition structure of the NFA, results from graph theory can help in obtaining shorter expressions. There we try to illustrate the key insights with the aid of examples, thereby avoiding the need for a deeper dive into graph theoretic concepts. We also explain the technique by which the recent lower bounds on regular expression size were obtained. In this part, the known upper and lower bounds match only in the sense that we can identify the rough order of magnitude. So we observe an interesting tension between algorithms with provable performance guarantees, other heuristics that are observed to behave better in experiments, and finally some lower bounds, which seize the expectations that we may have on practical algorithms.

2 Measures on Finite Automata and Regular Expressions

What can be said about the proper measure on finite automata and regular expressions? For finite automata there are two commonly accepted measures, namely the number of states and the number of

¹For convenience, parentheses in regular expressions are sometimes omitted and the concatenation is simply written as juxtaposition. The priority of operators is specified in the usual fashion: concatenation is performed before union, and star before both product and union.

transitions. The measure sc (nsc , respectively) counts the number of states of a deterministic (nondeterministic, respectively) finite automaton and tc (ntc , respectively) does the same for the number of transitions for the appropriate devices. Moreover, nsc_ε (ntc_ε , respectively) gives the number of states (transitions, respectively) in an ε -NFA. The following relations between these measures are well known—see also [84, 86, 94].

Theorem 1 *Let $L \subseteq \Sigma^*$ be a regular language. Then*

1. $nsc_\varepsilon(L) = nsc(L) \leq sc(L) \leq 2^{nsc(L)}$ and $tc(L) = |\Sigma| \cdot sc(L)$ and
2. $nsc(L) - 1 \leq ntc_\varepsilon(L) \leq ntc(L) \leq |\Sigma| \cdot (nsc(L))^2$,

where $sc(L)$, $tc(L)$ ($nsc(L)$, $ntc(L)$, respectively) refers to the minimum (nsc , ntc , respectively) among all DFAs (NFAs, respectively) accepting L . Similarly, $nsc_\varepsilon(L)$ ($ntc_\varepsilon(L)$, respectively) is the minimum nsc_ε (ntc_ε , respectively) among all ε -NFAs for the language L .

As it is defined above, deterministic transition complexity is not an interesting measure by itself, because it is directly related to sc , the deterministic state complexity. But the picture changes when deterministic transition complexity is defined in terms of partial DFAs. Here, a *partial DFA* is an NFA which transition function δ satisfies $|\delta(q, a)| \leq 1$, for all states $q \in Q$ and all alphabet symbols $a \in \Sigma$. A partial DFA cannot save more than one state compared to an ordinary DFA, but it can save a considerable number of transitions in some cases. This phenomenon is studied, e.g., in [30, 75, 76]. Further measures for the complexity of finite automata, in particular measures related to unambiguity and limited nondeterminism, can be found in [39, 40, 41, 59, 63, 70, 71, 72, 92, 93, 95].

Now let us come to measures on regular expressions. While there are the two commonly accepted measures for finite automata, there is no general agreement in the literature about the proper measure for regular expressions. We summarize some important ones: the measure $size$ is defined to be the total number of symbols (including \emptyset , ε , symbols from alphabet Σ , all operation symbols, and parentheses) of a completely bracketed regular expression (for example, used in [2], where it is called $length$). Another measure related to the reverse polish notation of a regular expression is rpn , which gives the number of nodes in the syntax tree of the expressions (parentheses are not counted). This measure is equal to the length of a (parenthesis-free) expression in post-fix notation [2]. The alphabetic width $awidth$ is the total number of alphabetic symbols from Σ (counted with multiplicity) [27, 83]. Relations between these measures have been studied, e.g., in [27, 28, 42, 66].

Theorem 2 *Let $L \subseteq \Sigma^*$ be a regular language. Then*

1. $size(L) \leq 3 \cdot rpn(L)$ and $size(L) \leq 8 \cdot awidth(L) - 3$,
2. $awidth(L) \leq \frac{1}{2} \cdot (size(L) + 1)$ and $awidth(L) \leq \frac{1}{2} \cdot (rpn(L) + 1)$, and
3. $rpn(L) \leq \frac{1}{2} \cdot (size(L) + 1)$ and $rpn(L) \leq 4 \cdot awidth(L) - 1$,

where $size(L)$ ($rpn(L)$, $awidth(L)$, respectively) refers to the minimum $size$ (rpn , $awidth$, respectively) among all regular expressions denoting L .

Further measures for the complexity of regular expressions can be found in [8, 27, 28, 49]. To our knowledge, these latter measures received far less attention to date.

In the remainder of this section we concentrate on two important measures on regular expression and finite automata that at first glance do not seem to be related to each other: *star height* and *cycle rank* or *loop complexity*. Both measures are very important, in particular, for the conversion of finite automata to regular expressions and for proving lower bound results on the latter. Intuitively, the star

height of an expression measures the nesting depth of Kleene-star operations. More precisely, for a regular expression, the *star height* is inductively defined by

$$\begin{aligned} \text{height}(\emptyset) &= \text{height}(\varepsilon) = \text{height}(a) = 0, \\ \text{height}(s+t) &= \text{height}(s \cdot t) = \max(\text{height}(s), \text{height}(t)), \end{aligned}$$

and

$$\text{height}(s^*) = 1 + \text{height}(s).$$

The star height of a regular language L , denoted by $\text{height}(L)$ is then defined as the minimum star height among all regular expressions describing L . The seminal work dealing with the star height of regular expressions [26] established a relation between the theory of regular languages and the theory of digraphs. The *cycle rank*, or *loop complexity*, of a digraph D is defined inductively by the following rules: (i) the cycle rank of an acyclic digraph is zero, (ii) cycle rank of a strongly connected component (SCC) of the digraph with at least one arc is 1 plus the minimum cycle rank among the digraphs obtainable from D by deleting a vertex, and (iii) the cycle rank of a digraph with multiple SCCs equals the maximum cycle rank among the sub-digraphs induced by these components. So, roughly speaking, the cycle rank of a digraph is large if the cycle structure of the digraph is intricate and highly connected. The following relation between cycle rank of automata and star height of regular languages became known as *Eggan's Theorem* [26, 97]:

Theorem 3 *The star height of a regular language L equals the minimum cycle rank among all ε -NFAs accepting L .*

An apparent difficulty with applying Eggan's Theorem is that the minimum is taken over infinitely many automata, and the cycle rank of the minimum DFA for the language does not always attain that minimum. That makes the star height a very intricate property of regular languages. Indeed, the decidability status of the star height problem was open for more than two decades, until a very difficult algorithm was given in [54]. For recent progress on algorithms for the star height problem, the reader is referred to [67]. From the above it is immediate that $\text{height}(L) \leq \text{nsc}(L)$. If the language is given as a regular expression, a result from [43] tells us a much sweeter truth:

Lemma 4 *Let $L \subseteq \Sigma^*$ be a regular language with alphabetic width n . Then $\text{height}(L) \leq 3 \log(n+1)$.*

The idea behind the proof of this lemma is that we can convert a regular expression into a ε -NFA of similar size. The cycle structure of that automaton is well-behaved; and thus its cycle rank is low compared to the size of the automaton. Then Eggan's Theorem is used to convert the automaton back into a regular expression of low star height.

We return to the relationship between required size and star height of regular expressions later on. Now let us turn our attention to the conversion of regular expressions into equivalent finite automata.

3 From Regular Expressions to Finite Automata

The conversion of regular expressions into small finite automata has been intensively studied for more than half a century. Basically the algorithms can be classified according to whether the output is an ε -NFA, NFA, or even a DFA. In principle one can distinguish between the following three major construction schemes and variants thereof:

1. *Thompson's construction* [100] and optimized versions, such as the *follow automaton* [66, 91],

2. construction of the *position automaton*, or *Glushkov automaton* [38, 83], and
3. computation of the (*partial derivative automaton*) [4, 14].

Further automata constructions from regular expressions can be found in, e.g., [6, 13, 19, 65, 31, 102]. We briefly explain some of these approaches in the course of action—for further readings on the subject we refer to [97].

Thompson’s construction [100] was popularized by the implementation of the UNIX command `grep` (globally search a regular expression and print). It amounts to the recursive connection of sub-automata *via* ε -transitions. These sub-automata are connected in parallel for the union, in series for the concatenation, and in an iterative fashion for the Kleene star. This yields an ε -NFA with a linear number of states and transitions. A structural characterization of the Thompson automaton in terms of the underlying digraph is given in [36, 37]. Thompson’s classical construction went through several stages of adaption and optimization. The construction with the least usage of ε -transitions was essentially given already in 1961 by Ott and Feinstein [91], which also can be found in [24, 77, 82]—see Figure 1. Later this




Figure 1: The inductive construction of Ott and Feinstein [91] yielding the precursor of the follow automaton $A_f(r)$ for a regular expression r .

construction was refined by Ilie and Yu [66] and promoted under the name *follow automaton*. In fact, the follow automaton is constructed from a regular expression r by recursively applying the construction of Ott and Feinstein and simultaneously improving on the use of ε -transitions in the following sense: (i) in the concatenation construction a ε -transition into the common state to both sub-automata leads to an appropriate state merging; similarly a state merging is done for an ε -transition leaving the common state, (ii) in the Kleene star construction, if the middle state is on a cycle of ε -transitions, all these transitions are removed, and all states of the cycle are merged, and (iii) after the construction is finished, a possible ε -transition from the start state is removed and both involved states are merged appropriately. Notice, that the automaton thus constructed may still contain ε -transitions. In order to amend the situation, an ε -removal procedure is applied: simply replace any sequence of an ε -transition followed by an a -transition directly by directly connecting the states on both ends of the sequence by a single a -transition directly. A final step takes care about the ε -transition to the final state. This results in the follow automaton $A_f(r)$ of [66], for the regular expression r .

Example 5 *Imagine a software buffer supporting the actions a (“add work packet”) and b (“remove work packet”), with a total capacity of n packets. Let r_n denote the regular expression for the action sequences that result in an empty buffer and never cause the buffer to exceed its capacity. Then*

$$r_1 = (ab)^* \quad \text{and} \quad r_n = (a \cdot r_{n-1} \cdot b)^*, \quad \text{for } n \geq 2.$$

Following the construction of the follow automaton as described in [66] results in the automaton depicted in Figure 2. Observe the constructed automaton is minimal, which is not the case in general. This is our




Figure 2: The follow automaton $A_f(r_n)$ accepting $L(r_n)$.

running example, where the behaviour of the state elimination technique described in the next section is discussed in more detail. \square

Preliminary bounds on the required size of a finite automaton equivalent to a given regular expression were given in [66]. Later, a tight bound in terms of reverse polish notation [51], and also a tight bound in terms of alphabetic width was found [42]. In the next theorem we summarize the results from [42, 51, 66]—here size of an automaton refers to the sum of the number of states and the number transitions:

Theorem 6 *Let $n \geq 1$, and r be a regular expression of alphabetic width n . Then size $\frac{22}{5}n$ is sufficient for an equivalent ε -NFA accepting $L(r)$. In terms of reverse polish length, the bound is $\frac{22}{15}(rpn(r) + 1) + 1$. Furthermore, there are infinitely many languages for which both bounds are tight.*

The aid for the tight bound in terms of the alphabetic width stated in the previous theorem is a certain normal form for regular expressions, which is a refinement of the *star normal form* from [13]. The definition reads as follows—transformation into strong star normal form preserves the described language, and is weakly monotone with respect to all usual size measures:

Definition 7 *The operators \circ and \bullet are defined on regular expressions² over alphabet Σ . The first operator is given by: $a^\circ = a$, for $a \in \Sigma$, $(r + s)^\circ = r^\circ + s^\circ$, $r^{2^\circ} = r^\circ$, $r^{*\circ} = r^\circ$; finally, $(r \cdot s)^\circ = r \cdot s$, if $\varepsilon \notin L(rs)$ and $r^\circ + s^\circ$ otherwise. The second operator is given by: $a^\bullet = a$, for $a \in \Sigma$, $(r + s)^\bullet = r^\bullet + s^\bullet$, $(r \cdot s)^\bullet = r^\bullet \cdot s^\bullet$, $r^{*\bullet} = r^{\circ*\bullet}$; finally, $r^{?^\bullet} = r^\bullet$, if $\varepsilon \in L(r)$ and $r^{?^\circ} = r^{\circ?}$ otherwise. The strong star normal form of an expression r is then defined as r^\bullet .*

What about the transformation of a regular expression into a finite automaton if ε -transitions are not allowed? One way to obtain an NFA directly is to perform the standard algorithm for removing ε -transitions, see, e.g., [61], which may increase the number of transitions at most quadratically. Another way is to directly implement the procedure during the recursive construction using non- ε -transitions to connect the sub-automata appropriately. Constructions of this kind can be found in, e.g., [3, 69]. For the conversion of ε -NFAs to NFAs the lower bound of [64] applies. There it was shown that there are infinitely many languages which are accepted by ε -NFAs with $O(n \cdot (\log n)^2)$ transitions, such that any NFA needs at least $\Omega(n^2)$ transitions. This lower bound is witnessed by a language over a growing size alphabet and shows that, in this case, the standard algorithm for removing ε -transitions cannot be improved significantly. For the case of binary alphabets, a lower bound of $\Omega(n \cdot 2^{c \cdot \sqrt{\log n}})$, for every $c < \frac{1}{2}$, was proved in [64] as well.

Another possibility to obtain ordinary NFAs is to directly construct the *position automaton*, also called the *Glushkov automaton* [38]—see also [83]. Intuitively, the states of this automaton correspond to the alphabetic symbols or, in other words, to positions between subsequent alphabetic symbols in the regular expression. Let us be more precise: assume that r is a regular expression over Σ of alphabetic

²Since \emptyset is only needed to denote the empty set, and the need for ε can be substituted by the operator $L^? = L \cup \{\varepsilon\}$, an alternative is to introduce also the $^?$ -operator and instead forbid the use of \emptyset and ε inside non-atomic expressions. This is sometimes more convenient, since one avoids unnecessary redundancy already at the syntactic level [42].

width n . In r we attach subscripts to each letter referring to its position (counted from left to right) in r . This yields a *marked* expression \bar{r} with distinct input symbols over an alphabet $\bar{\Sigma}$ that contains all letters that occur in \bar{r} . To simplify our presentation we assume that the same notation is used for unmarking, i.e., $\bar{\bar{r}} = r$. Then in order to describe the position automaton we need to define the following sets of positions on the marked expression. Let $\text{Pos}(r) = \{1, 2, \dots, \text{awidth}(r)\}$ and $\text{Pos}_0(r) = \text{Pos}(r) \cup \{0\}$. The position set First takes care of the possible beginnings of words in $L(\bar{r})$. It is inductively defined as follows:

$$\begin{aligned} \text{First}(\emptyset) &= \text{First}(\varepsilon) = \emptyset, \\ \text{First}(a_i) &= \{i\}, \\ \text{First}(s+t) &= \text{First}(s) \cup \text{First}(t), \\ \text{First}(s \cdot t) &= \begin{cases} \text{First}(s) \cup \text{First}(t) & \text{if } \varepsilon \in L(s) \\ \text{First}(s) & \text{otherwise,} \end{cases} \end{aligned}$$

and

$$\text{First}(s^*) = \text{First}(s).$$

Accordingly the position set Last takes care of the possible endings of words in $L(\bar{r})$. Its definition is similar to the definition of First , except for the concatenation, which reads as follows:

$$\text{Last}(s \cdot t) = \begin{cases} \text{Last}(s) \cup \text{Last}(t) & \text{if } \varepsilon \in L(t) \\ \text{Last}(t) & \text{otherwise.} \end{cases}$$

Finally, the set Follow takes care about the possible continuations in the words in $L(\bar{r})$. It is inductively defined as

$$\begin{aligned} \text{Follow}(\emptyset) &= \text{Follow}(\varepsilon) = \text{Follow}(a_i) = \emptyset \\ \text{Follow}(s+t) &= \text{Follow}(s) \cup \text{Follow}(t) \\ \text{Follow}(s \cdot t) &= \text{Follow}(s) \cup \text{Follow}(t) \cup \text{Last}(s) \times \text{First}(t) \end{aligned}$$

and

$$\text{Follow}(s^*) = \text{Follow}(s) \cup \text{Last}(s) \times \text{First}(s).$$

Then the position automaton for r is defined as $A_{\text{pos}}(r) = (\text{Pos}_0(r), \Sigma, \delta_{\text{pos}}, 0, F_{\text{pos}})$, where $\delta(0, a) = \{j \in \text{First}(\bar{r}) \mid a = \bar{a}_j\}$, for every $a \in \Sigma$ and $\delta(i, a) = \{j \mid (i, j) \in \text{Follow}(\bar{r}) \text{ and } a = \bar{a}_j\}$, for every $i \in \text{Pos}(r)$ and $a \in \Sigma$, and $F_{\text{pos}} = \text{Last}(\bar{r})$, if $\varepsilon \notin L(r)$, and $F_{\text{pos}} = \text{Last}(\bar{r}) \cup \{0\}$ otherwise.


Example 8 Consider the regular expression r_n from Example 5. If we mark the regular expression r_n , then we obtain $\bar{r}_n = (a_1(a_2(a_3 \dots b_{2n-2})^* b_{2n-1})^* b_{2n})^*$. Easy calculations show that the position sets read as follows:

$$\begin{aligned} \text{First}(\bar{r}_n) &= \{1\} \\ \text{Last}(\bar{r}_n) &= \{2n\} \end{aligned}$$

and

$$\text{Follow}(\bar{r}_n) = \{(i, i+1) \mid 1 \leq i < 2n\} \cup \{(i, 2n-i+1), (2n-i+1, i) \mid 1 \leq i \leq n\}$$

The position automaton on state set $\text{Pos}_0(r)$ is depicted in Figure 3. Here the set of final states is $F_{\text{pos}} = \{0, 2n\}$, since $\varepsilon \in L(r_n)$. Observe, that the follow automaton $A_f(r_n)$ can be obtained from $A_{\text{pos}}(r_n)$

Figure 3: The position automaton $A_{pos}(r_n)$ accepting $L(r_n)$.

by taking the quotient of automata, i.e., merging of states, with respect to the relation \equiv_f described in [66], which contains the elements $(i, 2n - i)$, for $0 \leq i \leq 2n$. This leads to the merging of states 0 and $2n$, states 1 and $2n - 1$, states 2 and $2n - 2$, up to states $n - 1$ and $n + 1$. \square

An immediate advantage of the position automaton is observed, e.g., in [1, 7]: for a regular expression r of alphabetic width n , for $n \geq 0$, the position automaton $A_{pos}(r)$ always has precisely $n + 1$ states. Simple examples, such as the singleton set $\{a^n\}$, show that this bound is tight. Nevertheless, several optimizations have been developed that give NFAs having often a smaller number of states, while the underlying constructions are mathematically sound refinements of the basic construction. A characterization of the position automaton is given in [16]. Moreover, structural comparisons between the position automaton with its refined versions, namely the *follow automaton*, the *partial derivative automaton* [4], or the *continuation automaton* [7] is given in [18, 66]. The partial derivative automaton is known under different names, such as *equation automaton* [85] or *Antimirov automaton* [4]. Further results on structural properties of these automata, when built from regular expressions in star normal form, can be found in [17, 20]. A quantitative comparison on the sizes of the the aforementioned NFAs for specific languages shows that they can differ a lot. The results listed in Table 1 are taken from [66]—here size of an automaton refers to the sum of the number of states and the number transitions. For comparison rea-

Expression	Finite Automaton			
	$A_f(\cdot)$	$A_{pd}(\cdot)$	$A_{pos}(\cdot)$	$A_{cfs}(\cdot)$
$r_1 = (a_1 + \varepsilon)^*$ and $r_{n+1} = (r_n + s_n)^*$ with $s_n = r_n[a_j \mapsto a_{j+2n-1}]$	$\Theta(r_n)$		$\Theta(r_n ^2)$	$\Theta(r_n \cdot (\log r_n)^2)$
$r_{n,m} = (\sum_{i=1}^n a_i)(\sum_{i=1}^n a_i + \sum_{i=1}^m b_i)^*$		$\Theta(r_{n,m})$	$\Theta(r_{n,m} ^2)$	$\Theta(r_{n,m} \cdot (\log r_{n,m})^2)$
$r_n = \sum_{i=1}^n a_i \cdot (b_1 + b_2 + \dots + b_n)^*$	$\Theta(r_n)$	$\Theta(r_n ^{1/2})$	$\Theta(r_n ^{3/2})$	$\Theta(r_n \cdot (\log r_n)^2)$
$r_n = (a_1 + \varepsilon) \cdot (a_2 + \varepsilon) \cdots (a_n + \varepsilon)$			$\Theta(r_n ^2)$	$\Theta(r_n \cdot (\log r_n)^2)$

Table 1: Comparing sizes of some automata constructions for specific languages from the literature—gray shading marks the smallest automaton. Here A_f refers to the follow automaton, A_{pd} to the partial derivative automaton, A_{pos} to the position automaton, and A_{cfs} to the common follow set automaton. Moreover, $|r_n|$ ($|r_{n,m}|$, respectively) refers to the alphabetic width of the regular expression r_n ($r_{n,m}$, respectively).

sons also the *common follow set automaton* A_{cfs} is listed—since the description of A_{cfs} is quite involved we refer the reader to [65]. There, this automaton was used to prove an upper bound on the number of transitions. The issue on transitions for NFAs, in particular when changing from an ε -NFA to an NFA, is discussed next.

Despite the mentioned optimizations, except for the common follow set automaton, all of these constructions share the same problem with respect to the number of transitions. An easy upper bound on the number of transitions in the position automaton is $O(n^2)$, independent of alphabet size. It is not hard to prove that the position automaton for the regular expression

$$r_n = (a_1 + \varepsilon) \cdot (a_2 + \varepsilon) \cdots (a_n + \varepsilon)$$

has $\Omega(n^2)$ transitions. It appears to be difficult to avoid such a quadratic blow-up in actual size if we stick to the NFA model. Also if we transform the expression first into a ε -NFA and perform the standard algorithm for removing ε -transitions, see, e.g., [61], we obtain no better result. This naturally raises the question of comparing the descriptive complexity of NFAs over regular expressions. For about forty years, it appears to have been considered as an unproven factoid that a quadratic number of transitions will be inherently necessary in the worst case (cf. [65]). A barely super-linear lower bound of $\Omega(n \log n)$ on the number of transitions of any NFA accepting the language of the expression r_n was proved [65]. More interestingly, the main result of that paper is an algorithm transforming a regular expression of size n into an equivalent NFA with at most $O(n \cdot (\log n)^2)$ transitions. See Figure 4 on how the algorithm of [65] saves transitions for regular expression r_n , explained for $n = 5$. In fact, this upper bound made their lower




Figure 4: Let $r_n = (a_1 + \varepsilon) \cdot (a_2 + \varepsilon) \cdots (a_n + \varepsilon)$ and $n = 5$. Position automaton $A_{pos}(r_5)$ (left) and its refined version the common follow set automaton $A_{cfs}(r_5)$ (right) accepting language $L(r_5)$; in both cases the dead state and all transitions leading to it are not shown. The automaton $A_{cfs}(r_5)$ is obtained as follows: the state 1 of $A_{pos}(r_5)$ is split such that the new state gets the outgoing transitions labeled with a_3, a_4 , and a_5 , and is finally identified with state 2, which can be done since it has the same outgoing transitions.

bound look reasonable at once! Shortly thereafter, an efficient implementation of that conversion algorithm was presented [52], and the lower bound was improved in [73] to $\Omega(n \cdot (\log n)^2 / \log \log n)$. Later work [98] established that any NFA accepting language $L(r_n)$ indeed must have at least $\Omega(n \cdot (\log n)^2)$ transitions. So the upper bound of $O(n \cdot (\log n)^2)$ from [65] is asymptotically tight:

Theorem 9 *Let $n \geq 1$ and r be a regular expression of alphabetic width n . Then $O(n \cdot (\log n)^2)$ transitions are sufficient for an NFA to accept $L(r)$. Furthermore, there are infinitely many languages for which this bound is tight.*

Notice that the example witnessing the lower bound is over an alphabet of growing size. For alphabets of size two, the upper bound was improved first [33] to $O(n \cdot \log n)$, and then even to $n \cdot 2^{O(\log^* n)}$, where \log^* denotes the iterated binary logarithm [98]. Moreover, a lower bound of $\Omega(n \cdot (\log k)^2)$ on the size of NFAs with k -letter input alphabet was shown in [98], too. Thus the question from [62] whether a conversion from regular expressions over a binary alphabet into NFAs of linear size is possible, is almost settled by now.

Theorem 10 *Let $n \geq 1$ and r be a regular expression of alphabetic width n over a binary alphabet. Then $n \cdot 2^{O(\log^* n)}$ transitions are sufficient for a NFA to accept $L(r)$.*

Next, let us briefly discuss the problem of converting regular expressions to DFAs. Again, this problem has been studied by many authors. The obvious way to obtain a DFA is by applying the well known *subset* or *power-set construction* [94]. Due to this construction the obtained DFA may be of exponential size. A more direct and convenient way is to use Brzozowski's derivatives of expressions [14]. A taxonomy comparing many different conversion algorithms is given in [102]. Regarding the descriptive complexity, a tight bound of $2^n + 1$ states in terms of alphabetic width is given in [69]. The mentioned work also establishes a matching lower bound, but for a rather nonstandard definition of size. In terms of alphabetic width, the best lower bound known to date is from [28]. Together, we have the following result:

Theorem 11 *Let $n \geq 1$ and r be a regular expression of alphabetic width n over a binary alphabet. Then $2^n + 1$ states are sufficient for a DFA to accept $L(r)$. In contrast, for infinitely many n there are regular expressions r_n of alphabetic width n over a binary alphabet, such that the minimal DFA accepting $L(r_n)$ has at least $\frac{5}{4}2^{\frac{n}{2}}$ states.*

Recent developments on the conversion of regular expressions to finite automata show an increasing attention on the study of descriptive complexity in the average case. For instance, in [89] it was shown that, when choosing the expression uniformly at random, the position automaton has $\Theta(n)$ transitions on average, where n refers to the nodes in the parse tree of the expression. A similar result holds w.r.t. alphabetic width, for the position automaton as well as for the partial derivative automaton [11]. A closer look reveals that the number of transitions in the partial derivative automaton is, on average, half the size of the number of transitions in the position automaton [11], for large alphabet sizes; this also holds for the number of states [10]. Results on the average size of ε -NFAs built from Thompson's construction and variants thereof [66, 99, 100] can be found in [12]—in their investigation the authors consider the follow automaton before the final ε -removal is done. Let us call this device *ε -follow automaton*. It turns out that the ε -follow automaton is superior to the other constructions considered. In particular, the number of ε -transitions asymptotically tends to zero, i.e., the ε -follow automaton approaches the follow-automaton.

Almost all of these results were obtained with the help of the framework of analytic combinatorics [29]. The idea to use this approach is quite natural. Recall, that the number of regular expressions of a certain size measured by, e.g., alphabetic width, can be counted by using generating functions—for more involved measures, one has to use multivariate generating functions. To this end one transforms a grammar describing regular expressions such as, e.g., the grammar devised in [50], into a generating function. Since the grammar describes a combinatorial class, the generating function can be obtained by the *symbolic method* of [29], and the coefficients of the power series can be estimated to give approximations of the measure under consideration.

Finally, let us note, that the results on the average size of automata depends on the probability distribution that is used for the average-case analysis. In [90] it was shown that the number of transitions of the position automaton is in $\Theta(n^2)$ under a distribution that is inspired from random binary search trees (BST-like model). To our knowledge, average case analysis under the BST-like model for other automata such as the follow automaton or the partial derivative automaton, has not been conducted so far.

4 From Finite Automata to Regular Expressions

There are a few classical algorithms for converting finite automata into equivalent regular expressions, namely

1. the *algorithm based on Arden's lemma* [5, 22], and
2. the *McNaughton-Yamada algorithm* [83], and
3. the *state elimination technique* [15].

These procedures look different at first glance. We briefly explain the main idea of these approaches—for a detailed description along with an explanation of the differences between the methods, the reader is referred to [97]. There it is shown, that all of the above approaches are more or less reformulations of the same underlying algorithmic idea, and they yield (almost) the same regular expressions.³

An algebraic approach to solve the conversion problem from finite automata to regular expressions is the *algorithm based on Arden's lemma* [5, 22]. It puts forward a set of language equations for a given finite automaton. Here, the i th equation describes the set X_i of words w such that the given automaton can go from the i th state to an accepting state on reading w . That system of equations can be resolved by eliminating the indeterminates X_i using a method that resembles Gaussian elimination. But we work in an algebraic structure different from a field, so for the elimination of variables, we have to resort to *Arden's lemma*:

Lemma 12 *Let Σ be an alphabet, and let $K, L \subseteq \Sigma^*$, where K does not contain the empty word ε . Then the set K^*L is the unique solution to the language equation $X = K \cdot X + L$, where X is the indeterminate.*

Now let us have a look on how Arden's lemma can be applied to our running example.

Example 13 *From the automaton depicted in Figure 2 one reads off the equations*

$$X_0 = a \cdot X_1 + \varepsilon, \quad X_i = a \cdot X_{i+1} + b \cdot X_{i-1}, \quad \text{for } 1 \leq i < n, \quad \text{and} \quad X_n = b \cdot X_{n-1}.$$

*Substituting the right hand side of X_n in the next to last equation and solving it by Arden's lemma results in $X_{n-1} = (ab)^*b \cdot X_{n-2}$. For short, $X_{n-1} = r_1 \cdot b \cdot X_{n-2}$, where r_i is defined as in Example 5. Next this solution is substituted into the equation for X_{n-2} . Solving for X_{n-2} gives us $X_{n-2} = r_2 \cdot b \cdot X_{n-3}$. Proceeding in this way up to the very first equation gives us $X_0 = a \cdot r_{n-1} \cdot b \cdot X_0 + \varepsilon$. The solution to the indeterminate X_0 is according to Arden's lemma $(a \cdot r_{n-1} \cdot b)^* \cdot \varepsilon = r_n$, by applying obvious simplifications. Hence, for instance, in case $n = 6$ we obtain $(a(a(a(a(ab)^*b)^*b)^*b)^*b)^*$. \square*

The *McNaughton-Yamada algorithm* [83] maintains a matrix with regular expression entries, where the rows and columns are the states of the given automaton. The iterative algorithm uses a ranking on the state set, and proceeds in n rounds, if n is the number of states in the given automaton A . In the matrix $(a_{jk})_{j,k}$ computed in round i , the entry a_{jk} is an expression describing the nonempty labels w of computations of A starting in j and ending in k , such that none of the intermediate states of the computation is ranked higher than i . From these expressions, it is not difficult to obtain a regular expression describing $L(A)$.

Example 14 *Running the McNaughton-Yamada algorithm on the automaton depicted in Figure 2 for $n = 3$ with the ranking 3, 2, 1, 0 starts with the following matrix:*

$$\begin{array}{c} 3 \quad 2 \quad 1 \quad 0 \\ 3 \begin{pmatrix} \emptyset & b & \emptyset & \emptyset \\ a & \emptyset & b & \emptyset \\ \emptyset & a & \emptyset & b \\ \emptyset & \emptyset & a & \emptyset \end{pmatrix} \\ 2 \\ 1 \\ 0 \end{array}$$

³Let us also mention that there is another algebraic algorithm from [22], which is based on the recursive decomposition of matrices into blocks. Here, the precise relation to the aforementioned algorithms remains to be investigated [97].

If $(a_{jk})_{j,k}$ denotes the matrix computed in round i , then the matrix $(b_{jk})_{j,k}$ for round $i+1$ can be computed using the rule

$$b_{jk} = a_{jk} + a_{ji}(a_{ii})^*a_{ik}$$

After the first round, the entry in the upper left corner of the matrix reads as $\emptyset + \emptyset\emptyset^*\emptyset$. It is of course helpful to simplify the intermediate regular expressions, by applying some obvious simplifications. As noted in [83], we can use in particular

$$b_{ij} = (a_{ii})^*a_{ij} \quad \text{and} \quad b_{ji} = a_{ji}(a_{ii})^*.$$

Then the matrix computed in the first round reads as

$$\begin{pmatrix} \emptyset & b & \emptyset & \emptyset \\ a & ab & b & \emptyset \\ \emptyset & a & \emptyset & b \\ \emptyset & \emptyset & a & \emptyset \end{pmatrix},$$

the one from the second round is

$$\begin{pmatrix} b(ab)^*a & b(ab)^* & b(ab)^*b & \emptyset \\ (ab)^*a & (ab)^*ab & (ab)^*b & \emptyset \\ a(ab)^*a & a(ab)^* & a(ab)^*b & b \\ \emptyset & \emptyset & \emptyset & \emptyset \end{pmatrix},$$

and the computation is continued in the same vein. Finally, the entry in the lower-right corner of the matrix reads as $(a(a(ab)^*b)^*b)^*a(a(ab)^*b)^*b$, and the desired regular expression describing L_3 is obtained by adding the empty word: $\varepsilon + (a(a(ab)^*b)^*b)^*a(a(ab)^*b)^*b$.

A few industrious readers, who have worked out the calculation of the previous example until the final matrix, may have observed that many of the intermediate expressions were actually not needed for the final result. Indeed, in a computer implementation [79, page 8] of the basic McNaughton-Yamada algorithm during the 1960s, the author notes: “a basic fault of the method is that it generates such cumbersome and so numerous expressions initially.” Below we discuss how the generation of unnecessary sub-expressions can be avoided.

We now come to an algorithm that we describe in greater detail, namely the *state elimination algorithm* [15]. This procedure maintains an extended finite automaton, whose transitions are labeled with regular expressions, rather than alphabet symbols. The computation of an NFA A can be thought of as reading the input word letter by letter, thereby nondeterministically changing its state with each letter in a way that is consistent with its transition table δ . On reading a word $w \in \Sigma$, we say that the finite automaton A can go on input w from state j to state k , if there is a computation on input w taking A from state j to k . Similarly, for a subset U of the state set Q of the automaton A , we say that A can go on input w from state j through U to state k , if there is a computation on input w taking A from state j to k , without going through any state outside U , except possibly j and k . With the rôles of j , k , and U fixed as above, we now define the language L_{jk}^U as the set of input words on which the automaton A can go from j to k through U . The state elimination scheme fixes an ordering on the state set Q . Starting with $U = \emptyset$, regular expressions denoting the languages L_{jk}^0 for all pairs $(j, k) \in Q \times Q$ can be easily read off from the transition table of A . Now an important observation is that for each state $i \in Q \setminus U$ holds

$$L_{jk}^{U \cup \{i\}} = L_{jk}^U \cup L_{ji}^U \cdot (L_{ii}^U)^* \cdot L_{ik}^U.$$

Letting i run over all states according to the ordering, we can grow the set U one by one, in each round computing the intermediate expressions $r_{jk}^{U \cup \{i\}}$ for all j and k . The final regular expression is obtained by utilizing the fact $L(A) = \bigcup_{f \in F} L_{q_0, f}^Q$.

As observed already by McNaughton and Yamada [83], we have $\text{awidth}(L_{jk}^0) \leq |\Sigma|$, and each round increases the alphabetic width of each intermediate sub-expression by a factor of at most 4. Another convenient trick is to modify the automaton, by adding a new initial state s and a new final state t to the automaton without altering the language, such that t is the single final state, and there are no transitions entering s or leaving t . Then s and t need not to be added to the set U . Instead, observe that $L(A) = L_{st}^U$, with $U = Q \setminus \{s, t\}$. We also note⁴ that the computation of r_{jk}^U needs to be carried out only for those j and k not in U . We thus obtain the following bound:

Theorem 15 *Let $n \geq 1$ and A be an n -state NFA over alphabet Σ . Then alphabetic width $|\Sigma| \cdot 4^n$ is sufficient for a regular expression describing $L(A)$. Such an expression can be constructed by state elimination.*

In contrast, the state elimination algorithm might suddenly yield a much simpler regular expression once we change the ordering in which the states are eliminated. We illustrate the influence of the elimination ordering on a small example.

Example 16 *Consider our software buffer from Example 5 for $n = 6$. Let $L_n := L(r_n)$. For illustration, a minimal DFA for L_6 is depicted in Figure 5. The two regular expressions*

$$(a(a(a(a(ab)^*b)^*b)^*b)^*b)^*$$

and

$$\varepsilon + a(ab + ba)^*b + a(ab + ba)^*aa(ab + ba + bb(ab + ba)^*aa + aa(ab + ba)^*bb)^*bb(ab + ba)^*b$$

both describe the language L_6 . The first expression is obtained by eliminating the states in the order 6, 5, 4, 3, 2, 1, and 0, while the second expression is produced by the order 0, 2, 4, 6, 1, 5, and 3. Note




Figure 5: A minimal DFA accepting the language $L_6 := L(r_6)$.

that the expressions have very different structure. The first is much shorter, but has star height 6, while the second, and longer expression, has star height 2. Indeed, in [81] it was shown that the minimum star height among all regular expressions denoting L_n equals $\lfloor \log(n+1) \rfloor$, so the star height of the second expression is optimal. The authors suspect that this language family exhibits a trade-off in the sense that the regular expressions for L_n cannot be simultaneously short and of low star height.

Perhaps the earliest reference mentioning the influence of the elimination ordering is from 1960. In [83], they proposed to identify the states that “bear the most traffic,” i.e., those vertices in the underlying graph with the highest degree, and to eliminate these states at last. Since then, various heuristics for computing elimination orderings that yield short regular expressions have been proposed in the literature. In [74], a simple greedy heuristic was devised. It was proposed to assign a measure to each

⁴The same trick applies for the McNaughton-Yamada algorithm: If the single initial and the single final state are not eliminated, we can erase the entries of the i th row and the i th column of the computed matrix in round i .

state, and this measure is recomputed each time when a state is eliminated. This measure indicates the priority in which the states are eliminated. Observe that eliminating a state tends to introduce new arcs in the digraph underlying the automaton. Thus we can order the states by a measure that is defined as the number of ingoing arcs times the number of outgoing arcs. In [23] a refined version of the same idea is proposed, which takes also the lengths of the intermediate expressions into account, instead of just counting the ingoing and outgoing arcs. Later, a different strategy for accounting the priority of a state was suggested: as measure function, simply take the number of cycles passing through a state. There are some automata, where this heuristic outperforms the one we previously described, but on most random DFAs the performance is comparable. For the heuristic based on counting the number of cycles, recomputing the measure after the elimination of each state does not make a big difference [87]. Another idea is to look for simple structures in finite automata, such as bridge states [53]. A bridge state typically exists if the language under consideration can be written as the concatenation of two nontrivial regular languages. Unfortunately, a random DFA almost surely contains no bridge states at all, as the number of states grows larger [87]. These and other heuristics were compared empirically on a large set of random DFAs as input in [48, 87]. Although there are also advanced strategies for choosing an elimination ordering, which have provable performance guarantees, the greedy heuristic from [23] performs best in most cases.

Beyond heuristics, we can use elimination orderings to prove nontrivial upper bounds on the conversion of DFAs over small alphabets into regular expressions. For the case of binary alphabets, a bound of $O(1.742^n)$ was given in [44], which was then improved to $O(1.682^n)$ in [25]. These bounds can be reached with state elimination by using appropriate elimination orderings. The latest record is $O(1.588^n)$, and the algorithm departs from pure state elimination, see [47].

Theorem 17 *Let $n \geq 1$ and A be an n -state DFA over a binary alphabet. Then size $O(1.588^n)$ is sufficient for a regular expression describing $L(A)$.*

Similar bounds, but with somewhat larger constants in place of 1.588, can be derived for larger alphabets. Moreover, the same holds for NFAs having a comparably low density of transitions.

We sketch how to establish a simpler upper bound than this, which after all gives $o(4^n)$ for all alphabets of constant size. To get things going, assume that we want to determine L_{jk}^U , and that the underlying sub-graph induced by U falls apart into two mutually disconnected sub-graphs A and B . Then on reading a word w , the automaton goes from j to k either through A or through B , and thus $L_{jk}^U = L_{jk}^A \cup L_{jk}^B$, and this is reflected by the regular expressions computed using state elimination. In particular, if the sub-graph induced by U is an independent set, i.e., a set of isolated vertices, in the underlying graph, then $L_{jk}^U = \bigcup_{i \in U} L_{jk}^{\{i\}}$. In this case, the blow-up factor incurred by eliminating U is linear in $|U|$, instead of exponential in $|U|$. For a DFA A over constant alphabet, the underlying graph has a linear number of edges. It is known that such graphs have an independent set of size cn , where c is a constant depending on the number of edges. Suppose that U is such an independent set. Then we partition the state set of A into an “easy” part U and a “hard” part $Q \setminus U$. Eliminating U increases the size of the intermediate expressions by a factor linear in $|U|$. Thereafter, eliminating the remaining $(1 - c)n$ states may incur a size blow-up by a factor of $4^{(1-c)n}$. Altogether, this gives a regular expression of alphabetic width in $|\Sigma| \cdot o(4^n)$ for $L(A)$.

Let us again take a look at an example.

Example 18 *For illustrating the above said, consider the language*

$$L_3 = (a_1 b_1)^* \sqcup (a_2 b_2)^* \sqcup (a_3 b_3)^*,$$




Figure 6: Automaton accepting the language $L_3 = (a_1b_1)^* \sqcup (a_2b_2)^* \sqcup (a_3b_3)^*$. The underlying graph is the 3-dimensional cube.

where the interleaving, or shuffle, of two languages L_1 and L_2 over alphabet Σ is

$$L \sqcup M = \{ w \in \Sigma^* \mid w \in x \sqcup y \text{ for some } x \in L \text{ and } y \in M \},$$

and the interleaving $x \sqcup y$ of two words x and y is defined as the set of all words of the form $x_1y_1x_2y_2 \cdots x_ny_n$, where $x = x_1x_2 \cdots x_n$, $y = y_1y_2 \cdots y_n$ with $x_i, y_i \in \Sigma^*$, for $n \geq 1$ and $1 \leq i \leq n$. Note that in this definition, some of the sub-words x_i and y_i can be empty.

The language L_3 can be accepted by a partial DFA over the state set $\{0, 1\}^3$, and whose transition function is given such that input a_i sets the i th bit left of the rightmost bit of the current state from 0 to 1, and input b_i resets the i th bit, again counting from right to left, of the current state from 1 to 0. All other transitions are undefined. The initial state is 000, which is also the single final state. Notice that the graph underlying this automaton is the 3-dimensional cube, with 8 vertices—see Figure 6. Generalizing this example to $d \geq 3$, the underlying graph of L_d is the d -dimensional hypercube, with 2^d many vertices.

It is well known that the d -dimensional hypercube is 2-colorable, and thus has an independent set that contains at least half of the vertices. Eliminating this independent set before the other vertices yields a regular expression of alphabetic width $O(n \cdot 2^n)$, which is way better than the trivial bound of $O(4^n)$.

We present another application of this idea. Planar finite automata are a special case of finite automata, which were first studied in [9]. To convert a planar finite automaton into a regular expression, one can look for a small set of vertices, whose removal leaves to mutually disconnected sub-graphs with vertex sets A and B . Then again, we have $L_{jk}^U = L_{jk}^A \cup L_{jk}^B$, and this is reflected by the regular expressions computed by state elimination. Since the sub-graphs induced by A and B are again planar, one can apply the trick recursively. Also for this special case, tight upper and lower bounds were found recently [28, 43, 46].

Theorem 19 *Let $n \geq 1$ and A be an n -state planar DFA or NFA over alphabet Σ . Then size $|\Sigma| \cdot 2^{O(\sqrt{n})}$ is sufficient for a regular expression describing $L(A)$. Such an expression can be constructed by state elimination.*

Taking this idea again a step further, one can arrive at a parametrization where the conversion problem from finite automata to regular expressions is fixed-parameter tractable, in the sense that the problem is

exponential in that parameter, but not in the size of the input. Recall that we have introduced the concept of cycle rank of a digraph in the course of discussing the star height in Section 2. Now for a digraph D , let D^{sym} denote the symmetric digraph obtained by replacing each arc in D with a pair of anti-parallel arcs. The *undirected cycle rank* of D is defined as the cycle rank of D^{sym} . If the conversion problem from finite automata to regular expressions is parametrized by the undirected cycle rank of the given automaton, one can prove the following bound [46]:

Theorem 20 *Let $n \geq 1$ and A be an n -state DFA or NFA over alphabet Σ , whose underlying digraph is of undirected cycle rank at most c , for some $c \geq 1$. Then size $|\Sigma| \cdot 4^c \cdot n$ is sufficient for a regular expression describing $L(A)$. Such an expression can be constructed by state elimination.*

Observe that fixed-parameter tractability also holds in the sense of computational complexity, since computing the undirected cycle rank is fixed-parameter tractable, see, e.g., [96]. A natural question is now whether we can find a similar parametrization in terms of cycle rank, instead of undirected cycle rank. Well, there are acyclic finite automata that require regular expressions of super-polynomial size [27, 49]. Notice that these automata have cycle rank 0. Hence the best we can hope for is a parametrization that is quasi-polynomial when the cycle rank is bounded. One can indeed obtain such an estimate [47], but the method is more technical, and no longer uses only state elimination. The upper bound in terms of directed cycle rank reads as follows:

Theorem 21 *Let $n \geq 1$ and A be an n -state DFA or NFA over alphabet Σ , whose underlying digraph is of cycle rank at most c , for some $c \geq 1$. Then size $|\Sigma| \cdot n^{O(c \cdot \log n)}$ is sufficient for a regular expression describing $L(A)$.*

But in the general case, the exponential blow-up when moving from finite automata to regular expressions is inherent, that is, independent of the conversion method. Already in the 1970s the existence of languages L_n was shown, that admit n -state finite automata, but require regular expressions of alphabetic width at least 2^{n-1} , for all $n \geq 1$, see [27]. Their witness language is over an alphabet of growing size, which is quadratic in the number of states. Their proof technique was tailored to the witness language involved. The question whether a comparable size blow-up can also occur for constant alphabet size [28] was settled only a few years ago. The answer was provided around the same time by two independent groups of researchers, who worked with different proof techniques, and gave different examples [35, 43].

How are such lower bounds established? We shall describe a general method, which has been used to prove lower bounds on regular expression size in various contexts [34, 43, 45, 55]. In the context of lower bounds for regular expression size, a more convenient formulation of Lemma 4 is the star height lemma, which reads as follows:

Lemma 22 *Let L be a regular language. Then $\text{awidth}(L) \geq 2^{\Omega(\text{height}(L))}$.*

That is, the minimum regular expression size of a regular language is at least exponential in the minimum required star height. But now this looks as if we have replaced one evil with another, since determining the star height is eminently difficult in general [67]. But there is an important special case, in which the star height can be determined more easily: a *partial* deterministic finite automaton is called *bideterministic*, if it has a single final state, and if the NFA obtained by reversing all transitions and exchanging the roles of initial and final state is again a partial DFA—notice that, by construction, this NFA in any case accepts the reversed language. A regular language L is *bideterministic* if there exists a bideterministic finite automaton accepting L . These languages form a proper subclass of the regular languages. For these languages, *McNaughton's Theorem* [80] states that the star height is equal to the cycle rank of the digraph underlying the minimal partial DFA.

Example 23 Define $K_m = \{w \in \{a, b\}^* \mid |w|_a \equiv 0 \pmod m\}$ and $L_n = \{w \in \{a, b\}^* \mid |w|_b \equiv 0 \pmod n\}$. For simplicity, assume $m \leq n$. It is straightforward to construct deterministic finite automata with m states (with n states, respectively) arranged in a directed cycle describing the languages K_m and L_n , respectively. By applying the standard product construction on these automata, we obtain a deterministic finite automaton A accepting the language $K_m \cap L_n$. The digraph underlying automaton A is the directed




Figure 7: Drawing of the discrete directed $(m \times n)$ -torus in the case where $m = 2$ and $n = 4$, induced by the automata for the languages K_m and L_n .

discrete torus. This digraph can be described as the Cartesian graph product of two directed cycles, see Figure 7 for illustration. The cycle rank of the $(m \times n)$ -torus is equal to m if $m = n$, and equal to $m + 1$ otherwise [43]. It is easily observed that the automaton A is bideterministic, hence the star height of $L(A)$ coincides with the cycle rank of its underlying digraph. By invoking the star height lemma, we can derive a lower bound of $2^{\Omega(m)}$ on the minimum regular expression size required for $L_m \cap K_n$.

For the succinctness gap between DFAs and regular expressions over binary alphabets, a lower bound of $2^{\Omega(\sqrt{n/\log n})}$ was reported in [35], while a parallel effort [43] resulted in an asymptotically tight lower bound of $2^{\Omega(n)}$. We have the following result:

Theorem 24 Let $n \geq 1$ and A be an n -state DFA or NFA over alphabet Σ . Then size $|\Sigma| \cdot 2^{\Theta(n)}$ is sufficient and necessary in the worst case for a regular expression describing $L(A)$. This already holds for alphabets with at least two letters.

Recall that the notation $2^{\Theta(n)}$ implies a lower bound of c^n , for some $c > 1$. The hidden constant in the lower bound for binary alphabets is much smaller compared to the lower bound of 2^{n-1} previously obtained in [27] for large alphabets. The upper bound from Theorem 17 implies that c can be at most 1.588 for alphabets of size two. Narrowing down the interval for the best possible c for various alphabet sizes is a challenge for further research.

We turn our attention to interesting special cases of regular languages, namely the *finite* and the *unary* regular languages. Here, the situation is significantly different, as we can harness specialized techniques which are more powerful than state elimination. Also, finite and unary languages have star height at most 1, and thus more tailored techniques than the star height lemma are needed to establish lower bounds. Indeed, the case of finite languages was already addressed in the very first paper on the descriptonal complexity of regular expressions [27]. They give a specialized conversion algorithm for finite languages, which is different from the state elimination algorithm. Their results imply that every n -state DFA accepting a finite language can be converted into an equivalent regular expression

of size $n^{O(\log n)}$. The method is quite interesting, since it is not based on state elimination, but rather on a clever application of the repeated squaring trick. They also provide a lower bound of $n^{\Omega(\log \log n)}$ when using an alphabet of size $O(n^2)$. The challenge of tightening this gap was settled more than thirty years later in [49], where a lower bound technique from communication complexity is adapted, which originated in the study of monotone circuit complexity.

Theorem 25 *Let $n \geq 1$ and A be an n -state DFA or NFA over alphabet Σ accepting a finite language. Then size $|\Sigma| \cdot n^{\Theta(\log n)}$ is sufficient and necessary in the worst case for a regular expression describing $L(A)$. This still holds for constant alphabets with at least two letters.*

The case of unary languages was discussed in [32, 78, 101]. Here the main idea is that one can exploit the simple cycle structure of unary DFAs and of unary NFAs in Chrobak normal form [21]. In the case of NFAs, elementary number theory helps to save a logarithmic factor of the quadratic upper bound [32]. The main results are summarized in the following theorem.

Theorem 26 *Let $n \geq 1$ and A be an n -state DFA over a unary alphabet. Then size $\Theta(n)$ is sufficient and necessary in the worst case for a regular expression describing $L(A)$. When considering NFAs, the upper bound changes to $O(n^2 / \log n)$.*

The tight bounds for the conversion of unary NFAs to regular expressions thus remain to be determined. The conversion problem has been studied also for a few other special cases of finite automata. Examples include finite automata whose underlying digraph is an acyclic series-parallel digraph [88], Thompson digraphs [36], and digraphs induced by Glushkov automata [16].

Acknowledgments. Thanks to Sebastian Jakobi for helpful comments and suggestions on an earlier draft of this paper.

References

- [1] A. Aho, R. Sethi & J. D. Ullman (1986): *Compilers: Principles, Techniques, and Tools*. Addison Wesley.
- [2] A. V. Aho, J. E. Hopcroft & J. D. Ullman (1974): *The Design and Analysis of Computer Algorithms*. Addison-Wesley.
- [3] A. V. Aho & J. D. Ullman (1972): *The Theory of Parsing, Translation and Compiling*. I, Prentice-Hall.
- [4] V. Antimirov (1996): *Partial derivatives of regular expressions and finite automaton constructions*. *Theoretical Computer Science* 155(2), pp. 291–319, doi:10.1016/0304-3975(95)00182-4.
- [5] D. N. Arden (1961): *Delayed-Logic and Finite-State Machines*. In T. Mott, editor: *Proceedings of the 1st and 2nd Annual Symposium on Switching Theory and Logical Design*, American Institute of Electrical Engineers, New York, Detroit, Michigan, USA, pp. 133–151, doi:10.1109/F0CS.1961.13.
- [6] A. Asperti, C. Sacerdoti Coen & E. Tassi (2010): *Regular Expressions, au point*. arXiv:1010.2604v1 [cs.FL].
- [7] G. Berry & R. Sethi (1986): *From Regular Expressions to Deterministic Automata*. *Theoretical Computer Science* 48(3), pp. 117–126, doi:10.1016/0304-3975(86)90088-5.
- [8] Ph. Bille & M. Thorup (2010): *Regular Expression Matching with Multi-Strings and Intervals*. In M. Charikar, editor: *Proceedings of the 21st Annual ACM-SIAM Symposium on Discrete Algorithms*, Society for Industrial and Applied Mathematics, Austin, Texas, USA, pp. 1297–1308, doi:10.1137/1.9781611973075.104.
- [9] R. V. Book & A. K. Chandra (1976): *Inherently Nonplanar Automata*. *Acta Informatica* 6(1), pp. 89–94, doi:10.1007/BF00263745.

- [10] S. Broda, A. Machiavelo, N. Moreira & R. Reis (2010): *On the Average Number of States of Partial Derivative Automata*. In Y. Gao, H. Lu, S. Seki & S. Yu, editors: *Proceedings of the 14th International Conference Developments in Language Theory, LNCS 6224*, Springer, London, Ontario, Canada, pp. 112–123, doi:10.1007/978-3-642-14455-4_12.
- [11] S. Broda, A. Machiavelo, N. Moreira & R. Reis (2012): *On the Average Size of Glushkov and Partial Derivative Automata*. *International Journal of Foundations of Computer Science* 23(5), pp. 969–984, doi:10.1142/S0129054112400400.
- [12] S. Broda, A. Machiavelo, N. Moreira & R. Reis (2014): *A Hitchhiker’s Guide to Descriptive Complexity Through Analytic Combinatorics*. *Theoretical Computer Science* 528, pp. 85–100, doi:10.1016/j.tcs.2014.02.013.
- [13] A. Brüggemann-Klein (1993): *Regular Expressions into Finite Automata*. *Theoretical Computer Science* 120, pp. 197–213, doi:10.1016/0304-3975(93)90287-4.
- [14] J. A. Brzozowski (1964): *Derivatives of regular expressions*. *Journal of the ACM* 11, pp. 481–494, doi:10.1145/321239.321249.
- [15] J. A. Brzozowski & E. J. McCluskey (1963): *Signal flow graph techniques for sequential circuit state diagrams*. *IEEE Transactions on Computers* C-12(2), pp. 67–76, doi:10.1109/PGEC.1963.263416.
- [16] P. Caron & D. Ziadi (2000): *Characterization of Glushkov automata*. *Theoretical Computer Science* 233(1–2), pp. 75–90, doi:10.1016/S0304-3975(97)00296-X.
- [17] J.-M. Champarnaud, F. Ouardi & D. Ziadi (2007): *Normalized Expressions and Finite Automata*. *International Journal of Algebra and Computation* 17(1), pp. 141–154, doi:10.1142/S021819670700355X.
- [18] J.-M. Champarnaud & D. Ziadi (2002): *Canonical derivatives, partial derivatives and finite automaton constructions*. *Theoretical Computer Science* 289(1), pp. 137–163, doi:10.1016/S0304-3975(01)00267-5.
- [19] C.H. Chang & R. Paige (1992): *From Regular Expressions to DFA’s Using Compressed NFA’s*. In A. Apostolico, M. Crochemore, Z. Galil & U. Manber, editors: *Proceedings of the 3rd Annual Symposium on Combinatorial Pattern Matching, LNCS 644*, Springer, Tucson, Arizona, USA, pp. 90–110, doi:10.1007/3-540-56024-6_8.
- [20] H. Chen (2010): *Finite Automata of Expressions in the Case of Star Normal Form and One-Unambiguity*. Technical Report ISCAS-LCS-10-11, Chinese Academy of Sciences, Institute of Software, State Key Laboratory of Computer Science, Beijing 100190 China.
- [21] M. Chrobak (1986): *Finite automata and unary languages*. *Theoretical Computer Science* 47, pp. 149–158, doi:10.1016/0304-3975(86)90142-8.
- [22] J. H. Conway (1971): *Regular Algebra and Finite Machines*. Chapman and Hall.
- [23] M. Delgado & J. Morais (2004): *Approximation to the Smallest Regular Expression for a Given Regular Language*. In M. Domaratzki, A. Okhotin, K. Salomaa & S. Yu, editors: *Proceedings of the 9th Conference on Implementation and Application of Automata, LNCS 3317*, Springer, Kingston, Ontario, Canada, pp. 312–314, doi:10.1007/978-3-540-30500-2_31.
- [24] D.-Z. Du & K.-I. Ko (2001): *Problem Solving in Automata, Languages, and Complexity*. John Wiley & Sons, doi:10.1002/0471224642.
- [25] K. Edwards & G. Farr (2012): *Improved Upper Bounds for Planarization and Series-Parallelization of Degree-Bounded Graphs*. *The Electronic Journal of Combinatorics* 19(2), p. #P25.
- [26] L. C. Eggan (1963): *Transition graphs and the star height of regular events*. *Michigan Mathematical Journal* 10, pp. 385–397, doi:10.1307/mmj/1028998975.
- [27] A. Ehrenfeucht & H. P. Zeiger (1976): *Complexity Measures for Regular Expressions*. *Journal of Computer and System Sciences* 12(2), pp. 134–146, doi:10.1016/S0022-0000(76)80034-7.
- [28] K. Ellul, B. Krawetz, J. Shallit & M.-W. Wang (2004): *Regular Expressions: New Results and Open Problems*. *Journal of Automata, Languages and Combinatorics* 9(2/3), pp. 233–256.

- [29] Ph. Flajolet & R. Sedgewick (2009): *Analytic Combinatorics*. Cambridge University Press, doi:10.1017/CB09780511801655.
- [30] Y. Gao, K. Salomaa & S. Yu (2011): *Transition Complexity of Incomplete DFAs*. *Fundamenta Informaticae* 110(1–4), pp. 143–158.
- [31] P. García, D. López, J. Ruiz & G. I. Álvarez (2011): *From regular expressions to smaller NFAs*. *Theoretical Computer Science* 412, pp. 5802–5807, doi:10.1016/j.tcs.2011.05.058.
- [32] P. Gawrychowski (2011): *Chrobak Normal Form Revisited, with Applications*. In B. Bouchou-Markhoff, P. Caron, J.-M. Champarnaud & D. Maurel, editors: *Proceedings of the 16th Conference on Implementation and Application of Automata, LNCS 6807*, Springer, Blois, France, pp. 142–153, doi:10.1007/978-3-642-22256-6_14.
- [33] V. Geffert (2003): *Translation of binary regular expressions into nondeterministic ϵ -free automata with $O(n \log n)$ transitions*. *Journal of Computer and System Sciences* 66(3), pp. 451–472, doi:10.1016/S0022-0000(03)00036-9.
- [34] W. Gelade (2010): *Succinctness of regular expressions with interleaving, intersection, and counting*. *Theoretical Computer Science* 411(31–33), pp. 2987–2998, doi:10.1016/j.tcs.2010.04.036.
- [35] W. Gelade & F. Neven (2008): *Succinctness of Complement and Intersection of Regular Expressions*. In S. Albers & P. Weil, editors: *Proceedings of the 25th International Symposium on Theoretical Aspects of Computer Science, Leibniz International Proceedings in Informatics 1*, Schloss Dagstuhl–Leibniz-Zentrum für Informatik, Dagstuhl, Germany, Bordeaux, France, pp. 325–336.
- [36] D. Giammarresi, J.-L. Ponty, D. Wood & D. Ziadi (2004): *A Characterization of Thompson Digraphs*. *Discrete Applied Mathematics* 134(1–3), pp. 317–337, doi:10.1016/S0166-218X(03)00299-3.
- [37] D. Giammarresi, J.-L. Pomy & D. Wood (1999): *Thompson Languages*. In J. Karhumäki, H. Maurer, G. Păun & G. Rozenberg, editors: *Jewels are Forever: Contributions on Theoretical Computer Science in Honor of Arto Salomaa*, Springer, pp. 16–24, doi:10.1007/978-3-642-60207-8_2.
- [38] V. M. Glushkov (1961): *The abstract theory of automata*. *Russian Mathematics Surveys* 16, pp. 1–53, doi:10.1070/RM1961v016n05ABEH004112.
- [39] J. Goldstine, M. Kappes, C. M. R. Kintala, H. Leung, A. Malcher & D. Wotschke (2002): *Descriptive Complexity of Machines with Limited Resources*. *Journal of Universal Computer Science* 8(2), pp. 193–234, doi:10.1142/9781848165458_0001.
- [40] J. Goldstine, C. M. R. Kintala & D. Wotschke (1990): *On Measuring Nondeterminism in Regular Languages*. *Information and Computation* 86(2), pp. 179–194, doi:10.1016/0890-5401(90)90053-K.
- [41] J. Goldstine, H. Leung & D. Wotschke (1992): *On the relation between ambiguity and nondeterminism in finite automata*. *Information and Computation* 100, pp. 261–170, doi:10.1016/0890-5401(92)90014-7.
- [42] H. Gruber & St. Gulan (2010): *Simplifying Regular Expressions*. In A. H. Dediu, H. Fernau & C. Martín-Vide, editors: *Proceedings of the 4th International Conference Language and Automata Theory and Applications, LNCS 6031*, Springer, Trier, Germany, pp. 285–296, doi:10.1007/978-3-642-13089-2_24.
- [43] H. Gruber & M. Holzer (2008): *Finite Automata, Digraph Connectivity, and Regular Expression Size*. In L. Aceto, I. Damgaard, L. A. Goldberg, M. M. Halldórsson, A. Ingólfssdóttir & I. Walkuwiewicz, editors: *Proceedings of the 35th International Colloquium on Automata, Languages and Programming, LNCS 5126*, Springer, Reykjavik, Iceland, pp. 39–50, doi:10.1007/978-3-540-70583-3_4.
- [44] H. Gruber & M. Holzer (2008): *Provably Shorter Regular Expressions from Deterministic Finite Automata (Extended Abstract)*. In M. Ito & M. Toyama, editors: *Proceedings of the 12th International Conference Developments in Language Theory, LNCS 5257*, Springer, Kyoto, Japan, pp. 383–395, doi:10.1007/978-3-540-85780-8_30.
- [45] H. Gruber & M. Holzer (2009): *Tight Bounds on the Descriptive Complexity of Regular Expressions*. In V. Diekert & D. Nowotka, editors: *Proceedings of the 13th International Conference Developments in Language Theory, LNCS 5583*, Springer, Stuttgart, Germany, pp. 276–287, doi:10.1007/978-3-642-02737-6_22.

- [46] H. Gruber & M. Holzer (2013): *Provably Shorter Regular Expressions From Finite Automata*. *International Journal of Foundations of Computer Science* 24(8), pp. 1255–1279, doi:10.1142/S0129054113500330.
- [47] H. Gruber & M. Holzer (2014): *Regular Expressions From Deterministic Finite Automata, Revisited*. IFIG Research Report 1403, Institut für Informatik, Justus-Liebig-Universität Gießen, Arndtstr. 2, D-35392 Gießen, Germany.
- [48] H. Gruber, M. Holzer & M. Tautschnig (2009): *Short Regular Expressions from Finite Automata: Empirical Results*. In S. Maneth, editor: *Proceedings of the 14th Conference on Implementation and Application of Automata, LNCS 5642*, Springer, Sydney, Australia, pp. 188–197, doi:10.1007/978-3-642-02979-0_22.
- [49] H. Gruber & J. Johannsen (2008): *Tight Bounds on the Descriptive Complexity of Regular Expressions*. In R. Amadio, editor: *Proceedings of the 11th Conference Foundations of Software Science and Computational Structures, LNCS 4962*, Springer, Budapest, Hungary, pp. 273–286, doi:10.1007/978-3-540-78499-9_20.
- [50] H. Gruber, J. Lee & J. Shallit (2012): *Enumerating regular expressions and their languages*. arXiv:1204.4982 [cs.FL].
- [51] St. Gulan & H. Fernau (2008): *An Optimal Construction of Finite Automata From Regular Expressions*. In R. Hariharan, M. Mukund & V. Vinay, editors: *Proceedings of the 28th Conference on Foundations of Software Technology and Theoretical Computer Science, Dagstuhl Seminar Proceedings 08002*, Internationales Begegnungs- und Forschungszentrum für Informatik (IBFI), Schloss Dagstuhl, Germany, Bangalore, India, pp. 211–222.
- [52] Ch. Hagenah & A. Muscholl (2000): *Computing ϵ -free NFA from regular expressions in $O(n \log^2(n))$ time*. *RAIRO–Informatique théorique et Applications / Theoretical Informatics and Applications* 34(5), pp. 257–277, doi:10.1051/ita:2000116.
- [53] Y.-S. Hand & D. Wood (2007): *Obtaining shorter regular expressions from finite-state automata*. *Theoretical Computer Science* 370(1–3), pp. 110–120, doi:10.1016/j.tcs.2006.09.025.
- [54] K. Hashiguchi (1988): *Algorithms for determining the relative star height and star height*. *Information and Computation* 78(2), pp. 124–169, doi:10.1016/0890-5401(88)90033-8.
- [55] M. Holzer & S. Jakobi (2011): *Chop Operations and Expressions: Descriptive Complexity Considerations*. In G. Mauri & A. Leporati, editors: *Proceedings of the 15th International Conference Developments in Language Theory, LNCS 6795*, Springer, Milan, Italy, pp. 264–275, doi:10.1007/978-3-642-22321-1_23.
- [56] M. Holzer & M. Kutrib (2009): *Nondeterministic Finite Automata—Recent Results on the Descriptive and Computational Complexity*. *International Journal of Foundations of Computer Science* 20(4), pp. 563–580, doi:10.1142/S0129054109006747.
- [57] M. Holzer & M. Kutrib (2010): *The Complexity of Regular(-Like) Expressions*. In Y. Gao, H. Lu, S. Seki & S. Yu, editors: *Proceedings of the 14th International Conference Developments in Language Theory, LNCS 6224*, Springer, London, Ontario, Canada, pp. 16–30, doi:10.1007/978-3-642-14455-4_3.
- [58] M. Holzer & M. Kutrib (2010): *Descriptive Complexity—An Introductory Survey*. In C. Martín-Vide, editor: *Scientific Applications of Language Methods*, World Scientific, pp. 1–58, doi:10.1142/9781848165458_0001.
- [59] M. Holzer & M. Kutrib (2010): *Descriptive Complexity of (Un)ambiguous Finite State Machines and Pushdown Automata*. In A. Kucera & I. Potapov, editors: *Proceedings of the 4th Workshop on Reachability Problems, LNCS 6227*, Springer, Brno, Czech Republic, pp. 1–23, doi:10.1007/978-3-642-15349-5_1.
- [60] M. Holzer & M. Kutrib (2011): *Descriptive and Computational Complexity of Finite Automata—A Survey*. *Information and Computation* 209(3), pp. 456–470, doi:10.1016/j.ic.2010.11.013.
- [61] J. E. Hopcroft & J. D. Ullman (1979): *Introduction to Automata Theory, Languages and Computation*. Addison-Wesley.

- [62] J. Hromkovič (2002): *Descriptive Complexity of Finite Automata: Concepts and Open Problems*. *Journal of Automata, Languages and Combinatorics* 7(4), pp. 519–531.
- [63] J. Hromkovič, J. Karhumäki, H. Klauck, G. Schnitger & S. Seibert (2002): *Communication complexity method for measuring nondeterminism in finite automata*. *Information and Computation* 172(2), pp. 202–217, doi:10.1006/inco.2001.3069.
- [64] J. Hromkovič & G. Schnitger (2005): *NFAs with and without ϵ -transitions*. In L. Caires, G. F. Italiano, L. Monteiro, C. Palamidessi & M. Yung, editors: *Proceedings of the 32nd International Colloquium Automata, Languages and Programming, LNCS 3580*, Springer, Lisbon, Portugal, pp. 385–396, doi:10.1007/11523468_32.
- [65] J. Hromkovič, S. Seibert & Th. Wilke (2001): *Translating Regular Expressions into Small ϵ -Free Automata*. *Journal of Computer and System Sciences* 62(4), pp. 565–588, doi:10.1006/jcss.2001.1748.
- [66] L. Ilie & S. Yu (2003): *Follow automata*. *Information and Computation* 186(1), pp. 140–162, doi:10.1016/S0890-5401(03)00090-7.
- [67] D. Kirsten (2005): *Distance desert automata and the star height problem*. *RAIRO—Informatique théorique et Applications / Theoretical Informatics and Applications* 39(3), pp. 455–509, doi:10.1051/ita:2005027.
- [68] S. C. Kleene (1956): *Representation of events in nerve nets and finite automata*. In C. E. Shannon & J. McCarthy, editors: *Automata studies, Annals of mathematics studies* 34, Princeton University Press, pp. 2–42.
- [69] E. Leiss (1981): *The complexity of restricted regular expressions and the synthesis problem for finite automata*. *Journal of Computer and System Sciences* 23(3), pp. 348–254, doi:10.1016/0022-0000(81)90070-2.
- [70] H. Leung (1998): *On Finite Automata with Limited Nondeterminism*. *Acta Informatica* 35(7), pp. 595–624, doi:10.1007/s002360050133.
- [71] H. Leung (1998): *Separating exponentially ambiguous finite automata from polynomially ambiguous finite automata*. *SIAM Journal on Computing* 27(4), pp. 1073–1082, doi:10.1137/S0097539793252092.
- [72] H. Leung (2005): *Descriptive complexity of NFA of different ambiguity*. *International Journal of Foundations of Computer Science* 16(5), pp. 975–984, doi:10.1142/S0129054105003418.
- [73] Y. Lifshits (2003): *A lower bound on the size of ϵ -free NFA corresponding to a regular expression*. *Information Processing Letters* 85(6), pp. 293–299, doi:10.1016/S0020-0190(02)00436-2.
- [74] S. Lombardy, Y. Régis-Gianas & J. Sakarovitch (2004): *Introducing VAUCANSON*. *Theoretical Computer Science* 328(1–2), pp. 77–96, doi:10.1016/j.tcs.2004.07.007.
- [75] E. Maia, N. Moreira & R. Reis (2013): *Incomplete Transition Complexity of Basic Operations on Finite Languages*. In S. Konstantinidis, editor: *Proceedings of the 18th International Conference on Implementation and Application of Automata, LNCS 7982*, Springer, Halifax, Nova Scotia, Canada, pp. 349–356, doi:10.1007/978-3-642-39274-0_31.
- [76] E. Maia, N. Moreira & R. Reis (2013): *Incomplete Transition Complexity of Some Basic Operations*. In P. v. Emde Boas, F. C. A. Groen, G. F. Italiano, J. R. Nawrocki & H. Sack, editors: *Proceedings of the 39th International Conference on Current Trends in Theory and Practice of Computer Science, LNCS 7741*, Springer, Špindlerøuv Mlýn, Czech Republic, pp. 319–331.
- [77] Z. Manna (1974): *Mathematical Theory of Computation*. McGraw-Hill.
- [78] A. Martinez (2002): *Efficient Computation of Regular Expressions from Unary NFAs*. In J. Dassow, M. Hobeberchts, H. Jürgensen & D. Wotschke, editors: *Pre-Proceedings of the 4th Workshop on Descriptive Complexity of Formal Systems, Report No. 586*, Department of Computer Science, The University of Western Ontario, Canada, London, Ontario, Canada, pp. 216–230.
- [79] H. V. McIntosh (1968): *REEX: A CONVERT Program to Realize the McNaughton-Yamada Analysis Algorithm*. Technical Report AIM-153, MIT Artificial Intelligence Laboratory.

- [80] R. McNaughton (1967): *The loop complexity of pure-group events*. *Information and Control* 11(1–2), pp. 167–176, doi:10.1016/S0019-9958(67)90481-0.
- [81] R. McNaughton (1969): *The loop complexity of regular events*. *Information Sciences* 1, pp. 305–328, doi:10.1016/S0020-0255(69)80016-2.
- [82] R. McNaughton (1982): *Elementary computability, formal languages, and automata*. Prentice-Hall.
- [83] Robert McNaughton & Hisao Yamada (1960): *Regular expressions and state graphs for automata*. *IRE Transactions on Electronic Computers* EC-9(1), pp. 39–47, doi:10.1109/TEC.1960.5221603.
- [84] A. R. Meyer & M. J. Fischer (1971): *Economy of description by automata, grammars, and formal systems*. In: *Proceedings of the 12th Annual Symposium on Switching and Automata Theory*, IEEE Computer Society Press, pp. 188–191, doi:10.1109/T-C.1971.223108.
- [85] B. G. Mirkin (1966): *An Algorithm for Constructing a Base in a Language of Regular Expressions*. *Engineering Cybernetics* 5, pp. 110–116.
- [86] F. R. Moore (1971): *On the bounds for state-set size in the proofs of equivalence between deterministic, nondeterministic, and two-way finite automata*. *IEEE Transaction on Computing* C-20, pp. 1211–1219, doi:10.1109/T-C.1971.223108.
- [87] N. Moreira, D. Nabais & R. Reis (2010): *State Elimination Ordering Strategies: Some Experimental Results*. In I. McQuillan & G. Pighizzini, editors: *Proceedings of the 12th Workshop on Descriptive Complexity of Formal Systems, EPTCS 31*, Saskatoon, Saskatchewan, Canada, pp. 139–148.
- [88] N. Moreira & R. Reis (2009): *Series-Parallel Automata and Short Regular Expressions*. *Fundamenta Informaticae* 91(3–4), pp. 611–629.
- [89] C. Nicaud (2009): *On the Average Size of Glushov’s Automaton*. In A. H. Dediu, A. M. Ionescu & C. Martín-Vide, editors: *Proceedings of the 3rd International Conference Language and Automata Theory and Applications, LNCS 5457*, Springer, Tarragona, Spain, pp. 626–637, doi:10.1007/978-3-642-00982-2_53.
- [90] C. Nicaud, C. Pivoteau & B. Razet (2010): *Average Analysis of Glushkov Automata under a BST-Like Model*. In K. Lodaya & M. Mahajan, editors: *Proceedings of the 30th IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science, Leibniz International Proceedings in Informatics 8*, Schloss Dagstuhl–Leibniz-Zentrum für Informatik, Dagstuhl, Germany, Chennai, India, pp. 388–399.
- [91] G. Ott & N. H. Feinstein (1961): *Design of Sequential Machines From Their Regular Expressions*. *Journal of the ACM* 8(4), pp. 585–600, doi:10.1145/321088.321097.
- [92] A. Palioudakis, K. Salomaa & S. Akl (2013): *Comparisons Between Measures of Nondeterminism on Finite Automata*. In H. Jürgensen & R. Reis, editors: *Proceedings of the 15th International Workshop Descriptive Complexity of Formal Systems, LNCS 8031*, Springer, London, Ontario, Canada, pp. 229–240, doi:10.1007/978-3-642-39310-5_21.
- [93] A. Palioudakis, K. Salomaa & S. G. Akl (2012): *State Complexity and Limited Nondeterminism*. In M. Kutrib, N. Moreira & R. Reis, editors: *Proceedings of the 14th Workshop on Descriptive Complexity of Formal Systems, LNCS 7386*, Springer, Braga, Portugal, pp. 252–265, doi:10.1007/978-3-642-31623-4_20.
- [94] M. O. Rabin & D. Scott (1959): *Finite Automata and Their Decision Problems*. *IBM Journal of Research and Development* 3, pp. 114–125, doi:10.1147/rd.32.0114.
- [95] B. Ravikumar & O. H. Ibarra (1989): *Relating the Type of Ambiguity of Finite Automata to the Succinctness of Their Representation*. *SIAM Journal on Computing* 18(6), pp. 1263–1282, doi:10.1137/0218083.
- [96] F. Reidl, P. Rossmanith, F. Sánchez Villamil & S. Sikdar (2014): *A Faster Parameterized Algorithm for Treedepth*. arXiv:1401.7540v3 [cs.DS].
- [97] J. Sakarovitch (2009): *Elements of Automata Theory*. Cambridge University Press, doi:10.1017/CB09781139195218.

- [98] G. Schnitger (2006): *Regular Expressions and NFAs Without ϵ -Transitions*. In B. Durand & W. Thomas, editors: *Proceedings of the 23th International Symposium on Theoretical Aspects of Computer Science*, LNCS 3884, Springer, Marseille, France, pp. 432–443.
- [99] S. Sippu & E. Soisalon-Soininen (1988): *Parsing Theory, Volume I: Languages and Parsing*. *EATCS Monographs on Theoretical Computer Science* 15, Springer, doi:10.1007/978-3-642-61345-6.
- [100] K. Thompson (1968): *Regular Expression Search Algorithm*. *Communications of the ACM* 11(6), pp. 419–422, doi:10.1145/363347.363387.
- [101] A. W. To (2009): *Unary finite automata vs. arithmetic progressions*. *Information Processing Letters* 109(17), pp. 1010–1014, doi:10.1016/j.ipl.2009.06.005.
- [102] B. W. Watson (1995): *Taxonomies and Toolkits of Regular Language Algorithms*. PhD thesis, Eindhoven University of Technology, Department of Mathematics and Computer Science, Den Dolech 2, 5612 AZ Eindhoven, The Netherlands.
- [103] S. Yu (2001): *State complexity of regular languages*. *Journal of Automata, Languages and Combinatorics* 6, pp. 221–234.