Regular Languages:
To Finite Automata and Beyond

Succinct Descriptions and Optimal Simulations

Luca Prigioniero ✉ 🏠 🆔

Dipartimento di Informatica, Università degli Studi di Milano, Milan, Italy

⎯ Abstract ⎯

It is well known that the class of regular languages coincides with the class of languages recognized by finite automata. Nevertheless, many other characterizations of this class in terms of computational devices and generative models are present in the literature. For example, by suitably restricting more powerful models such as context-free grammars, pushdown automata, and Turing machines, it is possible to obtain formal models that generate or recognize regular languages only. These restricted formalisms provide alternative representations of regular languages that may be significantly more concise than other models that share the same expressive power.

The goal of this work is to provide an overview of old and recent results on these formal systems from a descriptional complexity perspective, that is by considering the relationships between the sizes of such devices. We also present some results related to the investigation of the famous question posed by Sakoda and Sipser in 1978, concerning the size blowups from nondeterministic finite automata to two-way deterministic finite automata.

1 Introduction

The investigation of computational models operating under restrictions is one of classical topics of computer science.

In one of his pioneer papers, Chomsky introduced a hierarchy of classes of languages, also known as *Chomsky hierarchy*, obtained by applying increasing restrictions to general *grammars*, that characterize the class of *recursively enumerable* languages [7]. In this way he introduced the classes of *context-sensitive*, *context-free*, and *regular* languages.

For the same classes of languages, there also exist characterizations in terms of *computational devices*. Even in this case, bounding computational resources of general models, less powerful devices can be obtained. For example, while Turing machines (Tm for short), even in the variant with one tape only, characterize recursively enumerable languages, by restricting the working space they are allowed to use to the portion of the tape that initially contains the input, linear bounded automata are obtained, that are equivalent to context-sensitive grammars [33]. Also finite automata and pushdown automata (PDA), that are standard recognizers for regular and context-free languages, respectively, can be considered as particular Turing machines in which the capacity or access to the memory storage is limited.

Besides the standard models mentioned thus far, considering machines that make a restricted use of resources, it is possible to obtain alternative characterizations of the classes of the hierarchy. For example, Hennie proved that when the length of the computations, *i.e.*, the time, is linear in the input length, one-tape Turing machines are no more powerful than finite automata, that is, they recognize regular languages only [22].

As remarked by Chomsky, context-free languages have the property of being able to describe recursive structures such as, for instance, nested parentheses, arithmetic expressions, and typical programming language constructs. In terms of recognizing devices, this capability is typically implemented through the pushdown store, a memory structure in which the information is stored and recovered in a "last in–first out" way, which adds recursion to finite automata, so making the resulting model (pushdown automata), equivalent to context-free grammars (CFG) [8].
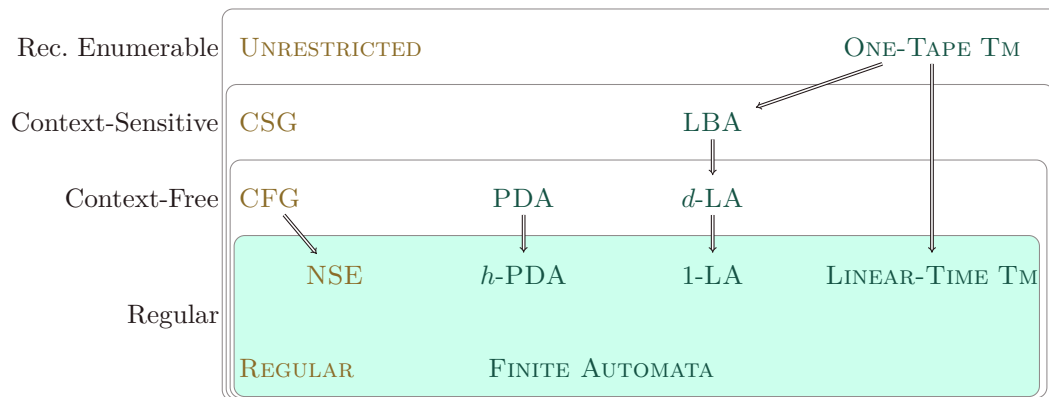
To emphasize the ability of context-free grammars to generate recursive sentential forms, Chomsky investigated the *self-embedding* property: a context-free grammar is self-embedding if it contains some variable which, in some sentential form, is able to reproduce itself enclosed between two nonempty strings [6]. Roughly speaking, this means that such a *self-embedded* variable can generate a "true" recursion that needs an auxiliary memory (typically a pushdown) to be implemented (in contrast with tail or head recursions, corresponding to the cases in which the two strings surrounding the variable are empty, that can be easily eliminated). Chomsky proved that, among all context-free grammars, only self-embedding ones can generate nonregular languages. Hence, *non-self-embedding grammars* (NSE) are no more powerful than finite automata [7, 6].

Counterpart devices for non-self-embedding grammars, for which the capability of recognizing recursive structures is limited by placing some restrictions on the size of the memory of the corresponding general model, are *constant-height pushdown automata* (h-PDA). More precisely, these devices are standard nondeterministic pushdown automata where the amount of available pushdown store is bounded by some constant $h \in \mathbb{N}$. Hence, the number of their possible configurations is finite, thus implying that they are no more powerful than finite automata.

By contrast to models that make use of space or time restrictions, Hibbard introduced *d-limited automata* (d-LA), that are obtained by limiting the writing capabilities of linear bounded automata allowing overwriting of each tape cell only the first d times that it is scanned, for some fixed $d \geq 0$ [23]. Nevertheless, the cell may be visited again and the information stored therein read arbitrarily many more times, but its contents is *frozen* for the remainder of the computation. Hibbard proved that, for each $d \geq 2$ this model characterizes context-free languages and showed the existence of an infinite hierarchy of deterministic d-limited automata, whose first level (*i.e.*, corresponding to deterministic 2-limited automata) has been later proved to coincide with the class of *deterministic context-free languages* [43]. (See [34] and references therein for further connections between limited automata and context-free languages.) Furthermore, as shown by Wagner and Wechsung, when $d = 1$, that is, when these devices are allowed to overwrite the contents of each tape cell during the first visit only, 1-limited automata (1-LA) are equivalent to finite automata [56]. It is a trivial observation that even when $d = 0$, and hence no writings on the tape are allowed, the model still recognizes regular languages only: 0-limited automata correspond to finite automata that can move their input head back and forth on the input tape, namely *two-way finite automata*.[1]

---

[1] Further technical details, properties, examples, and references about limited automata can be found in the recent survey by Pighizzini [41].

General devices and their restrictions characterizing regular languages and discussed in this work are depicted in Figure 1.

Rec. Enumerable	UNRESTRICTED			ONE-TAPE TM
Context-Sensitive	CSG		LBA	
Context-Free	CFG	PDA	d-LA	
	NSE	h-PDA	1-LA	LINEAR-TIME TM
Regular	REGULAR	FINITE AUTOMATA		

Figure 1 Models representing regular languages obtained by posing some restrictions on standard grammars (on the left, in gold) and recognizers (on the right, in green) characterizing the classes of the Chomsky hierarchy.

Descriptional Complexity of Formal Systems

In this work we shall focus on the models characterizing the bottom level of the Chomsky hierarchy: the class of regular languages. We give an overview of some recent results obtained in the area of descriptional complexity, specifically, the study of the capability of these devices of representing the same class of languages in a more, or less, concise way. More precisely, *descriptional complexity* is a branch of theoretical computer science whose goal is the investigation of the relationships between the sizes of the representations of formal systems that share the same computational power, or, in other words, the study of how concisely a system can describe a class of problems (or languages).

The devices characterizing the class of regular languages introduced at the beginning of the section are obtained by limiting some resources of more general models. Therefore, the main question we are interested about is *"How much does the limitation of one resource cost in terms of another resource, or, in other words, what are the upper and lower bounds of such trade-offs?"* [14].

Since our goal is comparing the sizes of the descriptions of devices and formal systems, for each model under consideration we evaluate its *size* as the total number of symbols used to describe it, or, in other words, to write down its description. In particular, to measure the size of recognizing devices, we consider the amount of memory required to store the "*algorithm*" they implement (their transition function), so, for example, the size of finite automata is bounded by a polynomial in the number of their states. On the other hand, the size of grammars is given by the number of symbols used to write down their derivation rules [32].

Given two different classes of computational models $\mathcal{M}_1$ and $\mathcal{M}_2$ characterizing regular languages, there are natural questions we are interested in:

- Does a function F exist, such that for all the regular languages L, the size of the smallest device of type $\mathcal{M}_1$ for L is upper bounded by the function F of the size of the smallest equivalent device of type $\mathcal{M}_2$? If F exists, it is an *upper bound* for the increase (*blow-up*)

in complexity when changing from a minimal model of type $\mathcal{M}_2$ for an arbitrary regular language to an equivalent minimal model of type $\mathcal{M}_1$.

- Do an infinite family of distinct regular languages L_i and a function f exist, such that for all $i \in \mathbb{N}$, the size of any device of type $\mathcal{M}_1$ for L_i is lower bounded by the function f of the size of the smallest equivalent device of type $\mathcal{M}_2$? If f exists, it is a *lower bound* for the increase in complexity when changing from a minimal model of type $\mathcal{M}_2$ to an equivalent minimal model of type $\mathcal{M}_1$ for infinitely many languages.

If there is no recursive function upper bounding the trade-off between two computational models $\mathcal{M}_1$ and $\mathcal{M}_2$, the trade-off is *non-recursive*. Furthermore, if the lower and the upper bounds coincide, we say that the bound is *optimal*. For more details about the area of descriptional complexity see, *e.g.*, the surveys by Goldstine et al., and Holzer and Kutrib [14, 24].

A classical problem in this field is the investigation of the relationships between deterministic and nondeterministic devices. It is well known that *one-way deterministic* finite automata (1DFA) are sufficient for regular languages. By allowing nondeterministic transitions, thus obtaining one-way nondeterministic finite automata (1NFA), the computational power does not increase [51]. A natural question concerning models that share the same computational power is the comparison of their size.

As an example, consider, for any fixed integer $n \geq 0$, the language $L_n = (0+1)^*1(0+1)^{n-1}$, composed by strings on the alphabet $\{0, 1\}$ whose n-th to last symbol is 1. A 1NFA $\mathcal{A}_n$ could recognize L_n by moving the input head forward on the input tape, until reaching the n-th symbol from the end, that is detected by performing a nondeterministic guess. If such a symbol is 1, the automaton checks whether the length of the remaining part of the input string is $n-1$. Therefore, it is easy to check that the number of states for a 1NFA accepting L_n is $n+1$. On the other hand, a 1DFA accepting L_n cannot guess which is the n-th symbol from the end of the input string. So, intuitively, it has to "*remember*" a factor representing the last n symbols read, by storing it in the finite control. When the end of the input is reached, the device checks that the leftmost symbol of the current factor is 1. It is easy to see that the number of the possible factors of length n is 2^n, and each of them is stored by using a state of the 1DFA, thus implying an exponential blowup in states with respect to the equivalent 1NFA $\mathcal{A}_n$.

Therefore, even if deterministic and nondeterministic finite automata characterize the same class of languages, one-way deterministic automata can require exponentially many states with respect to equivalent nondeterministic automata. Hence, there is an exponential size gap from one-way nondeterministic to one-way deterministic automata.

It is also known that even providing finite automata with the capability of moving the input head back and forth along the tape, thus obtaining *two-way* finite automata (2DFA and 2NFA), their recognizing power does not increase. In fact, Shepherdson and, independently, Rabin and Scott, proved this result by giving two constructions based on the analysis of the moves of the input head of the automata between the tape cells [51, 53]. Both constructions, given a two-way finite automaton, return an equivalent one-way deterministic finite automaton whose size is exponential in the square of the size of the automaton given in input.

At this point, one could ask *"Is it possible to exploit the ability to scan the input in a two-way fashion in finite automata to eliminate the nondeterminism?"*. Differently from the one-way case, the question about the size cost of the conversion of (one-way and two-way) nondeterministic finite automata into two-way deterministic finite automata, that was posed by Sakoda and Sipser, is open since 1978. In their paper, Sakoda and Sipser formulated the problem by the questions

1. For every 2NFA $\mathcal{A}$, does an equivalent 2DFA $\mathcal{A}'$ exist, such that $\mathcal{A}'$ has only polynomially more states than $\mathcal{A}$?

2. For every 1NFA $\mathcal{A}$, does an equivalent 2DFA $\mathcal{A}'$ exist, such that $\mathcal{A}'$ has only polynomially more states than $\mathcal{A}$?

They conjectured that, in both cases, the answer is negative and an exponential increase of the size is necessary.

The question has been solved in some special cases that can be grouped in three classes: by considering restrictions on the simulating machines (*e.g.*, *sweeping* [54], *oblivious* [27], *few reversals* two-way deterministic finite automata [28]), by considering restrictions on languages (*e.g.*, *unary case* [12]), by considering restrictions on the simulated machines (*e.g.*, *outer-nondeterministic* automata [10, 30]). However, in spite of all attempts, in the general case the question remains open.

The importance of this open problem is supported by its similarity to the well-known $P \overset{?}{=} NP$ problem [52] and by relationships with the LOGSPACE $\overset{?}{=}$ NLOGSPACE question [5, 13, 29, 30]. (See [40] for further details and references.)

Outline

This work is an excerpt of the dissertation [48] and an informal version appeared in [49]. It is organized as follows (a summary of some of the main results obtained is depicted in Figure 2).
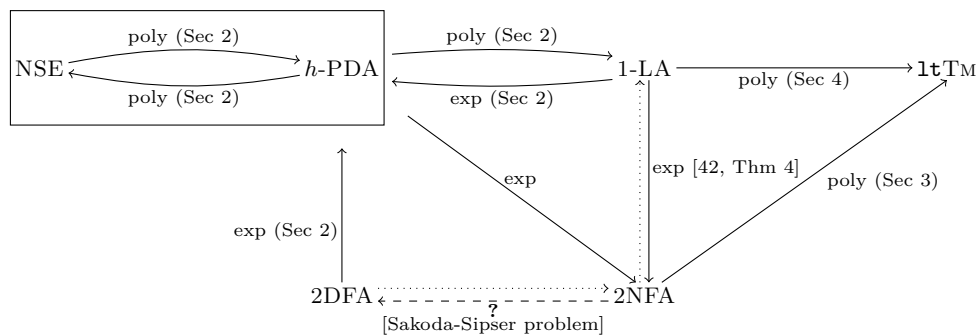

■ **Figure 2** Some of the main results discussed in the work. Dotted arrows denote trivial relationships, while the dashed arrow indicates the Sakoda and Sipser's question [52]. Linear-time Turing machines are denoted ltTM.

Section 2 presents some results on non-self-embedding grammars and their descriptional complexity [46], and their relations in size with constant-height pushdown automata and 1-limited automata [15].

Section 3 investigates machines one-tape Turing machines working in linear time. It is not decidable whether a one-tape Turing machine works in linear time, even if it is deterministic and restricted to use only the portion of the tape which initially contains the input, unless the machine is *weight-reducing*, *i.e.*, syntactically forced to operate in linear-time [16]. By relating the study of Turing machines working in linear time to the above-mentioned open question of Sakoda and Sipser, we present the costs of the conversion of nondeterministic finite automata into equivalent linear-time one-tape deterministic machines. A polynomial blowup from two-way nondeterministic finite automata into equivalent weight-reducing one-tape deterministic machines (that work in linear time) has been obtained.

The investigation about devices working in linear time is continued in Section 4, where the time complexity of 1-limited automata is analyzed from a descriptional complexity point of view [19]. Though the model recognizes regular languages only, it may use quadratic time in the input length. With a polynomial increase in size and preserving determinism, each 1-limited automaton can be transformed into a linear-time equivalent one. We also obtained polynomial transformations into linear-time Turing machines.

Section 5 turns the attention to pushdown automata [45]. In particular, we studied pushdown automata without any restriction on the pushdown height and we showed that it cannot be decided whether these devices accept using constant pushdown height, with respect to the input length, or not. Furthermore, in the case of acceptance in constant height, the height cannot be bounded by any recursive function in the size of the description of the machine. In contrast, in the restricted case of pushdown automata over a one-letter input alphabet, *i.e.*, unary pushdown automata, the above property becomes decidable. Moreover, if the height is bounded by a constant that does not depend on the input length, then it is at most exponential with respect to the size of the description of the pushdown automaton. This bound cannot be reduced. Finally, if a unary pushdown automaton uses non-constant height to accept, then the height should grow at least as the logarithm of the input length. This bound is optimal.

In conclusion, in Section 6 we briefly discuss some possible future research directions and open problems.

2   Non-Self-Embedding Grammars, Constant-Height Pushdown Automata, and Limited Automata

As mentioned in Section 1, Chomsky investigated the self-embedding property of context-free grammars, and proved that non-self-embedding grammars only generate regular languages. The proof of the result given by Chomsky is constructive: it provides a method for obtaining a finite automaton equivalent to a given non-self-embedding grammar [7, 6]. A different constructive proof of the same result was given by Anselmo, Giammarresi, and Varricchio, who showed that it is possible to decompose non-self-embedding grammars into regular grammars and then to iteratively apply regular substitutions in order to obtain equivalent finite automata [2]. In the same paper, the authors also presented an NSE grammar for the following family of languages. For any fixed integer n, let $U_n = \{\, a^{2^n} \,\}$ be the singleton composed by the unary string of length 2^n. U_n can be generated by an NSE grammar with variables $\{\, A_0, \ldots, A_n \,\}$ such that, for each $i = 1, \ldots, n$, the variable A_i generates two occurrences of the variable A_{i-1}, and the variable A_0 produces the terminal a. It is possible to see that the variable A_i derives the string a^{2^i}, for $i = 0, \ldots, n$. Hence, the language generated is U_n. Moreover, the size of the grammar is linear in the parameter n, while it is not hard to verify that the equivalent minimum deterministic finite automaton has $2^n + 1$ states. Actually, as a consequence of state lower bound presented by Mereghetti and Pighizzini, the same amount of states is also necessary to accept U_n by using a 2NFA [36]. Therefore, this language witnesses that the size gap between non-self-embedding grammars and equivalent finite automata is at least exponential.

It is worthwhile to mention that, as proved by Meyer and Fischer, for any recursive function f and arbitrarily large integer n, there exists a context-free grammar whose description has size n, the generated language is regular, and any equivalent finite automaton requires at least $f(n)$ states [37]. This means that it is not possible to obtain a recursive bound relating the size of context-free grammars generating regular languages with the number of states of equivalent finite automata. It is important to notice that the result by Meyer and Fischer

was obtained by considering grammars with a two-letter terminal alphabet. The unary case was studied by Pighizzini, Shallit, and Wang, who obtained optimal recursive bounds [47].

We recently proved that also in the case of non-self-embedding grammars the bounds are recursive, independently on the alphabet size [46]. In particular, by inspecting and refining the construction presented by Anselmo, Giammarresi, and Varricchio, we showed that

▶ **Result 2.1** ([46]). *Each non-self-embedding grammar of size n can be converted into equivalent nondeterministic and deterministic automata with $2^{O(n)}$ and $2^{2^{O(n)}}$ states, respectively.*

We also obtained a family of languages that witnesses that these gaps cannot be reduced. Furthermore, these gaps do not change if we allow the variables generating only unary strings (*i.e.*, strings consisting of occurrences of only one terminal) to be self-embedded.

Other formal models characterizing the class of regular languages and exhibiting gaps of the same order with respect to deterministic and nondeterministic automata are constant-height pushdown automata and 1-limited automata. So, it is natural to study the size relationships between non-self-embedding grammars, constant-height pushdown automata, and 1-limited automata, three models that restrict context-free acceptors to the level of regular recognizers.

The exponential and double exponential gaps from constant-height pushdown automata to nondeterministic and deterministic automata have been proven by Geffert, Mereghetti, and Palano [11]. Furthermore, Bednárová et al. showed the interesting result that the gap from nondeterministic to deterministic constant-height pushdown automata is double exponential also [3]. We remind the reader that both non-self-embedding grammars and constant-height pushdown automata are restrictions of the corresponding general models, where true recursions are not possible. By adapting the standard transformation from PDA to CFGs, and by modifying the decomposition of NSE grammars presented by Anselmo, Giammarresi, and Varricchio, we proved that

▶ **Result 2.2** ([15]). *Non-self-embedding grammars and constant-height pushdown automata are polynomially related in size.*

Also 1-limited automata can be significantly smaller than equivalent finite automata. The equivalence between 1-limited automata and finite automata has been investigated from the descriptional complexity point of view by Pighizzini and Pisoni, who proved that each 1-limited automaton $\mathcal{A}$ with n states can be simulated by a one-way deterministic automaton with a number of states doubly exponential in a polynomial in n. Furthermore, in the worst case, double exponentially many states are necessary for this simulation. The cost reduces to a single exponential when $\mathcal{A}$ is deterministic [42]. The lower bounds in this result have been obtained by providing witness languages defined over a binary alphabet.

We investigated the unary case, for which it was an open question whether the same bounds held. In particular, we obtained that

▶ **Result 2.3** ([44]). *There is an exponential gap between unary deterministic 1-limited automata and two-way nondeterministic finite automata.*

To this end, we showed that, for each $n > 1$, the singleton language $U_n = \{\, a^{2^n} \,\}$ can be recognized by a deterministic 1-limited automaton having $2n + 1$ states and a description of size $O(n^2)$. Since the same language requires $2^n + 1$ states to be accepted by a one-way nondeterministic automaton, it turns out that the state gap between deterministic 1-limited automata and one-way nondeterministic automata in the unary case is the same as in the binary case. It is an easy observation that the gap does not reduce if we want to convert unary deterministic 1-limited automata into two-way nondeterministic automata.

We also considered the relationships between 1-limited automata and unary context-free grammars. The cost of the conversion of these grammars into finite automata has been investigated: exponential gaps have been proven by Pighizzini, Shallit, and Wang [47]. With the help of a result presented by Okhotin [38], we proved that

▶ **Result 2.4** ([44]). *Each unary context-free grammar $\mathcal{G}$ can be converted into an equivalent* 1-*limited automaton whose description has a size that is polynomial in the size of $\mathcal{G}$.*

Let us now turn our attention to the size relationships between 1-limited automata and non-self-embedding grammars.

We obtained a construction transforming each non-self-embedding grammar into a 1-limited automaton of polynomial size. In particular, given an input w, the 1-LA nondeterministically generates a *compression* of a derivation tree of w. Then, it verifies the validity of such a guess. As a consequence,

▶ **Result 2.5** ([15]). *Each constant-height pushdown automaton can be transformed into an equivalent* 1-*limited automaton of polynomial size.*

Even the conversion of deterministic constant-height pushdown automata into deterministic 1-limited automata costs polynomial in size. For the converse transformation, we showed that an exponential size is necessary. Indeed, consider the following family of languages. For any fixed integer $n > 0$, let P_n be the language of the powers of any string u of length n over the alphabet $\{0,1\}$, *i.e.*, $P_n = \{ u^k \mid u \in \{0,1\}^n, k \geq 0 \}$. P_n can be accepted by a two-way deterministic finite automaton (and, hence, by a 1-LA) with $O(n)$ states, but requires exponentially many states to be accepted even by an unrestricted pushdown automaton, which is forced to store the word u in its finite control. From the cost of the conversion of 1-limited automata into nondeterministic automata, it turns out that for the conversion of 1-limited automata into non-self-embedding grammars an exponential size is also sufficient.

▶ **Result 2.6** ([15]). *The size cost of the conversion of* 1-*limited automata into non-self-embedding grammars and constant-height pushdown automata is exponential.*

Bednárová et al. raised the question of the cost of the conversion of deterministic h-PDA into 1NFAs [3]. To this regard, it is possible to observe that, for any integer $n \geq 0$, the language $S_n = \left(a^{2^n} \right)^*$ is accepted by a deterministic h-PDA of size polynomial in n (a constant-height pushdown automaton for S_n with a constant number of states, $h = n$, and $2n+1$ pushdown symbols can be found in [45]) but, as discussed at the beginning of this section, it requires an exponential number of states to be accepted by a finite automaton. Hence, we can conclude that the cost of both simulations from two-way automata to h-PDA and from h-PDA to two-way automata is at least exponential.

More details about the results presented in this section can be found in [15, 44, 46].

3   Two-Way Automata and One-Tape Machines

In this section we present some results on one-tape Turing machines working in linear time and on the relationships between the sizes of their descriptions with those of equivalent finite automata. We consider the following variants of one-tape *deterministic* Turing machines:

End-marked machines. We say that a device is *end-marked* if each input string is surrounded by two special symbols called the left and the right endmarkers and the machine is restricted to use only the portion of the tape that initially contains the input (plus the endmarkers), that cannot be left by the head.

Weight-reducing Turing machines. Roughly speaking, in *weight-reducing* Turing machines each overwriting is decreasing with respect to some fixed order on the working alphabet. As a consequence, after overwriting a cell with a minimal symbol, such a machine cannot visit the cell again. By this condition, in weight-reducing Turing machines the number of visits to each tape cell is bounded by a constant. However, weight-reducing machines could have non-halting computations which, hence, visit infinitely many tape cells.

Linear-time Turing machines. A deterministic Turing machine is said to be *linear-time* if, over each input w, its computation halts within $O(|w|)$ steps.

Hennie machines. A Hennie machine is a linear-time Turing machine which is, furthermore, end-marked.

Weight-Reducing Hennie machines. These devices are defined by combining previous conditions. Observe that each end-marked weight-reducing Turing machine can execute a number of steps which is at most linear in the length of the input. Consequently, end-marked weight-reducing deterministic Turing machines are necessarily Hennie machines.

It is useful to recall that it cannot be decided whether or not a one-tape Turing machine works in linear time.[2] Furthermore, there is no recursive function bounding the size blowup from one-tape Turing machines working in linear time to equivalent finite automata. We proved that these results remain true even in the restricted case of end-marked machines.

▶ **Result 3.1** ([16, 18]). *It is undecidable whether an end-marked Turing machine works in linear time.*

To overcome the above-mentioned "negative" results, we considered weight-reducing machines, that can be seen as a syntactical restriction of one-tape Turing machines. In this case, it can be decided whether a deterministic Turing machine is weight reducing. These devices can have non-halting computations. However, they work in linear time as soon as they are halting. In fact, we showed that it is possible to decide whether a weight-reducing machine is halting. As a consequence,

▶ **Result 3.2** ([16, 18]). *It is possible to decide whether a weight-reducing machine works in linear time.*

Moreover, by a polynomial size increase, each weight-reducing machine can be transformed into an equivalent one which always halts and works in linear time. The same blowup is easily extended to weight-reducing Hennie machines.

▶ **Result 3.3** ([16, 18]). *With a polynomial size increase, any weight-reducing machine can be made halting, so working in linear time. Furthermore, the cost of the simulation of weight-reducing machines by* 1DFA *is doubly exponential.*

Furthermore, we proved that

▶ **Result 3.4** ([16, 18]). *Each linear-time machine $\mathcal{T}$ can be converted into an equivalent weight-reducing one whose size is bounded by a polynomial function in the size and in the execution time of $\mathcal{T}$.*

We also related the study of these restricted variants of Turing machines accepting only regular languages to the Sakoda and Sipser question. In this case, we propose a new approach

---

[2] We mention that it is decidable, though, whether or not a machine makes at most $cm + d$ steps on input of length m, for any *fixed* $c, d > 0$ [9].

in which we considered, as target devices, linear-time one-tape *deterministic* Turing machines. It turned out that each 2NFA $\mathcal{A}$ can be simulated by a one-tape deterministic Turing machine which works in linear time (with respect to the input length) and which has a polynomial size with respect to the size of $\mathcal{A}$. We point out that the resulting machine can use extra space, besides the tape segment which initially contained the input. Next, the machine is halting and weight reducing, thus implying a linear execution time. Hence,

▶ **Result 3.5** ([16, 17]). *Nondeterminism can be eliminated from* 2NFA *with at most a polynomial size increase, obtaining a linear execution time in the input length, provided the ability to rewrite tape cells and to use some extra space.*

We then investigated what happens by removing the latter possibility, namely if the machine does not have any further tape storage, *i.e.*, it is a Hennie machine. We proved that even under this restriction it is still possible to obtain a machine of polynomial size, namely

▶ **Result 3.6** ([16, 17]). *Each* 2NFA *can be transformed into an equivalent Hennie machine of polynomial size.*

However, the machine resulting from our construction is not weight reducing, unless we require that it agrees with the given 2NFA only on sufficiently long inputs. This problem does not occur in the unary case, where we proved that

▶ **Result 3.7** ([16, 17]). *Each unary* 2NFA *can be simulated by a weight-reducing Hennie machine of polynomial size.*

Similar results have been obtained for the transformation of 1NFAs into variants of one-tape deterministic machines.

For details about the results discussed in this section we address the reader to [16, 17, 18].

4   Limited Automata: A Time Constraint

As proved by Hennie, deterministic one-tape Turing machines operating in linear time recognize exactly the class of regular languages [22]. The result has later been extended to the nondeterministic case by Pighizzini and Tadaki, Yamakami, and Lin [39, 55]. Here, *operating in linear time* means that every computation has length linearly bounded in the input length. In particular, linear-time machines are necessarily halting – see [39] for investigations of alternative linear-time restrictions. The above-mentioned result implies that every Hennie machine is equivalent to some finite automaton. From the opposite point of view, this means that providing two-way finite automata with the ability to overwrite the tape cells does not extend the expressiveness of the model, as long as the time is linearly bounded in the length of the input.

However, it is undecidable, given an end-marked deterministic Turing machine (*i.e.*, a deterministic linear bounded automaton), to check whether it works in linear time over all input strings, or, in other words, whether it is actually a deterministic Hennie machine [16, 17, 50]. To avoid this drawback, Průša proposed the weight-reducing variant of Hennie machines, in which the time limitation is syntactic. As a consequence, the number of visits of a cell by the head is bounded by some constant (*i.e.*, not depending on the input length), hence the device works in linear time over every input string.

By contrast to Hennie machines, in d-limited automata the head is allowed to visit a cell after the d-th visit, even if it cannot rewrite the contents anymore. This allows to use super-linear time.

As an example let us consider, for each fixed integer $n \geq 0$, the language

$$Q_n = \{ x_0 x_1 \cdots x_k \mid k > 0, \text{ for each } i\colon x_i \in \Sigma^n \text{ and for some } j \neq 0\colon x_j = x_0 \}$$

on the alphabet $\Sigma = \{0, 1\}$. A deterministic 1-LA $\mathcal{A}_n$ may recognize Q_n as follows. It first scans the factor x_0, overwriting each input symbol with a marked copy. Then, $\mathcal{A}_n$ repeats a subroutine which overwrites a factor x_j with the marker $\sharp \notin \Sigma$, while checking whether x_j equals x_0 or not. This can be achieved as follows. Before overwriting the ℓ-th symbol of x_j, first, $\mathcal{A}_n$, with the help of a counter modulo n, moves the head leftward to the position ℓ of x_0 and stores the unmarked scanned symbol σ in its finite control; second, it moves the head rightward until reaching the position ℓ of x_j, namely, the leftmost position that has not been overwritten so far. At this point, $\mathcal{A}_n$ compares the scanned symbol (*i.e.*, the ℓ-th symbol of x_j) with σ (*i.e.*, the ℓ-th symbol of x_0). When $\mathcal{A}_n$ finds out that a complete factor x_j matches x_0, it reaches the end of the input checking that has length multiple of n.

It is possible to implement $\mathcal{A}_n$ with a number of states linear in n and $\#\Sigma+1$ working symbols. Since for each position of a factor x_i, $i > 0$, the head has to move back to the factor x_0, it is possible to observe that $\mathcal{A}_n$ works in quadratic time in the length of the input string.

Therefore, also in the case $d = 1$, d-limited automata can operate in super-linear time. This contrasts with Hennie machines which operate in linear time by definition. The question we addressed is whether this ability of 1-limited automata with respect to Hennie machines yields a gap between the two models in terms of the size of their representations.

We proved that operating in super-linear time is not essential for 1-LA, if allowing a polynomial increase in the number of states. In other words,

▶ **Result 4.1** ([19]). *With a polynomial increase in size, each* 1*-limited automaton can be transformed into an equivalent linear-time* 1*-limited automaton, or, alternatively, into a weight-reducing Hennie machine.*

Furthermore, the obtained device is deterministic when the original machine is deterministic as well. This is achieved by extending the exponential-cost simulation of 1-LA by 2NFA given in [42] (which in turn extends Shepherdson's classic conversion of 2DFA into 1DFA [53]) with a method for storing and accessing a carefully chosen subcollection of the many "Shepherdson tables" that a 1NFA would need to remember in its states: the simulating automaton can both store the tables (despite the 1-limitation) and access them efficiently (in both size and time).

For details about the results discussed in this section please refer to [19].

5 Pushdown Automata and Space Restrictions

As discussed in Section 2, constant-height pushdown automata, that are pushdown automata in which the maximum height of the pushdown is limited by some constant, allow more succinct representations of regular languages than finite automata [11], and are polynomially related in size with their natural generative counterpart, NSE grammars, roughly context-free grammars without "true" recursion [7].

In this section we turn our attention to standard pushdown automata, namely with an unrestricted pushdown store, that, however, are able to accept their inputs by making use only of a constant amount of the pushdown store. More precisely, we say that a pushdown automaton $\mathcal{M}$ *accepts in constant height* h, for some given h, if for each word in the language accepted by $\mathcal{M}$ there exists one accepting computation in which the maximum height reached

by the store is bounded by h. Notice that this does not prevent the existence of accepting or rejecting computations using an unbounded pushdown height.

It is a simple observation that a pushdown automaton $\mathcal{M}$ accepting in constant height h can be converted into an equivalent constant-height pushdown automaton: in any configuration it is enough to keep track of the current height in order to stop and reject when a computation tries to exceed the height limit. The description of the resulting constant-height pushdown automaton has size polynomial in h and in the size of the description of $\mathcal{M}$.

While studying these size relationships, we tried to understand *how large the height h of the pushdown can be with respect to the size of the description of $\mathcal{M}$*. We discovered that h can be arbitrarily large. Indeed, adapting an argument presented by Meyer and Fischer to prove non recursive trade-offs between the size of PDA accepting regular languages and the number of states of equivalent automata [37], we showed that

▶ **Result 5.1** ([45]). *There is no recursive function bounding the maximal height reached by the pushdown store in a pushdown automaton accepting in constant height, with respect to the size of its description.*

With the same argument, we obtained that

▶ **Result 5.2** ([45]). *There is no recursive function bounding the size blowup from PDA accepting in constant height to finite automata.*

Moreover, using a technique introduced by Hartmanis, based on suitable encodings of single-tape Turing machine computations [21], we also proved that

▶ **Result 5.3** ([45]). *It cannot be decided whether a pushdown automaton accepts in constant height or not.*

We also investigated the unary case. By studying the structure of the computations of unary pushdown automata, we were able to prove that, in contrast to the general case,

▶ **Result 5.4** ([45]). *It can be decided whether or not unary pushdown automata accept in constant height.*

Furthermore, we proved that

▶ **Result 5.5** ([45]). *If a unary pushdown automaton $\mathcal{M}$ accepts in height h, constant with respect to the input length, then h is bounded by an exponential function in the size of $\mathcal{M}$.*

By presenting a suitable family of pushdown automata, we showed that this bound is optimal.

Let us turn our attention to pushdown automata that accept using height which is not constant in the input length, in order to investigate how the pushdown height grows. In particular, we asked if there exists a minimum growth of the pushdown height, with respect to the length of the input, when it is not constant. The answer to this question is already known and it derives from results on Turing machines: the height of the store should grow at least as a double logarithmic function [1]. This lower bound cannot be increased, because a matching upper bound has been recently obtained in [4]. As a consequence of the constructions obtained for automata accepting unary languages, we were able to prove that

▶ **Result 5.6** ([45]). *In the unary case, the height of the store should grow at least as a logarithmic function, and this lower bound cannot be further increased.*

For more details about the results discussed in this section we refer the reader to [45].

6 Future Work and Open Problems

In this work we summarized some old and recent results related to the area of descriptional complexity, and, in this regard, we focused on the class of regular languages. To conclude, we discuss possible directions for future research in this field.

First of all, it is worth mentioning that there exist other models characterizing the class of regular languages besides the ones analyzed here. One example are the well-known *regular expressions*, widely discussed in classical textbooks (see, *e.g.* [26]). From regular expressions we can derive a more succinct representation of regular languages, by using *straight-line programs*, namely programs representing directed acyclic graphs, whose internal nodes correspond to the basic regular operations (*i.e.*, union, concatenation, and star). Descriptional complexity of straight-line programs has been analyzed. It has been proved that straight-line programs are polynomially related in size with constant height pushdown automata [11]. Anyways, it would be interesting to deepen the study of descriptional complexity of models *"derived"* from regular expressions.

As widely discussed, the question posed by Sakoda and Sipser in 1978 about the elimination of nondeterminism from finite automata using the two-way motion is still open. We plan to continue the investigation on this question by considering models that have the same computational power as finite automata and by studying the relations in sizes between these nondeterministic devices and deterministic machines. For example, as remarked by Pighizzini in a recent survey, at the moment, direct simulations of 1-limited automata by *deterministic* 1-limited automata and by two-way deterministic automata are not known [41]. It could be interesting to know if, when simulating unary and non-unary 1-limited automata by two-way (instead of one-way) deterministic finite automata, the cost reduces from a double exponential to a simple exponential.

It could be also interesting to study "relaxed" versions of the problem of Sakoda and Sipser, in which the simulating machine is a deterministic 1-limited automaton (*i.e.*, a deterministic two-way automaton with the capability of rewriting the contents of tape cells during the first visit).

Moreover, following the research line started in [10], one could deepen the investigation on the Sakoda and Sipser problem in case of simulated devices that perform a limited use of nondeterminism. In this regard, it is possible to consider several restrictions like, for example, on the number of nondeterministic choices along the computation (see, *e.g.* [25]), or on the number of total, or accepting, computations (also known as *degree of ambiguity* [20, 31, 35]).

───── **References** ─────

1   Maris Alberts. Space complexity of alternating Turing machines. In *Fundamentals of Computation Theory (FCT) '85*, volume 199 of *Lecture Notes in Computer Science*, pages 1–7. Springer, 1985. `doi:10.1007/BFb0028785`.

2   Marcella Anselmo, Dora Giammarresi, and Stefano Varricchio. Finite automata and non-self-embedding grammars. In *Conference on Implementation and Application of Automata (CIAA) 2002*, volume 2608 of *Lecture Notes in Computer Science*, pages 47–56. Springer, 2002. `doi:10.1007/3-540-44977-9_4`.

3   Zuzana Bednárová, Viliam Geffert, Carlo Mereghetti, and Beatrice Palano. Removing nondeterminism in constant height pushdown automata. *Information and Computation*, 237:257–267, 2014. `doi:10.1016/j.ic.2014.03.002`.

4   Zuzana Bednárová, Viliam Geffert, Klaus Reinhardt, and Abuzer Yakaryilmaz. New results on the minimum amount of useful space. *International Journal of Foundations of Computer Science*, 27(2):259–282, 2016. `doi:10.1142/S0129054116400098`.

5   Piotr Berman and Andrzej Lingas. On the complexity of regular languages in terms of finite automata. Technical Report 304, Polish Academy of Sciences, 1977.

6   Noam Chomsky. A note on phrase structure grammars. *Information and Control*, 2(4):393–395, 1959. `doi:10.1016/S0019-9958(59)80017-6`.

7   Noam Chomsky. On certain formal properties of grammars. *Information and Control*, 2(2):137–167, 1959. `doi:10.1016/S0019-9958(59)90362-6`.

8   Noam Chomsky. Context-Free Grammars and Pushdown Storage. Technical Report 65, Massachusetts Institute of Technology, Research Laboratory of Electronics, 1962. URL: `https://dspace.mit.edu/bitstream/handle/1721.1/53697/RLE_QPR_065_XVII.pdf`.

9   David Gajser. Verifying time complexity of Turing machines. *Theoretical Computer Science*, 600:86–97, 2015. `doi:10.1016/j.tcs.2015.07.028`.

10 Viliam Geffert, Bruno Guillon, and Giovanni Pighizzini. Two-way automata making choices only at the endmarkers. *Information and Computation*, 239:71–86, 2014. `doi:10.1016/j.ic.2014.08.009`.

11 Viliam Geffert, Carlo Mereghetti, and Beatrice Palano. More concise representation of regular languages by automata and regular expressions. *Information and Computation*, 208(4):385–394, 2010. `doi:10.1016/j.ic.2010.01.002`.

12 Viliam Geffert, Carlo Mereghetti, and Giovanni Pighizzini. Converting two-way nondeterministic unary automata into simpler automata. *Theoretical Computer Science*, 295(1–3):189–203, 2003. `doi:10.1016/S0304-3975(02)00403-6`.

13 Viliam Geffert and Giovanni Pighizzini. Two-way unary automata versus logarithmic space. *Information and Computation*, 209(7):1016–1025, 2011. `doi:10.1016/j.ic.2011.03.003`.

14 Jonathan Goldstine, Martin Kappes, Chandra M. R. Kintala, Hing Leung, Andreas Malcher, and Detlef Wotschke. Descriptional complexity of machines with limited resources. *Journal of Universal Computer Science*, 8(2):193–234, 2002. `doi:10.3217/jucs-008-02-0193`.

15 Bruno Guillon, Giovanni Pighizzini, and Luca Prigioniero. Non-self-embedding grammars, constant-height pushdown automata, and limited automata. *International Journal of Foundations of Computer Science*, 31(8):1133–1157, 2020. `doi:10.1142/S0129054120420071`.

16 Bruno Guillon, Giovanni Pighizzini, Luca Prigioniero, and Daniel Průša. Two-way automata and one-tape machines - read only versus linear time. In *Developments in Language Theory (DLT) 2018*, volume 11088 of *Lecture Notes in Computer Science*, pages 366–378. Springer, 2018. `doi:10.1007/978-3-319-98654-8_30`.

17 Bruno Guillon, Giovanni Pighizzini, Luca Prigioniero, and Daniel Průša. Converting nondeterministic two-way automata into small deterministic linear-time machines. *CoRR*, abs/2103.05485, 2021. `arXiv:2103.05485`.

18 Bruno Guillon, Giovanni Pighizzini, Luca Prigioniero, and Daniel Průša. Weight-reducing Turing machines. *CoRR*, abs/2103.05486, 2021. `arXiv:2103.05486`.

19 Bruno Guillon and Luca Prigioniero. Linear-time limited automata. *Theoretical Computer Science*, 798:95–108, 2019. `doi:10.1016/j.tcs.2019.03.037`.

20 Yo-Sub Han, Arto Salomaa, and Kai Salomaa. Ambiguity, nondeterminism and state complexity of finite automata. *Acta Cybernetica*, 23(1):141–157, 2017. `doi:10.14232/actacyb.23.1.2017.9`.

21 Juris Hartmanis. Context-free languages and Turing machine computations. In *Mathematical Aspects of Computer Science*, volume 19 of *Proceedings of Symposia in Applied Mathematics*, pages 42–51. American Mathematical Society, 1967.

22 Fred C. Hennie. One-tape, off-line Turing machine computations. *Information and Control*, 8(6):553–578, 1965. `doi:10.1016/S0019-9958(65)90399-2`.

23 Thomas N. Hibbard. A generalization of context-free determinism. *Information and Control*, 11(1/2):196–238, 1967. `doi:10.1016/S0019-9958(67)90513-X`.

24 Markus Holzer and Martin Kutrib. Descriptional complexity — an introductory survey. In *Scientific Applications of Language Methods*, pages 1–58. Imperial College Press, 2010. `doi:10.1142/9781848165458_0001`.

25 Markus Holzer and Martin Kutrib. One-time nondeterministic computations. *International Journal of Foundations of Computer Science*, 30(6-7):1069–1089, 2019. `doi:10.1142/S012905411940029X`.

26 John E. Hopcroft and Jeffrey D. Ullman. *Introduction to Automata Theory, Languages and Computation.* Addison-Wesley, 1979.

27 Juraj Hromkovič and Georg Schnitger. Nondeterminism versus determinism for two-way finite automata: Generalizations of Sipser's separation. In *International Colloquium on Automata, Languages, and Programming (ICALP) 2003*, volume 2719 of *Lecture Notes in Computer Science*, pages 439–451. Springer, 2003. `doi:10.1007/3-540-45061-0_36`.

28 Christos A. Kapoutsis. Nondeterminism is essential in small two-way finite automata with few reversals. *Information and Computation*, 222:208–227, 2013. `doi:10.1016/j.ic.2012.11.001`.

29 Christos A. Kapoutsis. Two-way automata versus logarithmic space. *Theory of Computing Systems*, 55(2):421–447, 2014. `doi:10.1007/s00224-013-9465-0`.

30 Christos A. Kapoutsis and Giovanni Pighizzini. Two-way automata characterizations of L/poly versus NL. *Theory of Computing Systems*, 56(4):662–685, 2015. `doi:10.1007/s00224-014-9560-x`.

31 Chris Keeler and Kai Salomaa. Nondeterminism growth and state complexity. In *Descriptional Complexity of Formal Systems (DCFS) 2019*, volume 11612 of *Lecture Notes in Computer Science*, pages 210–222. Springer, 2019. `doi:10.1007/978-3-030-23247-4_16`.

32 Alica Kelemenová. Complexity of normal form grammars. *Theoretical Computer Science*, 28:299–314, 1984. `doi:10.1016/0304-3975(83)90026-9`.

33 Sige-Yuki Kuroda. Classes of languages and linear-bounded automata. *Information and Control*, 7(2):207–223, 1964. `doi:10.1016/S0019-9958(64)90120-2`.

34 Martin Kutrib, Giovanni Pighizzini, and Matthias Wendlandt. Descriptional complexity of limited automata. *Information and Computation*, 259(2):259–276, 2018. `doi:10.1016/j.ic.2017.09.005`.

35 Hing Leung. Descriptional complexity of nfa of different ambiguity. *International Journal of Foundations of Computer Science*, 16(5):975–984, 2005. `doi:10.1142/S0129054105003418`.

36 Carlo Mereghetti and Giovanni Pighizzini. Two-way automata simulations and unary languages. *Journal of Automata, Languages and Combinatorics*, 5(3):287–300, 2000. `doi:10.25596/jalc-2000-287`.

37 Albert R. Meyer and Michael J. Fischer. Economy of description by automata, grammars, and formal systems. In *Switching Automata Theory (SwAT) 1971*, pages 188–191. IEEE Computer Society, 1971. `doi:10.1109/SWAT.1971.11`.

38 Alexander Okhotin. Non-erasing variants of the Chomsky-Schützenberger theorem. In *Developments in Language Theory (DLT) 2012*, volume 7410 of *Lecture Notes in Computer Science*, pages 121–129. Springer, 2012. `doi:10.1007/978-3-642-31653-1_12`.

39 Giovanni Pighizzini. Nondeterministic one-tape off-line Turing machines. *Journal of Automata, Languages and Combinatorics*, 14(1):107–124, 2009. `doi:10.25596/jalc-2009-107`.

40 Giovanni Pighizzini. Two-way finite automata: Old and recent results. *Fundamenta Informaticae*, 126(2-3):225–246, 2013. `doi:10.3233/FI-2013-879`.

41 Giovanni Pighizzini. Limited automata: Properties, complexity and variants. In *Descriptional Complexity of Formal Systems (DCFS) 2019*, volume 11612 of *Lecture Notes in Computer Science*, pages 57–73. Springer, 2019. `doi:10.1007/978-3-030-23247-4_4`.

42 Giovanni Pighizzini and Andrea Pisoni. Limited automata and regular languages. *International Journal of Foundations of Computer Science*, 25(7):897–916, 2014. `doi:10.1142/S0129054114400140`.

43 Giovanni Pighizzini and Andrea Pisoni. Limited automata and context-free languages. *Fundamenta Informaticae*, 136(1-2):157–176, 2015. `doi:10.3233/FI-2015-1148`.

44 Giovanni Pighizzini and Luca Prigioniero. Limited automata and unary languages. *Information and Computation*, 266:60–74, 2019. `doi:10.1016/j.ic.2019.01.002`.

45   Giovanni Pighizzini and Luca Prigioniero. Pushdown automata and constant height: De-cidability and bounds. In *Descriptional Complexity of Formal Systems (DCFS) 2019*, vol-ume 11612 of *Lecture Notes in Computer Science*, pages 260–271. Springer, 2019. `doi:` `10.1007/978-3-030-23247-4_20`.

46   Giovanni Pighizzini and Luca Prigioniero. Non-self-embedding grammars and descriptional complexity. *Fundamenta Informaticae*, 180(1-2):103–122, 2021. To appear. `doi:10.3233/` `FI-2021-2036`.

47   Giovanni Pighizzini, Jeffrey O. Shallit, and Ming-wei Wang. Unary context-free grammars and pushdown automata, descriptional complexity and auxiliary space lower bounds. *Journal of Computer and System Sciences*, 65(2):393–414, 2002. `doi:10.1006/jcss.2002.1855`.

48   Luca Prigioniero. *Regular Languages: To Finite Automata and Beyond - Succinct Descriptions and Optimal Simulations.* PhD thesis, Università degli Studi di Milano, Dipartimento di Informatica, 2020. `doi:10.13130/prigioniero-luca_phd2020-01-31`.

49   Luca Prigioniero. Regular languages: To finite automata and beyond succinct descriptions and optimal simulations. *Bulletin of EATCS*, 131, 2020. URL: `http://eatcs.org/beatcs/` `index.php/beatcs/article/view/619`.

50   Daniel Průša. Weight-reducing Hennie machines and their descriptional complexity. In *Language and Automata Theory and Applications (LATA) 2014*, volume 8370 of *Lecture Notes in Computer Science*, pages 553–564, 2014. `doi:10.1007/978-3-319-04921-2_45`.

51   Michael O. Rabin and Dana Scott. Finite automata and their decision problems. *IBM Journal of Research and Development*, 3(2):pp. 114–125, 1959. `doi:10.1147/rd.32.0114`.

52   William J. Sakoda and Michael Sipser. Nondeterminism and the size of two way finite automata. In *Symposium on Theory of Computing (SToC) 1978*, pages 275–286. ACM, 1978. `doi:10.1145/800133.804357`.

53   John C. Shepherdson. The reduction of two-way automata to one-way automata. *IBM Journal of Research and Development*, 3(2):198–200, 1959. `doi:10.1147/rd.32.0198`.

54   Michael Sipser. Halting space-bounded computations. *Theoretical Computer Science*, 10(3):335–338, 1980. `doi:10.1016/0304-3975(80)90053-5`.

55   Kohtaro Tadaki, Tomoyuki Yamakami, and Jack C. H. Lin. Theory of one-tape linear-time Turing machines. *Theoretical Computer Science*, 411(1):22–43, 2010. `doi:10.1016/j.tcs.` `2009.08.031`.

56   Klaus W. Wagner and Gerd Wechsung. *Computational Complexity.* D. Reidel Publishing Company, Dordrecht, 1986.