



The State Complexity of Lexicographically Smallest Words and Computing Successors

Lukas Fleischer^(✉) and Jeffrey Shallit

School of Computer Science, University of Waterloo, 200 University Avenue West,
Waterloo, ON N2L 3G1, Canada
{lukas.fleischer, shallit}@uwaterloo.ca

Abstract. Given a regular language L over an ordered alphabet Σ , the set of lexicographically smallest (resp., largest) words of each length is itself regular. Moreover, there exists an unambiguous finite-state transducer that, on a given word $w \in \Sigma^*$, outputs the length-lexicographically smallest word larger than w (henceforth called the L -successor of w). In both cases, naïve constructions result in an exponential blowup in the number of states. We prove that if L is recognized by a DFA with n states, then $2^{\Theta(\sqrt{n \log n})}$ states are sufficient for a DFA to recognize the subset $S(L)$ of L composed of its lexicographically smallest words. We give a matching lower bound that holds even if $S(L)$ is represented as an NFA. We then show that the same upper and lower bounds hold for an unambiguous finite-state transducer that computes L -successors.

1 Introduction

One of the most basic problems in formal language theory is the problem of enumerating the words of a language L . Since, in general, L is infinite, language enumeration is often formalized in one of the following two ways:

1. A function that maps an integer $n \in \mathbb{N}$ to the n -th word of L .
2. A function that takes a word and maps it to the next word in L .

Both descriptions require some linear ordering of the words in order for them to be well-defined. Usually, *radix order* (also known as length-lexicographical order) is used. Throughout this work, we focus on the second formalization.

While enumeration is non-computable in general, there are many interesting special cases. In this paper, we investigate the case of fixed regular languages, where successors can be computed in linear time [1, 2, 9]. Moreover, Frougny [7] showed that for every regular language L , the mapping of words to their successors in L can be realized by a finite-state transducer. Later, Angrand and Sakarovitch refined this result [3], showing that the successor function of any regular language is a finite union of functions computed by sequential

transducers that operate from right to left. However, to the best of our knowledge, no upper bound on the size of smallest transducer computing the successor function was known.

In this work, we consider transducers operating from left to right, and prove that the optimal upper bound for the size of transducers computing successors in L is in $2^{\Theta(\sqrt{n \log n})}$, where n is the size of the smallest DFA for L .

The construction used to prove the upper bound relies heavily on another closely related result. Many years before Frougny published her proof, it had already been shown that if L is a regular language, the set of all lexicographically smallest (resp., largest) words of each length is itself regular; see, e.g., [11, 12]. This fact is used both in [3] and in our construction. In [12], it was shown that if L is recognized by a DFA with n states, then the set of all lexicographically smallest words is recognized by a DFA with 2^{n^2} states. While it is easy to improve this upper bound to $n2^n$, the exact state complexity of this operation remained open. We prove that $2^{\Theta(\sqrt{n \log n})}$ states are sufficient and that this upper bound is optimal. We also prove that nondeterminism does not help with recognizing lexicographically smallest words, i.e., the corresponding lower bound still holds if the constructed automaton is allowed to be nondeterministic.

The key component to our results is a careful investigation of the structure of lexicographically smallest words. This is broken down into a series of technical lemmas in Sect. 3, which are interesting in their own right. Some of the other techniques are similar to those already found in [3], but need to be carried out more carefully to achieve the desired upper bound.

For some related results, see [5, 10].

2 Preliminaries

We assume familiarity with basic concepts of formal language theory and automata theory; see [8, 13] for a comprehensive introduction. Below, we introduce concepts and notation specific to this work.

Ordered Words and Languages. Let Σ be a finite ordered alphabet. Throughout the paper, we consider words ordered by *radix order*, which is defined by $u < v$ if either $|u| < |v|$ or there exist factorizations $u = xay$, $v = xbz$ with $|y| = |z|$ and $a, b \in \Sigma$ such that $a < b$. We write $u \leq v$ if $u = v$ or $u < v$. In this case, the word u is *smaller* than v and the word v is *larger* than u .

For a language $L \subseteq \Sigma^*$ and two words $u, v \in \Sigma^*$, we say that v is the *L-successor* of u if $v \in L$ and $w \notin L$ for all $w \in \Sigma^*$ with $u < w < v$. Similarly, u is the *L-predecessor* of v if $u \in L$ and $w \notin L$ for all $w \in \Sigma^*$ with $u < w < v$. A word is *L-minimal* if it has no *L-predecessor*. A word is *L-maximal* if it has no *L-successor*. Note that every nonempty language contains exactly one *L-minimal* word. It contains a (unique) *L-maximal* word if and only if L is finite. A word $u \in \Sigma^*$ is *L-length-preserving* if it is not *L-maximal* and the *L-successor* of u has length $|u|$. Words that are not *L-length-preserving* are called *L-length-increasing*.

Note that by definition, an L -maximal word is always L -length-increasing. For convenience, we sometimes use the terms *successor* (resp., *predecessor*) instead of Σ^* -*successor* (resp., Σ^* -*predecessor*).

For a given language $L \subseteq \Sigma^*$, the set of all smallest words of each length in L is denoted by $S(L)$. It is formally defined as follows:

$$S(L) = \{u \in L \mid \forall v \in L: v < u \implies |v| < |u|\}.$$

Similarly, we define $B(L)$ to be the set of all L -length-increasing words:

$$B(L) = \{u \in L \mid \forall v \in L: v > u \implies |v| > |u|\}.$$

A language $L \subseteq \Sigma^*$ is *thin* if it contains at most one word of each length, i.e., $|L \cap \Sigma^n| \in \{0, 1\}$ for all $n \geq 1$. It is easy to see that for every language $L \subseteq \Sigma^*$, the languages $S(L)$ and $B(L)$ are thin.

Finite Automata and Transducers. A *nondeterministic finite automaton* (NFA for short) is a 5-tuple $(Q, \Sigma, \cdot, q_0, F)$ where Q is a finite set of *states*, Σ is a finite alphabet, $q_0 \in Q$ is the *initial state*, $F \subseteq Q$ is the set of *accepting states* and $\cdot : Q \times \Sigma \rightarrow 2^Q$ is the *transition function*. We usually use the notation $q \cdot a$ instead of $\cdot(q, a)$, and we extend the transition function to $2^Q \times \Sigma^*$ by letting $X \cdot \varepsilon = X$ and $X \cdot wa = \bigcup_{q \in X \cdot w} q \cdot a$ for all $X \subseteq Q$, $w \in \Sigma^*$, and $a \in \Sigma$. For a state $q \in Q$ and a word $w \in \Sigma^*$, we also use the notation $q \cdot w$ instead of $\{q\} \cdot w$ for convenience. A word $w \in \Sigma^*$ is *accepted* by the NFA if $q_0 \cdot w \cap F \neq \emptyset$. We sometimes use the notation $p \xrightarrow{a} q$ to indicate that $q \in p \cdot a$. An NFA is *unambiguous* if for every input, there exists at most one accepting run. Unambiguous NFA are also called unambiguous finite state automata (UFA). A *deterministic finite automaton* (DFA for short) is an NFA $(Q, \Sigma, \cdot, q_0, F)$ with $|q \cdot a| = 1$ for all $q \in Q$ and $a \in \Sigma$. Since this implies $|q \cdot w| = 1$ for all $w \in \Sigma^*$, we sometimes identify the singleton $q \cdot w$ with the only element it contains.

A *finite-state transducer* is a nondeterministic finite automaton that additionally produces some output that depends on the current state, the current letter and the successor state. For each transition, we allow both the input and the output letter to be empty. Formally, it is a 6-tuple $(Q, \Sigma, \Gamma, \cdot, q_0, F)$ where Q is a finite set of states, Σ and Γ are finite alphabets, $q_0 \in Q$ is the initial state and $F \subseteq Q$ is the set of accepting states, and $\cdot : Q \times (\Sigma \cup \{\varepsilon\}) \rightarrow 2^{Q \times (\Gamma \cup \{\varepsilon\})}$ is the *transition function*. One can extend this transition function to the product $2^Q \times \Sigma^*$. To this end, we first define the ε -closure of a set $T \subseteq Q \times \Sigma^*$ as the smallest superset C of T with $\{(q \cdot \varepsilon, w) \mid (q, w) \in C\} \subseteq C$. We then define $X \cdot \varepsilon$ to be the ε -closure of $\{(q, \varepsilon) \mid q \in X\}$ and $X \cdot wa$ to be the ε -closure of $\{(q', ub) \mid (q, u) \in X \cdot w, (q', b) \in q \cdot a\}$ for all $X \subseteq Q$, $w \in \Sigma^*$ and $a \in \Sigma$. We sometimes use the notation $p \xrightarrow{a|b} q$ to indicate that $(q, b) \in p \cdot a$. A finite-state transducer is *unambiguous* if, for every input, there exists at most one accepting run.

3 The State Complexity of $S(L)$

It is known that if L is a regular language, then both $S(L)$ and $B(L)$ are also regular [11, 12]. In this section, we investigate the state complexity of the operations $L \mapsto S(L)$ and $L \mapsto B(L)$ for regular languages. Since the operations are symmetric, we focus on the former. To this end, we first prove some technical lemmas. The first lemma is a simple observation that helps us investigate the structure of words in $S(L)$.

Lemma 1. *Let $x, u, y, v, z \in \Sigma^*$ with $|u| = |v|$. Then $xuuyz < xuyvz$ or $xyvuz < xuyvz$ or $xuuyz = xuyvz = xyvuz$.*

Proof. Note that uy and yv are words of the same length. If $uy < yv$, then $xuuyz < xuyvz$. Similarly, $uy > yv$ immediately yields $xuyvz > xyvuz$. The last case is $uy = yv$, which implies $xuuyz = xuyvz = xyvuz$. \square

Using this observation, we can generalize a well-known factorization technique for regular languages to minimal words. For a DFA with state set Q , a state $q \in Q$ and a word $w = a_1 \cdots a_n \in \Sigma^*$, we define

$$\text{tr}(q, w) = (q, q \cdot a_1, \dots, q \cdot a_1 \cdots a_n)$$

to be the sequence of all states that are visited when starting in state q and following the transitions labeled by the letters from w .

Lemma 2. *Let \mathcal{A} be a DFA over Σ with n states and with initial state q_0 . Then for every word $w \in \Sigma^*$, there exists a factorization $w = u_1v_1^{i_1} \cdots u_kv_k^{i_k}$ with $u_1, v_1, \dots, u_k, v_k \in \Sigma^*$ and $i_1, \dots, i_k \geq 1$ such that, for all $j \in \{1, \dots, k\}$, the following hold:*

- (a) $q_0 \cdot u_1v_1^{i_1} \cdots u_{j-1}v_{j-1}^{i_{j-1}}u_j = q_0 \cdot u_1v_1^{i_1} \cdots u_{j-1}v_{j-1}^{i_{j-1}}u_jv_j$,
- (b) $|u_jv_j| \leq n$, and
- (c) v_j is not a prefix of $u_{j+1}v_{j+1}^{i_{j+1}} \cdots u_kv_k^{i_k}$.

Additionally, if $w \in S(L(\mathcal{A}))$, this factorization can be chosen such that

- (d) the lengths $|v_j|$ are pairwise disjoint (i.e., $|\{v_1, \dots, v_k\}| = k$) and
- (e) there exists at most one $j \in \{1, \dots, k\}$ with $i_j > n$.

Proof. To construct the desired factorization, initialize $j := 1$ and $q := q_0$ and follow these steps:

1. If $w = \varepsilon$, we are done. If $w \neq \varepsilon$ and the states in $\text{tr}(q, w)$ are pairwise distinct, let $u_j = w$ and $v_j = \varepsilon$ and we are done. Otherwise, factorize $w = xy$ with $|x|$ minimal such that $\text{tr}(q_0, x)$ contains exactly one state twice, i.e., $|x|$ distinct states in total.
2. Choose the unique factorization $x = uv$ such that $q \cdot u = q \cdot uv$ and $v \neq \varepsilon$.
3. Let $q := q \cdot x$ and $w := y$.

4. If $j > 1$ and $u = \varepsilon$ and $v = v_{j-1}$, increment i_{j-1} and go back to step 1. Otherwise, let $u_j := u$, $v_j := v$ and $j := j + 1$; then go back to step 1.

This factorization satisfies the first three properties by construction. It remains to show that if $w \in S(L(\mathcal{A}))$, then Properties (d) and (e) are satisfied as well.

Let us begin with Property (d). For the sake of contradiction, assume that there exist two indices a, b with $a < b$ and $|v_a| = |v_b|$. Note that by construction, v_a and v_b must be nonempty. Moreover, by Property (a), the words

$$\begin{aligned} w' &:= u_1 v_1^{i_1} \cdots u_a v_a^{i_a+1} \cdots u_b v_b^{i_b-1} \cdots u_k v_k^{i_k} \text{ and} \\ w'' &:= u_1 v_1^{i_1} \cdots u_a v_a^{i_a-1} \cdots u_b v_b^{i_b+1} \cdots u_k v_k^{i_k} \end{aligned}$$

both belong to $L(\mathcal{A})$. However, since $w \in S(L(\mathcal{A}))$, neither w' nor w'' can be strictly smaller than w . Using Lemma 1, we obtain that $w' = w$. This contradicts Property (c).

Property (e) can be proved by using the same argument: Assume that there exist indices a, b with $a < b$ and $i_a, i_b > n$. The words $v_a^{|v_b|}$ and $v_b^{|v_a|}$ have the same lengths. We define

$$\begin{aligned} w' &:= u_1 v_1^{i_1} \cdots u_a v_a^{i_a+|v_b|} \cdots u_b v_b^{i_b-|v_a|} \cdots u_k v_k^{i_k}, \\ w'' &:= u_1 v_1^{i_1} \cdots u_a v_a^{i_a-|v_b|} \cdots u_b v_b^{i_b+|v_a|} \cdots u_k v_k^{i_k}, \end{aligned}$$

and obtain $w = w'$, which is a contradiction as above. □

The existence of such a factorization almost immediately yields our next technical ingredient.

Lemma 3. *Let \mathcal{A} be a DFA with $n \geq 3$ states. Let q_0 be the initial state of \mathcal{A} and let $w \in S(L(\mathcal{A}))$. Then there exists a factorization $w = xy^i z$ with $i \in \mathbb{N}$, $|xz| \leq n^3$ and $|y| \leq n$ such that $q_0 \cdot xy = q_0 \cdot x$. In particular, $xy^* z \subseteq L(\mathcal{A})$.*

Proof. Let $w = u_1 v_1^{i_1} \cdots u_k v_k^{i_k}$ be a factorization that satisfies all properties in the statement of Lemma 2. Suppose first that all exponents i_j are at most n . Using Properties (b) and (d), we obtain $k \leq n + 1$ and the maximum length of w is achieved when all lengths $\ell \in \{0, \dots, n\}$ are present among the factors v_j and the corresponding u_j have lengths $n - |v_j|$. This yields

$$|w| \leq \sum_{\ell=0}^n (n - \ell + n\ell) = n(n + 1) + (n - 1) \sum_{\ell=1}^n \ell = n^2 + n + \frac{(n - 1)n(n + 1)}{2} \leq n^3$$

where the last inequality uses $n \geq 3$. Therefore, we may set $x := w$, $y := \varepsilon$ and $z := \varepsilon$.

If not all exponents are at most n , by Property (e), there exists a unique index j with $i_j > n$. In this case, let $x := u_1 v_1^{i_1} \cdots u_{j-1} v_{j-1}^{i_{j-1}} u_j$, $y := v_j$ and $z := u_{j+1} v_{j+1}^{i_{j+1}} \cdots u_k v_k^{i_k}$. The upper bound $|xz| \leq n^3$ still follows by the argument above, and $|y| \leq n$ is a direct consequence of Property (b). Moreover, $w \in L(\mathcal{A})$ and Property (a) together imply that $xy^* z \subseteq L(\mathcal{A})$. □

For the next lemma, we need one more definition. Let \mathcal{A} be a DFA with initial state q_0 . Two tuples (x, y, z) and (x', y', z') are *cycle-disjoint with respect to \mathcal{A}* if the sets of states in $\text{tr}(q_0 \cdot x, y)$ and $\text{tr}(q_0 \cdot x', y')$ are either equal or disjoint.

Lemma 4. *Let \mathcal{A} be a DFA with $n \geq 3$ states and initial state q_0 . Let (x, y, z) and (x', y', z') be tuples that are not cycle-disjoint with respect to \mathcal{A} such that*

$$q_0 \cdot x = q_0 \cdot xy, \quad q_0 \cdot x' = q_0 \cdot x'y', \quad |xz|, |x'z'| \leq n^3 \text{ and } |y|, |y'| \leq n.$$

*Then either $xy^*z \cap S(L(\mathcal{A}))$ or $x'(y')^*z' \cap S(L(\mathcal{A}))$ only contains words of length at most $n^3 + n^2$.*

Proof. Since the tuples are not cycle-disjoint with respect to \mathcal{A} , we can factorize $y = uv$ and $y' = u'v'$ such that $q_0 \cdot xu = q_0 \cdot x'u'$.

Note that since $q_0 \cdot xuv = q_0 \cdot x$, the sets of states in $\text{tr}(q_0 \cdot x, uv)$ and $\text{tr}(q_0 \cdot xu, (vu)^i)$ coincide for all $i \geq 1$. By the same argument, the sets of states in $\text{tr}(q_0 \cdot x', u'v')$ and $\text{tr}(q_0 \cdot x'u', (v'u')^i)$ coincide for all $i \geq 1$.

If the powers $(vu)^{|y'|}$ and $(v'u')^{|y|}$ were equal, then $\text{tr}(q_0 \cdot xu, (vu)^{|y'|})$ and $\text{tr}(q_0 \cdot x'u', (v'u')^{|y|})$ coincide. By the previous observation, this would imply that the tuples (x, y, z) and (x', y', z') are cycle-disjoint, a contradiction. We conclude $(vu)^{|y'|} \neq (v'u')^{|y|}$.

By symmetry, we may assume that $(vu)^{|y'|} < (v'u')^{|y|}$. But then, for every word of the form $x'(y')^i z' \in L(\mathcal{A})$ with $i > |y|$, there exists a strictly smaller word $x'u'(vu)^{|y'|}(v'u')^{i-|y|-1}v'z'$ in $L(\mathcal{A})$. To see that this word indeed belongs to $L(\mathcal{A})$, note that $q_0 \cdot x'u'vu = q_0 \cdot xuvu = q_0 \cdot xu = q_0 \cdot x'u'$. This means that all words in $x'(y')^*z' \cap S(L)$ are of the form $x'(y')^i z'$ with $i \leq |y|$. \square

The previous lemmas now allow us to replace any language L by another language that has a simple structure and approximates L with respect to $S(L)$.

Lemma 5. *Let \mathcal{A} be a DFA over Σ with $n \geq 3$ states. Then there exist an integer $k \leq n^4 + n^3$ and tuples $(x_1, y_1, z_1), \dots, (x_k, y_k, z_k) \in (\Sigma^*)^3$ such that the following properties hold:*

- (i) $S(L(\mathcal{A})) \subseteq \bigcup_{i=1}^k x_i y_i^* z_i \subseteq L(\mathcal{A})$,
- (ii) $|x_i z_i| \leq n^3 + n^2$ for all $i \in \{1, \dots, k\}$, and
- (iii) $\sum_{\ell \in Y} \ell \leq n$ where $Y = \{|y_1|, \dots, |y_k|\}$.

Proof. If we ignore the required upper bound $k \leq n^4 + n^3$ and Property (iii) for now, the statement follows immediately from Lemma 3 and the fact that there are only finitely many different tuples (x, y, z) with $|xz| \leq n^3$ and $|y| \leq n$. We start with such a finite set of tuples $(x_1, y_1, z_1), \dots, (x_k, y_k, z_k)$ and show that we can repeatedly eliminate tuples until at most $n^4 + n^3$ cycle-disjoint tuples remain. The desired upper bound $\sum_{\ell \in Y} \ell \leq n$ then follows automatically.

In each step of this elimination process, we handle one of the following cases:

- If there are two distinct tuples (x_i, y_i, z_i) and (x_j, y_j, z_j) with $|x_i z_i| = |x_j z_j|$ and $y_i = y_j$, there are two possible scenarios. If $x_i z_i < x_j z_j$, then for every word in $x_j y_j^* z_j$ there exists a smaller word in $x_i y_i^* z_i$ and we can remove (x_j, y_j, z_j) from the set of tuples. By the same argument, we can remove the tuple (x_i, y_i, z_i) if $y_i = y_j$ and $x_i z_i > x_j z_j$.
- Now consider the case that there are two distinct tuples (x_i, y_i, z_i) and (x_j, y_j, z_j) with $|x_i z_i| = |x_j z_j|$ and $|y_i| = |y_j|$ but $y_i \neq y_j$. We first check whether $x_i z_i < x_j z_j$. If true, we add the tuple (x_i, ε, z_i) , otherwise we add (x_j, ε, z_j) . If $x_i y_i < x_j y_j$, we know that each word in $x_j y_j^+ z_j$ has a smaller word in $x_i y_i^+ z_j$, and we remove the tuple (x_j, y_j, z_j) . Otherwise, we can remove (x_i, y_i, z_i) by the same argument.
- The last case is that there exist two tuples (x_i, y_i, z_i) and (x_j, y_j, z_j) that are not cycle-disjoint. By Lemma 4, we can remove at least one of these tuples and replace it by multiple tuples of the form (x, ε, z) . Note that the newly introduced tuples might be of the form (x, ε, z) with $|xz| > n^3$ but Lemma 4 asserts that they still satisfy $|xz| \leq n^3 + n^2$.

Note that we introduce new tuples of the form (x, ε, z) during this elimination process. These new tuples are readily eliminated using the first rule.

After iterating this elimination process, the remaining tuples are pairwise cycle-disjoint and the pairs $(|x_i z_i|, |y_i|)$ assigned to these tuples (x_i, y_i, z_i) are pairwise disjoint. Properties (ii) and (iii) yield the desired upper bound on k . \square

Remark 1. While $S(L)$ can be approximated by a language of the simple form given in Lemma 5, the language $S(L)$ itself does not necessarily have such a simple description. An example of a regular language L where $S(L)$ does not have such a simple form is given in the proof of Theorem 2.

The last step is to investigate languages L of the simple structure described in the previous lemma and show how to construct a small DFA for $S(L)$.

Lemma 6. *Let $n \in \mathbb{N}$. Let $L = \bigcup_{i=1}^k x_i y_i^* z_i$ with $k \leq n^4 + n^3$ and $|x_i z_i| \leq n^3 + n^2$ for all $i \in \{1, \dots, k\}$ and $\sum_{\ell \in Y} \ell \leq n$ where $Y = \{|y_1|, \dots, |y_k|\}$. Then $S(L)$ is recognized by a DFA with $2^{\mathcal{O}(\sqrt{n \log n})}$ states.*

Proof. We describe how to construct a DFA of the desired size that recognizes the language $S(L)$. This DFA is the product automaton of multiple components.

In one component (henceforth called the *counter component*), we keep track of the length of the processed input as long as at most $n^3 + n^2$ letters have been consumed. If more than $n^3 + n^2$ letters have been consumed, we only keep track of the length of the processed input modulo all numbers $|y_i|$ for $i \in \{1, \dots, k\}$.

For each $i \in \{1, \dots, k\}$, there is an additional component (henceforth called the *i -th activity component*). In this component, we keep track of whether the currently processed prefix u of the input is a prefix of a word in $x_i y_i^*$, whether u is a prefix of a word in $x_i y_i^* z_i$ and whether $u \in x_i y_i^* z_i$. Note that if some prefix of the input is not a prefix of a word in $x_i y_i^* z_i$, no longer prefix of the input can

be a prefix of a word in $x_i y_i^* z_i$. The information stored in the counter component suffices to compute the possible letters of $x_i y_i^* z_i$ allowed to be read in each step to maintain the prefix invariants.

It remains to describe how to determine whether a state is final. To this end, we use the following procedure. First, we determine which sets of the form $x_i y_i^* z_i$ the input word leading to the considered state belongs to. These languages are called the *active languages* of the state. They can be obtained from the activity components of the state. If there are no active languages, the state is immediately marked as not final. If the length of the input word w leading to the considered state is $n^3 + n^2$ or less, we can obtain $|w|$ from the counter component and reconstruct w from the set of active languages. If the length of the input is larger than $n^3 + n^2$, we cannot fully recover the input from the information stored in the state. However, we can determine the shortest word w with $|w| > n^3 + n^2$ such that $|w|$ is consistent with the length information stored in the counter component and w itself is consistent with the set of active languages. In either case, we then compute the set A of all words of length $|w|$ that belong to any (possibly not active) language $x_i y_i^* z_i$ with $1 \leq i \leq k$. If w is the smallest word in A , the state is final, otherwise it is not final.

The desired upper bound on the number of states follows from known estimates on the least common multiple of a set of natural numbers with a given sum; see e.g., [6]. □

We can now combine the previous lemmas to obtain an upper bound on the state complexity of $S(L)$.

Theorem 1. *Let L be a regular language that is recognized by a DFA with n states. Then $S(L)$ is recognized by a DFA with $2^{\mathcal{O}(\sqrt{n \log n})}$ states.*

Proof. By Lemma 5, we know that there exists a language L' of the form described in the statement of Lemma 6 with $S(L) \subseteq L' \subseteq L$. Since $L' \subseteq L$ implies $S(L') \subseteq S(L)$ and since $S(S(L)) = S(L)$, this also means that $S(L') = S(L)$. Lemma 6 now shows that there exists a DFA of the desired size. □

To show that the result is optimal, we provide a matching lower bound.

Theorem 2. *There exists a family of DFA $(\mathcal{A}_n)_{n \in \mathbb{N}}$ over a binary alphabet such that \mathcal{A}_n has n states and every NFA for $S(L(\mathcal{A}_n))$ has $2^{\Omega(\sqrt{n \log n})}$ states.*

Proof. For $i \in \{1, \dots, k\}$, let p_i be the i -th prime number and let $p = p_1 \cdots p_k$. We define a language

$$L = 1^* \cup \bigcup_{1 \leq i \leq k} 1^i 0^{k-i+1} \{1, 1^2, \dots, 1^{p_i-1}\} (1^{p_i})^*.$$

It is easy to see that L is recognized by a DFA with $k^2 + p_1 + \dots + p_k$ states. We show that $S(L)$ is not recognized by any NFA with less than p states. From known estimates on the prime numbers (e.g., [4, Sec. 2.7]), this suffices to prove our claim.

Let \mathcal{A} be a NFA for $S(L)$ and assume, for the sake of contradiction, that \mathcal{A} has less than p states. Note that since for each $i \in \{1, \dots, k\}$, the integer p is a multiple of p_i , the language L does not contain any word of the form $1^i 0^{k-i+1} 1^p$. Therefore, the word 1^{k+1+p} belongs to $S(L)$ and by assumption, an accepting path for this word in \mathcal{A} must contain a loop of some length $\ell \in \{1, \dots, p-1\}$. But then $1^{k+1+p+\ell}$ is accepted by \mathcal{A} , too. However, since $1 \leq \ell < p$, there exists some $i \in \{1, \dots, k\}$ such that p_i does not divide ℓ . This means that p_i also does not divide $p + \ell$. Thus, $1^i 0^{k-i+1} 1^{p+\ell} \in L$, contradicting the fact that $1^{k+1+p+\ell}$ belongs to $S(L)$. \square

Combining the previous two theorems, we obtain the following corollary.

Corollary 1. *Let L be a language that is recognized by a DFA with n states. Then, in general, $2^{\Theta(\sqrt{n \log n})}$ states are necessary and sufficient for a DFA or NFA to recognize $S(L)$.*

By reversing the alphabet ordering, we immediately obtain similar results for largest words.

Corollary 2. *Let L be a language that is recognized by a DFA with n states. Then, in general, $2^{\Theta(\sqrt{n \log n})}$ states are necessary and sufficient for a DFA or NFA to recognize $B(L)$.*

4 The State Complexity of Computing Successors

One approach to efficient enumeration of a regular language L is constructing a transducer that reads a word and outputs its L -successor [3, 7]. We consider transducers that operate from left to right. Since the output letter in each step might depend on letters that have not yet been read, this transducer needs to be nondeterministic. However, the construction can be made *unambiguous*, meaning that for any given input, at most one computation path is accepting and yields the desired output word. In this paper, we prove that, in general, $2^{\Theta(\sqrt{n \log n})}$ states are necessary and sufficient for a transducer that performs this computation.

Our proof is split into two parts. First, we construct a transducer that only maps L -length-preserving words to their corresponding L -successors. All other words are rejected. This construction heavily relies on results from the previous section. Then we extend this transducer to L -length-increasing words by using a technique called *padding*. For the first part, we also need the following result.

Theorem 3. *Let $L \subseteq \Sigma^*$ be a thin language that is recognized by a DFA with n states. Then the languages*

$$L_{\leq} = \{v \in \Sigma^* \mid \exists u \in L: |u| = |v| \text{ and } v \leq u\} \text{ and}$$

$$L_{\geq} = \{v \in \Sigma^* \mid \exists u \in L: |u| = |v| \text{ and } v \geq u\}$$

are recognized by UFA with $2n$ states.

Proof. Let $\mathcal{A} = (Q, \Sigma, \cdot, q_0, F)$ be a DFA for L and let $n = |Q|$. We construct a UFA with $2n$ states for L_{\leq} . The statement for L_{\geq} follows by symmetry.

The state set of the UFA is $Q \times \{0, 1\}$, the initial state is $(q_0, 0)$ and the set of final states is $F \times \{0, 1\}$. The transitions are

$$\begin{aligned} (q, 0) &\xrightarrow{a} (q \cdot a, 0) && \text{for all } q \in Q \text{ and } a \in \Sigma, \\ (q, 0) &\xrightarrow{a} (q \cdot b, 1) && \text{for all } q \in Q \text{ and } a, b \in \Sigma \text{ with } a < b, \\ (q, 1) &\xrightarrow{a} (q \cdot b, 1) && \text{for all } q \in Q \text{ and } a, b \in \Sigma. \end{aligned}$$

It is easy to verify that this automaton indeed recognizes L_{\leq} . To see that this automaton is unambiguous, consider an accepting run of a word w of length ℓ . Note that the sequence of first components of the states in this run yield an accepting path of length ℓ in \mathcal{A} . Since $L(\mathcal{A})$ is thin, this path is unique. Therefore, the sequence of first components is uniquely defined. The second components are then uniquely defined, too: they are 0 up to the first position where w differs from the unique word of length ℓ in L , and 1 afterwards. \square

For a language $L \subseteq \Sigma^*$, we denote by $B_{\geq}(L)$ the language of all words from Σ^* such that there exists no strictly larger word of the same length in L . Combining Theorem 1 and Theorem 3, the following corollary is immediate.

Corollary 3. *Let L be a language that is recognized by a DFA with n states. Then there exists a UFA with $2^{\mathcal{O}(\sqrt{n \log n})}$ states that recognizes the language $B_{\geq}(L)$.*

For a language $L \subseteq \Sigma^*$, we define

$$X(L) = \{u \in \Sigma^* \mid \forall v \in L: |u| \neq |v|\}.$$

If L is regular, it is easy to construct an NFA for the complement of $X(L)$, henceforth denoted as $\overline{X(L)}$. To this end, we take a DFA for L and replace the label of each transition with all letters from Σ . This NFA can also be viewed as an NFA over the unary alphabet $\{\Sigma\}$; here, Σ is interpreted as a letter, not a set. It can be converted to a DFA for $\overline{X(L)}$ by using Chrobak's efficient determinization procedure for unary NFA [6]. The resulting DFA can then be complemented to obtain a DFA for $X(L)$:

Corollary 4. *Let L be a language that is recognized by a DFA with n states. Then there exists a DFA with $2^{\mathcal{O}(\sqrt{n \log n})}$ states that recognizes the language $X(L)$.*

We now use the previous results to prove an upper bound on the size of a transducer performing a variant of the L -successor computation that only works for L -length-preserving words.

Theorem 4. *Let L be a language that is recognized by a DFA with n states. Then there exists an unambiguous finite-state transducer with $2^{\mathcal{O}(\sqrt{n \log n})}$ states that rejects all L -length-increasing words and maps every L -length-preserving word to its L -successor.*

Proof. Let $\mathcal{A} = (Q, \Sigma, \cdot, q_0, F)$ be a DFA for L and let $n = |Q|$. For every $q \in Q$, we denote by \mathcal{A}_q the DFA that is obtained by making q the new initial state of \mathcal{A} . We use \mathcal{A}_q^S to denote DFA with $2^{\mathcal{O}(\sqrt{n \log n})}$ states that recognizes the language $S(L(\mathcal{A}_q))$. These DFA exist by Theorem 1. Moreover, by Corollary 3, there exist UFA with $2^{\mathcal{O}(\sqrt{n \log n})}$ states that recognize the languages $B_{\geq}(L(\mathcal{A}_q))$. We denote these UFA by \mathcal{A}_q^B . Similarly, we use \mathcal{A}_q^X to denote DFA with $2^{\mathcal{O}(\sqrt{n \log n})}$ states that recognize $X(L(\mathcal{A}_q))$. These DFA exist by Corollary 4.

In the finite-state transducer, we first simulate \mathcal{A} on a prefix u of the input, copying the input letters in each step, i.e., producing the output u . At some position, after having read a prefix u leading up to the state $q := q_0 \cdot u$, we nondeterministically decide to output a letter b that is strictly larger than the current input letter a . From then on, we guess an output letter in each step and start simulating multiple automata in different components. In one component, we simulate \mathcal{A}_{q-a}^B on the remaining input. In another component, we simulate \mathcal{A}_{q-b}^S on the guessed output. In additional components, for each $c \in \Sigma$ with $a < c < b$, we simulate \mathcal{A}_{q-c}^X on the input. The automata in all components must accept in order for the transducer to accept the input.

The automaton \mathcal{A}_{q-a}^B verifies that there is no word in L that starts with the prefix ua , has the same length as the input word and is strictly larger than the input word. The automaton \mathcal{A}_{q-b}^S verifies that there is no word in L that starts with the prefix ub , has the same length as the input word and is strictly smaller than the output word. It also certifies that the output word belongs to L . For each letter c , the automaton \mathcal{A}_{q-c}^X verifies that there is no word in L that starts with the prefix uc and has the same length as the input word.

Together, the components ensure that the guessed output is the unique successor of the input word, given that it is L -length-preserving. It is also clear that L -length-increasing words are rejected, since the \mathcal{A}_{q-b}^S -component does not accept for any sequence of nondeterministic choices. \square

The construction given in the previous proof can be extended to also compute L -successors of L -length-increasing words. However, this requires some quite technical adjustments to the transducer. Instead, we use a technique called *padding*. A very similar approach appears in [3, Prop. 5.1].

We call the smallest letter of an ordered alphabet Σ the *padding symbol* of Σ . A language $L \subseteq \Sigma^*$ is \diamond -*padded* if \diamond is the padding symbol of Σ and $L = \diamond^* K$ for some $K \subseteq (\Sigma \setminus \{\diamond\})^*$. The key property of padded languages is that all words prefixed by a sufficiently long block of padding symbols are L -length-preserving.

Lemma 7. *Let \mathcal{A} be a DFA over Σ with n states such that $L(\mathcal{A})$ is a \diamond -padded language. Let $\Gamma = \Sigma \setminus \{\diamond\}$ and let $K = L(\mathcal{A}) \cap \Gamma^*$. Let $u \in \Gamma^*$ be a word that is not K -maximal. Then the $L(\mathcal{A})$ -successor of $\diamond^n u$ has length $|\diamond^n u|$.*

Proof. Let v be the K -successor of u . By a standard pumping argument, we have $|u| \leq |v| \leq |u| + n$. This means that $\diamond^{n+|u|-|v|} v$ is well-defined and belongs to $L(\mathcal{A})$. Note that this word is strictly greater than $\diamond^n u$ and has length $|\diamond^n u|$. Thus, the $L(\mathcal{A})$ -successor of $\diamond^n u$ has length $|\diamond^n u|$, too. \square

We now state the main result of this section.

Theorem 5. *Let \mathcal{A} be a deterministic finite automaton over Σ with n states. Then there exists an unambiguous finite-state transducer with $2^{\mathcal{O}(\sqrt{n \log n})}$ states that maps every word to its $L(\mathcal{A})$ -successor.*

Proof. We extend the alphabet by adding a new padding symbol \diamond and convert \mathcal{A} to a DFA for \diamond^*L by adding a new initial state. The language L' accepted by this new DFA is \diamond -padded. By Theorem 4 and Lemma 7, there exists an unambiguous transducer of the desired size that maps every word from $\diamond^{n+1}\Sigma^*$ to its successor in L' . It is easy to modify this transducer such that all words that do not belong to $\diamond^{n+1}\Sigma^*$ are rejected. We then replace every transition that reads a \diamond by a corresponding transition that reads the empty word instead. Similarly, we replace every transition that outputs a \diamond by a transition that outputs the empty word instead. Clearly, this yields the desired construction for the original language L . A careful analysis of the construction shows that the transducer remains unambiguous after each step. \square

We now show that this construction is optimal up to constants in the exponent. The idea is similar to the construction used in Theorem 2.

Theorem 6. *There exists a family of deterministic finite automata $(\mathcal{A}_n)_{n \in \mathbb{N}}$ such that \mathcal{A}_n has n states whereas the smallest unambiguous transducer that maps every word to its $L(\mathcal{A}_n)$ -successor has $2^{\Omega(\sqrt{n \log n})}$ states.*

Proof. Let $k \in \mathbb{N}$. Let p_1, \dots, p_k be the k smallest prime numbers such that $p_1 < \dots < p_k$ and let $p = p_1 \cdots p_k$. We construct a deterministic finite automaton \mathcal{A} with $2 + p_1 + \dots + p_k$ states such that the smallest transducer computing the desired mapping has at least p states. From known estimates on the prime numbers (e.g., [4, Sec. 2.7]), this suffices to prove our claim.

The automaton is defined over the alphabet $\Sigma = \{1, \dots, k\} \cup \{\#\}$. It consists of an initial state q_0 , an error state q_{err} , and states (i, j) for $i \in \{1, \dots, k\}$ and $j \in \{0, \dots, p_i - 1\}$ with transitions defined as follows:

$$q_0 \cdot a = \begin{cases} (a, 0), & \text{for } a \in \{1, \dots, k\}; \\ q_{err}, & \text{if } a = \#; \end{cases}$$

$$(i, j) \cdot a = \begin{cases} (i, j + 1 \bmod p_i), & \text{if } a = \#; \\ q_{err}, & \text{for } a \in \{1, \dots, k\}. \end{cases}$$

The set of accepting states is $\{(i, 0) \mid 1 \leq i \leq k\}$. The language $L(\mathcal{A})$ is the set of all words of the form $i\#^j$ with $1 \leq i \leq k$ such that j is a multiple of p_i .

Assume, to get a contradiction, that there exists an unambiguous transducer with less than p states that maps w to the smallest word in $L(\mathcal{A})$ strictly greater than w . Consider an accepting run of this transducer on some input of the form $2\#^{\ell p}$ with $\ell \in \mathbb{N}$ large enough such that the run contains a cycle. Clearly, since $\ell p + 1$ and p are coprime, the output of the transducer has to be $2\#^{\ell p + 2}$. We fix one cycle in this run.

If the number of $\#$ read in this cycle does not equal the number of $\#$ output in this cycle, by using a pumping argument, we can construct a word of the form $2\#^j$ that is mapped to a word of the form $i\#^{j'}$ with $|j' - j| > 2$. This contradicts the fact that $2\#^{2\mathbb{N}}$ is a subset of $L(\mathcal{A})$. Therefore, we may assume that both the number of letters read and output on the cycle is $r \in \{1, \dots, p-1\}$.

Again, by a pumping argument, this implies that $2\#\ell_{p+jr}$ is mapped to $2\#\ell_{p+jr+2}$ for every $j \in \mathbb{N}$. Since $r < p$, at least one of the prime numbers p_i is coprime to r . Therefore, we can choose j such that $jr + 1 \equiv 0 \pmod{p_i}$. However, this means that $p_i\#\ell_{p+jr+1}$ belongs to $L(\mathcal{A})$, contradicting the fact that the transducer maps $2\#\ell_{p+jr}$ to $2\#\ell_{p+jr+2}$. \square

Combining the two previous theorems, we obtain the following corollary.

Corollary 5. *Let L be a language that is recognized by a DFA with n states. Then, in general, $2^{\Theta(\sqrt{n \log n})}$ states are necessary and sufficient for an unambiguous finite-state transducer that maps words to their L -successors.*

References

1. Ackerman, M., Mäkinen, E.: Three new algorithms for regular language enumeration. In: Ngo, H.Q. (ed.) COCOON 2009. LNCS, vol. 5609, pp. 178–191. Springer, Heidelberg (2009). https://doi.org/10.1007/978-3-642-02882-3_19
2. Ackerman, M., Shallit, J.: Efficient enumeration of words in regular languages. *Theoret. Comput. Sci.* **410**(37), 3461–3470 (2009)
3. Angrand, P.-Y., Sakarovitch, J.: Radix enumeration of rational languages. *RAIRO - Theoret. Inform. Appl.* **44**(1), 19–36 (2010)
4. Bach, E., Shallit, J.: *Algorithmic Number Theory*. MIT Press, Cambridge (1996)
5. Berthé, V., Frougny, C., Rigo, M., Sakarovitch, J.: On the cost and complexity of the successor function. In: Arnoux, P., Bédaride, N., Cassaigne, J. (eds.) *Proceedings of WORDS 2007*, Technical Report, Institut de mathématiques de Luminy, pp. 43–56 (2007)
6. Chrobak, M.: Finite automata and unary languages. *Theoret. Comput. Sci.* **47**, 149–158 (1986). Erratum, 302:497–498, 2003
7. Frougny, C.: On the sequentiality of the successor function. *Inf. Comput.* **139**(1), 17–38 (1997)
8. Hopcroft, J.E., Motwani, R., Ullman, J.D.: *Introduction to Automata Theory, Languages, and Computation*, 3rd edn. Addison-Wesley Longman Publishing Co., Inc., New York (2006)
9. Mäkinen, E.: On lexicographic enumeration of regular and context-free languages. *Acta Cybern.* **13**(1), 55–61 (1997)
10. Okhotin, A.S.: On the complexity of the string generation problem. *Discret. Math. Appl.* **13**, 467–482 (2003)
11. Sakarovitch, J.: Deux remarques sur un théorème de S. Eilenberg. *RAIRO - Theoret. Inform. Appl.* **17**(1), 23–48 (1983)
12. Shallit, J.: Numeration systems, linear recurrences, and regular sets. *Inf. Comput.* **113**(2), 331–347 (1994)
13. Shallit, J.: *A Second Course in Formal Languages and Automata Theory*. Cambridge University Press, Cambridge (2008)