

Algorithmics on SLP-Compressed Strings: A Survey

Markus Lohrey

Universität Leipzig, Institut für Informatik, Germany
lohrey@informatik.uni-leipzig.de

October 2, 2012

Abstract

Results on algorithmic problems on strings that are given in a compressed form via straight-line programs are surveyed. A straight-line program is a context-free grammar that generates exactly one string. In this way, exponential compression rates can be achieved. Among others, we study pattern matching for compressed strings, membership problems for compressed strings in various kinds of formal languages, and the problem of querying compressed strings. Applications in combinatorial group theory and computational topology and to the solution of word equations are discussed as well. Finally, extensions to compressed trees and pictures are considered.

1 Introduction

The topic of this paper are algorithms on compressed strings. The goal of such algorithms is to check properties of compressed strings and thereby beat a straight forward “decompress-and-check” strategy. There are three main applications for algorithms of this kind.

- In many areas, large string data have to be not only stored in compressed form, but the initial data has to be processed and analyzed as well. Here, it makes sense to design algorithms that directly operate on the compressed string representation in order to save the time and space for (de)compression. Such a scenario can be found for instance in large genom databases or XML processing.
- Large and often highly compressible strings may appear as intermediate data structures in algorithms. Here, one may try to store a compressed representation of these intermediate data structures and to process this representation. This may lead to more efficient algorithms. Examples for this strategy can be found for instance in combinatorial group theory [57, 58, 95, 97, 125], computational topology [37, 122, 124], interprocedural analysis [52], and bisimulation checking [61, 79].
- In some situations it makes sense to compute in a first phase a compressed representation of an input string, which makes regularities in the string explicit. These regularities may be exploited in a second phase for speeding up an algorithm. This principle is known as *acceleration by compression*. It was recently applied in order to speed up the Viterbi algorithm for analyzing hidden Markov models [83] as well as speeding up edit distance computation [31, 59].

When we talk about algorithms on compressed strings, we have to make precise the compressed representation we want to use. Such a representation should have two properties: (i) it should cover many compression schemes from practice and (ii) it should be mathematically easy to handle. Straight-line programs (SLPs) have both of these properties. An SLP is a context-free grammar that generates exactly one word. In an SLP, repeated subpatterns in a string have to be represented only once by introducing a non-terminal for the pattern. It is not difficult to see that with an SLP

consisting of n productions, a string of length 2^n can be generated by repeated doubling. In this sense, an SLP can indeed be seen as a compressed representation of the string it generates.

Several dictionary-based compressed representations, like for instance run-length encoding and the family of Lempel-Ziv (LZ) encodings [144], can be converted efficiently into straight-line programs of similar size. Moreover, a certain variant of LZ77-encoding turns out to be equivalent to SLPs in the following sense: From the LZ77-encoding of a string one can compute in polynomial time an SLP for the string, and, vice versa, from an SLP for a string one can compute in polynomial time the LZ77-encoding of the string [115]. This implies that complexity results can be transferred from SLP-encoded input strings to LZ77-encoded input strings. We will discuss the relationship between LZ77 and SLPs in more detail in Section 3.

The problem of computing a small SLP for a given input string is known as *grammar-based compression*; it will be briefly discussed in Section 4. Although computing a size-minimal SLP for a given input string is not possible in polynomial time unless $P = NP$ [23], there are several algorithms that compute for a given input string of length n an SLP that is only by a factor $\log(n)$ larger than a minimal grammar [23, 119, 121].

In Sections 5–10 we consider algorithmic problems for SLP-compressed strings. The following types of algorithmic problems on compressed strings will be studied:

- Comparing compressed strings (Section 5): Are two compressed strings equal or similar in a certain sense?
- Pattern matching for compressed strings (Section 6 and 8): Does a compressed string occur (in a certain sense) as a pattern in another compressed string?
- Membership problems for compressed strings in regular and context-free languages (Section 11): Does a compressed string belong to a certain regular or context-free language?
- Querying problems for compressed strings (Section 10): What is the symbol at a given position in a compressed string?

In Section 7 we discuss the relationship between SLPs and leaf languages (and important tool in complexity theory), which is a useful technique for proving lower bounds on the complexity of algorithmic problems for SLP-compressed strings.

In Section 11 we present applications in combinatorial group theory. In particular, we survey results on the *compressed word problem* for groups. The compressed word problem for the finitely generated group G asks whether an SLP-encoded input string over the generators of the group represents the group identity. It turned out that the compressed word problem for a group G is the right tool in order to attack the (non-compressed) word problem for (finitely generated subgroups of) the automorphism group of G . For instance, along this line it was shown that the word problem for the automorphism group of a free group can be solved in polynomial time [125].

Similar to group theory, SLP-techniques can also be used in order to show that several problems in computational topology can be solved in polynomial time; this is the topic of Section 13. These problems mainly concern curves in 2-dimensional surfaces that are represented by so called normal coordinates. With the help of SLPs, it was shown for instance that Dehn twists and geometric intersection numbers of curves that are given in normal coordinates can be computed in polynomial time [124].

Finally, in Section 15 we briefly discuss algorithmic problems for extensions of SLPs to trees (Section 15.1) and 2-dimensional pictures (Section 15.2).

Previous surveys on algorithms for compressed strings can be found in [42, 116, 120].

2 Preliminaries

2.1 General notations

Let Γ be a finite alphabet. The *empty word* is denoted by ε . Let $s = a_1 a_2 \cdots a_n \in \Gamma^*$ be a word over Γ , where $a_i \in \Gamma$ for $1 \leq i \leq n$. The set of symbols occurring in s is $\text{alph}(s) = \{a_1, \dots, a_n\}$.

The *length* of s is $|s| = n$. For $a \in \Gamma$ and $1 \leq i \leq j \leq n$ let $|s|_a = |\{k \mid a_k = a\}|$ be the number of a 's in s . We also write $s[i] = a_i$, and $s[i, j] = a_i a_{i+1} \cdots a_j$. If $i > j$ we set $s[i, j] = \varepsilon$. Any word $s[i, j]$ is called a *factor* of s .

We assume some background in complexity theory [110]. In particular, the reader should be familiar with the complexity classes L (deterministic logarithmic space) P (deterministic polynomial time), NP (nondeterministic polynomial time), coNP (languages, whose complement belongs to NP), PSPACE (polynomial space), EXPTIME (deterministic exponential time), and NEXPTIME (nondeterministic exponential time). As usual, when talking about space bounded classes (e.g. logarithmic space), we only count the required space on the work tape (but not the space occupied by the input). A problem $A \subseteq \Sigma^*$ is *logspace many-one reducible* to a problem $B \subseteq \Gamma^*$ if there exists a logspace computable function $f : \Sigma^* \rightarrow \Gamma^*$ such that for all $x \in \Sigma^*$: $x \in A$ if and only if $f(x) \in B$. When talking about efficient algorithms (e.g., linear time algorithms) one has to specify the underlying computational model. Unless otherwise specified, we always refer to the *random access model* (RAM) with uniform cost measure. The latter means that arithmetic operations on arbitrary numbers can be carried out in constant time. This might seem unrealistic. On the other hand, all algorithms mentioned in this paper only have to store numbers with $O(n)$ many bits, where n is the length of the input.

We also assume some basic knowledge in automata theory. In particular, the reader should be familiar with regular languages and context-free languages. Background can be found in the classical text book [62].

2.2 Straight-line programs

Definition 1 (straight-line program (SLP)) *A straight-line program (SLP) over the terminal alphabet Γ is a context-free grammar $\mathbb{A} = (V, \Gamma, S, P)$ (V is the set of nonterminals, Γ is the set of terminals, $S \in V$ is the initial nonterminal, and $P \subseteq V \times (V \cup \Gamma)^*$ is the set of productions) such that the following two conditions hold:*

- (1) *For every $A \in V$ there exists exactly one production of the form $(A, \alpha) \in P$ for $\alpha \in (V \cup \Gamma)^*$.*
- (2) *The relation $\{(A, B) \mid (A, \alpha) \in P, B \in \text{alph}(\alpha)\}$ is acyclic.*

A production (A, α) is also written as $A \rightarrow \alpha$. Clearly, the language generated by the SLP \mathbb{A} consists of exactly one word that is denoted by $\text{eval}(\mathbb{A})$. More generally, from every nonterminal $A \in V$ we can generate exactly one word that is denoted by $\text{eval}_{\mathbb{A}}(A)$ (thus $\text{eval}(\mathbb{A}) = \text{eval}_{\mathbb{A}}(S)$). We omit the index \mathbb{A} if the underlying SLP is clear from the context.

The *derivation tree* of the SLP $\mathbb{A} = (V, \Gamma, S, P)$ is a finite rooted ordered tree, where every node is labelled with a symbol from $V \cup \Gamma$. The root is labelled with the initial nonterminal S and every node that is labelled with a symbol from Γ is a leaf of the derivation tree. A node that is labelled with a nonterminal A such that $(A \rightarrow \alpha_1 \cdots \alpha_n) \in P$ (where $\alpha_1, \dots, \alpha_n \in V \cup \Gamma$) has n children that are labeled from left to right with $\alpha_1, \dots, \alpha_n$.

The size of the SLP $\mathbb{A} = (V, \Gamma, S, P)$ is $|\mathbb{A}| = \sum_{(A, \alpha) \in P} |\alpha|$. Every SLP can be transformed in linear time into an equivalent SLP in *Chomsky normal form*, where every production has the form (A, a) with $a \in \Gamma$ or (A, BC) with $B, C \in V$.

Example 2 *Consider the SLP \mathbb{A} over the terminal alphabet $\{a, b\}$ that consists of the following productions: $A_1 \rightarrow b$, $A_2 \rightarrow a$, and $A_i \rightarrow A_{i-1}A_{i-2}$ for $3 \leq i \leq 7$. The start nonterminal is A_7 . Then $\text{eval}(\mathbb{A}) = \text{abaababaabaab}$, which is the 7th Fibonacci word. The SLP \mathbb{A} is in Chomsky normal form and $|\mathbb{A}| = 12$.*

A simple induction shows that for every SLP \mathbb{A} of size m one has $|\text{eval}(\mathbb{A})| \leq O(3^{m/3})$ [23, proof of Lemma 1]. On the other hand, it is straightforward to define an SLP \mathbb{B} in Chomsky normal form of size $2n$ such that $|\text{eval}(\mathbb{B})| \geq 2^n$. Hence, an SLP can be seen as a compressed representation of the string it generates, and exponential compression rates can be achieved in this way.

One may also allow exponential expressions of the form A^i for $A \in V$ and $i \in \mathbb{N}$ in right-hand sides of productions. Here the number i is coded binary. Such an expression can be replaced by a

sequence of ordinary productions, where the length of that sequence is bounded polynomially in the length of the binary coding of i .

The following construction from [85] is often useful for proving lower complexity bounds, see e.g. the end of Section 5.

Example 3 Let $\bar{w} = (w_1, \dots, w_n)$ be a tuple of natural numbers. For a bit vector $\bar{x} \in \{0, 1\}^n$ of length n let us define $\bar{x} \cdot \bar{w} = x_1 w_1 + x_2 w_2 + \dots + x_n w_n$. Let $\bar{1}_k$ be the constant-1 vector $(1, 1, \dots, 1)$ of length k and let $s_k = \bar{1}_k \cdot (w_1, \dots, w_k) = w_1 + \dots + w_k$ for $1 \leq k \leq n$. Let $s = s_n = w_1 + w_2 + \dots + w_n$. Finally, define the string

$$S(\bar{w}) = \prod_{\bar{x} \in \{0, 1\}^n} a^{\bar{x} \cdot \bar{w}} b a^{s - \bar{x} \cdot \bar{w}}.$$

Here, the product $\prod_{\bar{x} \in \{0, 1\}^n}$ means that we concatenate all string $a^{\bar{x} \cdot \bar{w}} b a^{s - \bar{x} \cdot \bar{w}}$ for $\bar{x} \in \{0, 1\}^n$ and the order of concatenation is the lexicographic order on $\{0, 1\}^n$, where the right-most bit has highest significance. For example, we have

$$\begin{aligned} S(2, 3, 5) &= ba^{10} a^2 b a^8 a^3 b a^7 a^5 b a^5 a^5 b a^5 a^7 b a^3 a^8 b a^2 a^{10} b \\ &= ba^{12} b a^{11} b a^{12} b a^{10} b a^{12} b a^{11} b a^{12} b. \end{aligned}$$

Let us show that there is an SLP \mathbb{A} of size polynomial in $\sum_{i=1}^n \log(w_i)$ for the string $S(\bar{w})$. Note that $\sum_{i=1}^n \log(w_i)$ is roughly the length of the binary encoding of the tuple \bar{w} . The productions of \mathbb{A} are:

$$\begin{aligned} A_1 &\rightarrow ba^{s+w_1} b \\ A_{k+1} &\rightarrow A_k a^{s-s_k+w_{k+1}} A_k \quad (1 \leq k \leq n-1) \end{aligned}$$

Let A_n be the start nonterminal of \mathbb{A} .

We prove by induction on k that

$$\text{eval}_{\mathbb{A}}(A_k) = \left(\prod_{\bar{x} \in \{0, 1\}^k \setminus \{\bar{1}_k\}} a^{\bar{x} \cdot \bar{w}_k} b a^{s - \bar{x} \cdot \bar{w}_k} \right) a^{s_k} b.$$

The case $k = 1$ is clear, since $\text{eval}_{\mathbb{A}}(A_1) = ba^{s+w_1} b = ba^s a^{s_1} b = a^0 b a^{s-0} a^{s_1} b$. For $k+1 \leq n$ we obtain the following:

$$\begin{aligned} &\left(\prod_{\bar{x} \in \{0, 1\}^{k+1} \setminus \{\bar{1}_{k+1}\}} a^{\bar{x} \cdot \bar{w}_{k+1}} b a^{s - \bar{x} \cdot \bar{w}_{k+1}} \right) a^{s_{k+1}} b = \\ &\underbrace{\left(\prod_{\bar{x} \in \{0, 1\}^k} a^{\bar{x} \cdot \bar{w}_k} b a^{s - \bar{x} \cdot \bar{w}_k} \right)}_{\text{eval}_{\mathbb{A}}(A_k) a^{s-s_k}} \underbrace{\left(\prod_{\bar{x} \in \{0, 1\}^k \setminus \{\bar{1}_k\}} a^{\bar{x} \cdot \bar{w}_k + w_{k+1}} b a^{s - \bar{x} \cdot \bar{w}_k - w_{k+1}} \right)}_{a^{w_{k+1}} \text{eval}_{\mathbb{A}}(A_k)} a^{s_{k+1}} b = \\ &\text{eval}_{\mathbb{A}}(A_k) a^{s-s_k+w_{k+1}} \text{eval}_{\mathbb{A}}(A_k) = \text{eval}_{\mathbb{A}}(A_{k+1}). \end{aligned}$$

For $k = n$ we finally get

$$\text{eval}(\mathbb{A}) = \text{eval}_{\mathbb{A}}(A_n) = \prod_{\bar{x} \in \{0, 1\}^n} a^{\bar{x} \cdot \bar{w}} b a^{s - \bar{x} \cdot \bar{w}}.$$

This concludes Example 3.

Let us state some simple algorithmic problems that can be easily solved in polynomial time:

- (1) Given an SLP \mathbb{A} (w.l.o.g. in Chomsky normal form), calculate $|\text{eval}(\mathbb{A})|$: One introduces for each nonterminal A of the SLP \mathbb{A} a variable n_A which takes values in \mathbb{N} . For a production $A \rightarrow a$ of \mathbb{A} we set $n_A = 1$, and for a production $A \rightarrow BC$ of \mathbb{A} we set $n_A = n_B + n_C$. In this way, the length $|\text{eval}(\mathbb{A})|$ can be computed with $|\mathbb{A}|$ many additions. Moreover, the number of bits of each value n_A is bounded by $O(\log(|\text{eval}_{\mathbb{A}}(A)|)) \leq O(|\mathbb{A}|)$.
- (2) Given an SLP \mathbb{A} (w.l.o.g. in Chomsky normal form) and a number $1 \leq i \leq |\text{eval}(\mathbb{A})|$, calculate $\text{eval}(\mathbb{A})[i]$ (this problem is in fact P-complete [82], see Section 10): Basically, the algorithm walks down in the derivation tree for \mathbb{A} to position i . At each stage, the algorithm stores a number p and a nonterminal A of \mathbb{A} such that $1 \leq p \leq |\text{eval}_{\mathbb{A}}(A)|$. Initially, $p = i$, and A is the initial nonterminal of \mathbb{A} . If the unique production for A is $A \rightarrow BC$, then there are two cases: If $1 \leq p \leq |\text{eval}_{\mathbb{A}}(B)|$, then we continue with position p and the nonterminal B . Otherwise, we have $|\text{eval}_{\mathbb{A}}(B)| + 1 \leq p \leq |\text{eval}_{\mathbb{A}}(A)|$, and we continue with position $p - |\text{eval}_{\mathbb{A}}(B)|$ (this number can be computed using the algorithm from the previous point) and the nonterminal C . Finally, if the unique production for A is $A \rightarrow a$, then we must have $p = 1$ and we output the terminal symbol a .
- (3) Given an SLP \mathbb{A} over the terminal alphabet Γ and a homomorphism $\rho : \Gamma^* \rightarrow \Sigma^*$, calculate an SLP \mathbb{B} such that $\text{eval}(\mathbb{B}) = \rho(\text{eval}(\mathbb{A}))$: We obtain \mathbb{B} by simply replacing every occurrence of a terminal symbol a in \mathbb{A} by the string $\rho(a)$.

Straight-line programs can also be viewed as particular *circuits*. Given an algebraic structure \mathcal{A} with operations f_1, \dots, f_n (of arbitrary arity) and constants c_1, \dots, c_m , a circuit over \mathcal{A} is a finite directed acyclic graph D , where

- every node of indegree 0 is labelled with one of the constants c_1, \dots, c_m ,
- every node of indegree $n \geq 1$ is labelled with an operation f_i of arity n ,
- the incoming edges for a node are linearly ordered, and
- there is a unique node of outdegree 0 (the output node of D).

Such a circuit computes an element of \mathcal{A} . An SLP over the terminal alphabet Γ is nothing else than a circuit over the free monoid Γ^* , where the only operation is concatenation and the constants are the symbols from the alphabet Γ . It should be noted that in the literature, the term “straight-line program” often refers to circuits over the ring of integers or, more generally, polynomial rings, see e.g. [3, 63, 70].

For some applications, an extension of SLPs called *composition systems* in [43] are useful.¹ A composition system \mathbb{A} is defined as an SLP but may also contain productions of the form $X \rightarrow Y[i, j]$ for $i, j \in \mathbb{N}$ with $1 \leq i \leq j$. Assume that the string $w = \text{eval}_{\mathbb{A}}(Y)$ is already defined. Then we let $\text{eval}_{\mathbb{A}}(X) = w[i, \max\{|w|, j\}]$. A composition system is in Chomsky normal form if all productions have the form $X \rightarrow a$, $X \rightarrow YZ$ or $X \rightarrow Y[i, j]$ for nonterminals X, Y , and Z and a terminal symbol a . The following result was shown by Hagenah in his PhD thesis [55] (in German), see also [125].

Theorem 4 *From a given composition system \mathbb{A} in Chomsky normal form with n nonterminals one can compute in time $O(n^2)$ an SLP \mathbb{B} of size $O(n^2)$ such that $\text{eval}(\mathbb{B}) = \text{eval}(\mathbb{A})$.*

3 Straight-line programs and dictionary-based compression

As already remarked in the Introduction, there is a tight relationship between straight-line programs and dictionary-based compression, in particular, the LZ77 compression scheme. There are various variants of LZ77 compression. The following variant is used for instance in [119]:

¹The formalism in [43] differs in some minor details. Moreover, composition systems are called *collage systems* in [74] and interval grammars in [55].

The *LZ77-factorization* of a non-empty word $s \in \Sigma^+$ is defined as $s = f_1 f_2 \cdots f_m$, where for all $1 \leq i \leq m$, f_i is longest non-empty prefix of $f_i f_{i+1} \cdots f_m$ that is a factor of $f_1 f_2 \cdots f_{i-1}$ in case such a prefix exists, otherwise f_i is the first symbol of $f_i f_{i+1} \cdots f_m$. Then, the *LZ77-encoding* of w is the sequence $c_1 c_2 \cdots c_m$ such that for all $1 \leq i \leq m$, either $|f_i| = 1$ and $c_i = f_i$ or $|f_i| \geq 2$, $f_i = (f_1 f_2 \cdots f_{i-1})[p, q]$ and $f_i \neq (f_1 f_2 \cdots f_{i-1})[k, k + q - p]$ for all $1 \leq k < p$.

Theorem 5 ([23, 119]) *For every SLP \mathbb{A} , the number of factors in the LZ-factorization of $\text{eval}(\mathbb{A})$ is bounded by $|\mathbb{A}|$.*

Theorem 6 ([23, 119]) *From the LZ77-factorization of a given string $w \in \Sigma^*$, one can compute an SLP of size $O(\log(|w|/m) \cdot m)$ for w in time $O(\log(|w|/m) \cdot m)$, where m is the number of factors in the LZ77-factorization of w .*

Rytter [119] uses SLPs, whose derivation tree is an AVL-tree for the proof of Theorem 6. An AVL-tree is a binary tree such that for every node v the height of the left subtree of v and the height of the right subtree of v differ by at most 1. The construction of Charikar et al. [23] is based on so called α -balanced SLPs. An SLP \mathbb{A} is α -balanced, where $0 < \alpha \leq 1/2$, if the right-hand side of every production has length two and for every production $A \rightarrow \beta\gamma$ of \mathbb{A} (where β and γ are nonterminals or terminal symbols) the following holds: If m (resp., n) is the length of the string derived from β (resp., γ), then

$$\frac{\alpha}{1 - \alpha} \leq \frac{m}{n} \leq \frac{1 - \alpha}{\alpha}.$$

Using the fact that the number of factors in the LZ77-factorization of an SLP-compressed word $\text{eval}(\mathbb{A})$ is bounded by the size of the SLP \mathbb{A} (Theorem 6) as well as the results from Section 6 on compressed pattern matching, one can also show that the LZ77-encoding of $\text{eval}(\mathbb{A})$ can be computed in polynomial time from \mathbb{A} . Hence, SLPs and LZ77-encodings are equivalent with respect to polynomial time transformations.

4 Grammar-based compression

As explained in the Introduction, in this paper we mainly deal with algorithmic problems on strings, which are already given in compressed form as an SLP. In this section, we will briefly consider the problem of compression itself, i.e., the computation of a small SLP for a given input string. This problem is known as grammar-based compression, and is considered for instance in [23, 101, 119, 121]. For a string w let $\text{opt}(w)$ be the size of a minimal SLP for w . Thus, there exists an SLP \mathbb{A} with $\text{eval}(\mathbb{A}) = w$, $|\mathbb{A}| = \text{opt}(w)$ and for every SLP \mathbb{B} with $\text{eval}(\mathbb{B}) = w$, $|\mathbb{B}| \geq |\mathbb{A}|$ holds. The following theorem provides a lower bound for grammar-based compression.

Theorem 7 ([23]) *Unless $\text{P} = \text{NP}$ there is no polynomial time algorithm with the following specification:*

- *The input consists of a string w over some alphabet Σ .*
- *The output is a grammar \mathbb{A} such that $\text{eval}(\mathbb{A}) = w$ and $|\mathbb{A}| \leq 8569/8568 \cdot \text{opt}(w)$.*

Theorem 7 is shown in [23] using a reduction from the vertex cover problem for graphs with maximal degree 3. In Theorem 7 it is important that the alphabet Σ is not fixed. It is an open problem whether there is a polynomial time algorithm that computes for a binary string $w \in \{0, 1\}^*$ an SLP \mathbb{A} such that $\text{eval}(\mathbb{A}) = w$ and $|\mathbb{A}| = \text{opt}(w)$. Some partial results can be found in [8].

The best known upper bound for grammar-based compression is provided by the following results that was shown independently by Charikar et al. [23] and Rytter [119]:

Theorem 8 ([23, 119]) *There is an $O(\log(|\Sigma|) \cdot n)$ -time algorithm that computes for a given word $w \in \Sigma^*$ of length n an SLP \mathbb{A} such that $\text{eval}(\mathbb{A}) = w$ and $|\mathbb{A}| \leq O(\log(n/\text{opt}(w)) \cdot \text{opt}(w))$ (i.e., the approximation ratio of the algorithm is $O(\log(n/\text{opt}(w)))$).*

The algorithms of Charikar et al. and Rytter follow a similar strategy. First, the LZ77-encoding of the input string w is computed. This is possible in time $O(\log(|\Sigma|) \cdot |w|)$ using suffix trees, see e.g. [53]. Let m be the number of factors in the LZ77-factorization of w . By Theorem 5, we have $m \leq \text{opt}(w)$. In a second step, the LZ77-encoding of w is transformed into an SLP for w of size at most $O(\log(|w|/m) \cdot m) \leq O(\log(|w|/\text{opt}(w)) \cdot \text{opt}(w))$ using Theorem 6.

Rytter’s technique [119] can also be used to transform an SLP \mathbb{A} for a string of length n in time $O(|\mathbb{A}| \cdot \log(n))$ into an equivalent SLP of height $O(\log(n))$, where the height of an SLP is the height of the corresponding derivation tree. Hence, there is a very efficient algorithm for making SLPs balanced.

Let us remark that the existence of a polynomial time algorithm that computes from a given string w an SLP \mathbb{A} for w of size $o(\log(|w|)/\log \log(|w|)) \cdot \text{opt}(w)$ would imply a major progress on the problem of computing shortest addition chains, see [23]. In this problem, a sequence k_1, \dots, k_m of binary encoded natural numbers is given, and one is looking for a smallest circuit over $(\mathbb{N}, +)$ such that every k_i is the value of some gate. Currently the best approximation algorithm for this problem is by Yao [143] (from 1976) and its approximation ratio is $O(\log n / \log \log n)$, where $n = k_1 + \dots + k_m$.

Some optimizations of Rytter’s algorithm can be found in [17]. Another grammar-based compression algorithm with linear running time (for a fixed alphabet) and an approximation ratio of $O(\log(n/\text{opt}(w)))$ was presented by Sakamoto [121]. His algorithm is based on the RePair compressor [78], which is another popular grammar-based compressor.

The algorithms from [23, 119, 121] all use space $\Omega(n)$ in order to generate an SLP for a string of length n , which might be a problem for large input texts. A grammar-based compression algorithm with an approximation ratio of $O(\log^2(n))$ that runs in linear time and needs no more space than the size of the output grammar was recently presented in [101].

Grammar-based compression using a weak form of context-sensitive grammars (so called Σ -sensitive grammars) is discussed in [102].

5 Compressed equality checking

The most basic task for SLP-compressed strings is equality checking: Given two SLPs \mathbb{A} and \mathbb{B} , does $\text{eval}(\mathbb{A}) = \text{eval}(\mathbb{B})$ hold? Clearly, a simple decompress-and-compare strategy is very inefficient. It takes exponential time to compute $\text{eval}(\mathbb{A})$ and $\text{eval}(\mathbb{B})$. Nevertheless a polynomial time algorithm exists. This was independently discovered by Hirshfeld, Jerrum, and Moller [61], Mehlhorn, Sundar, and Uhrig [105], and Plandowski [112].

Theorem 9 ([61, 105, 112]) *The following problem can be solved in polynomial time:*

input: Two SLPs \mathbb{A} and \mathbb{B}

output: $\text{eval}(\mathbb{A}) = \text{eval}(\mathbb{B})$?

In [105], the result is not explicitly stated but follows immediately from the main result. Mehlhorn et al. provide an efficient data structure D for a finite set of strings that supports the following operations:

- Add a string consisting of the symbol a to D .
- For two string x and y in D , add the string xy (here, one may have $x = y$) to D . Moreover, the strings x and y remain in D (the latter property is called persistence).
- For a string x in D , and a position k in x , split x into its length- k prefix and the remaining suffix and add both strings to D . Again, x remains in D .
- Check whether two strings x and y from D are identical.

The idea is to compute for each variable a signature, which is a small number, and that allows to do the equality test in constant time by comparing signatures. A signature corresponds to

a nonterminal of an SLP. The signature of a string is computed by iteratively breaking up the sequence into small blocks, which are encoded by integers using a pairing function. Adding the concatenation xy of two string x and y to the data structure needs time $O(\log n(\log m \log^* m + \log n))$ for the m^{th} operation, where n is the length of the resulting string (hence, $\log(n) \leq m$). This leads to a cubic time algorithm for checking equality of SLP-compressed strings. An improvement of the data structure from [105] was obtained by Alstrup, Brodal, and Rauhe [5]; see [4] for a long version. There, the complexity of adding the concatenation xy of two string x and y to the data structure is improved to $O(\log n(\log^* m + \log |\Sigma|))$, where n and m have the same meaning as above and Σ is the underlying alphabet. This leads to the complexity of $O(n^2 \log^* n)$ (where $n = |\mathbb{A}| + |\mathbb{B}|$) for checking equality of SLP-compressed strings, which is currently the best upper bound. One can view the algorithms from [5, 105] as efficient methods for transforming an SLP \mathbb{A} into a canonical SLP \mathbb{A}' such that $\text{eval}(\mathbb{A}) = \text{eval}(\mathbb{B})$ if and only if \mathbb{A}' and \mathbb{B}' are identical. One should note that the machine model in [5] is a RAM that allows to compute bitwise XOR and AND of machine words in constant time.

In contrast to [105], the polynomial time algorithms of Hirshfeld et al. [61] and Plandowski [112] use combinatorial properties of words, in particular the periodicity lemma of Fine and Wilf [39]. This lemma states that if p and q are periods of a string w (i.e., $w[i] = w[i+p]$ and $w[j] = w[j+q]$ for all positions $1 \leq i \leq |w| - p$ and $1 \leq j \leq |w| - q$) and $p + q \leq |w|$ then also the greatest common divisor of p and q is a period of w . The algorithm from [61] achieves a running time of $O(n^4)$, where $n = |\mathbb{A}| + |\mathbb{B}|$.

Both, Plandowski and Hirshfeld et al. use Theorem 9 as a tool to solve another problem. Plandowski derives from Theorem 9 a polynomial time algorithm for testing whether two given morphisms (between free monoids) agree on a given context-free language. Hirshfeld et al. use Theorem 9 in order to check bisimilarity of two normed context-free processes (a certain class of infinite state systems) in polynomial time.

The equality problem for SLP-compressed strings can also be viewed as a very restricted case of the equality problem for context-free languages, where the two context-free languages are singleton sets. One might try to generalize Theorem 9 to larger classes of context-free grammars. Equality of arbitrary context-free languages is a well-known undecidable problem [62]. An important class in between SLPs and general context-free grammars are *acyclic context-free grammars* (also known as non-recursive context-free grammars). These are context-free grammars such that from a nonterminal A one cannot derive in ≥ 1 steps a word containing A . In other words, we only keep condition (2) in Definition 1. Acyclic context-free grammars generate only finite sets of words. Let G be an acyclic context-free grammar of size n . Then, every word generated by G has length $2^{O(n)}$, and $|L(G)| \leq 2^{2^{O(n)}}$.

Unfortunately, Theorem 9 cannot be generalized to acyclic context-free grammars, as the following result from [64] shows:

Theorem 10 ([64]) *It is NEXPTIME-complete to check for two given acyclic context-free grammars G_1 and G_2 , whether $L(G_1) \neq L(G_2)$.*

In [108], it was shown that the problem of checking $L(G_1) \cap L(G_2) = \emptyset$ for two given acyclic context-free grammars G_1 and G_2 is PSPACE-complete.

Let us finally mention a related result of Lifshits [82], which when compared with Theorem 9 shows again the quite subtle borderline between tractability and intractability for problems on compressed strings. A function $f : \Sigma^* \rightarrow \mathbb{N}$ belongs to $\#\text{P}$ if there exists a nondeterministic polynomial time bounded Turing-machine M such that for every $x \in \Sigma^*$, $f(x)$ equals the number of accepting paths of M on input x . A function $f : \Sigma^* \rightarrow \mathbb{N}$ is $\#\text{P}$ -complete if it belongs to $\#\text{P}$ and for every $\#\text{P}$ -function $g : \Gamma^* \rightarrow \mathbb{N}$ there is a logspace computable mapping $h : \Gamma^* \rightarrow \Sigma^*$ such that $h \circ f = g$. Functions that are $\#\text{P}$ -complete are computationally very powerful. By Toda's [135] famous result, every language from the polynomial time hierarchy can be decided in deterministic polynomial time with the help of a $\#\text{P}$ -function, briefly: $\text{PH} \subseteq \text{P}^{\#\text{P}}$. For two strings u and v of the same length m , the *Hamming-distance* $d_H(u, v)$ is the number of positions $i \in \{1, \dots, m\}$ such that $u[i] \neq v[i]$.

Theorem 11 ([82]) *The mapping $(\mathbb{A}, \mathbb{B}) \mapsto d_H(\text{eval}(\mathbb{A}), \text{eval}(\mathbb{B}))$, where \mathbb{A} and \mathbb{B} are SLPs is #P-complete.*

This result can be shown using the construction from Example 3. Standard arguments show that the function that maps a tuple of binary encoded natural numbers w_1, \dots, w_n, t to the number of tuples $\bar{x} \in \{0, 1\}^n$ such that $\bar{x} \cdot (w_1, \dots, w_n) = t$ is #P-complete (this can be shown by the same reduction that reduces the 3SAT problem to the subsetsum problem); we use here the notation from Example 3. Now, given a tuple of binary encoded natural numbers w_1, \dots, w_n, t , we construct in logspace two SLPs \mathbb{A} and \mathbb{B} such that the following holds (as in Example 3 let $s = w_1 + \dots + w_n$):

$$\begin{aligned} \text{eval}(\mathbb{A}) &= \prod_{\bar{x} \in \{0, 1\}^n} a^{\bar{x} \cdot \bar{w}} b a^{s - \bar{x} \cdot \bar{w}} \\ \text{eval}(\mathbb{B}) &= (a^t b a^{s-t})^{2^n} \end{aligned}$$

The SLP \mathbb{B} can easily be constructed, and Example 3 shows how to construct \mathbb{A} (in logspace). Now, it is an easy observation that $d_H(\text{eval}(\mathbb{A}), \text{eval}(\mathbb{B}))$ is exactly the number of tuples $\bar{x} \in \{0, 1\}^n$ such that $\bar{x} \cdot (w_1, \dots, w_n) = t$.

6 Compressed pattern matching

A natural generalization of checking equality of two strings is pattern matching. In the classical *pattern matching problem* it is asked for given strings p (usually called the pattern) and t (usually called the text), whether p is a factor of t . There are dozens of linear time algorithms for this problem on uncompressed strings, most notably the well-known Knuth-Morris-Pratt algorithm [75]. It is therefore natural to ask, whether a polynomial time algorithm for pattern matching on SLP-compressed strings exists; this problem is sometimes called *fully compressed pattern matching* and is defined as follows:

input: Two SLPs \mathbb{P} and \mathbb{T}

question: Is $\text{eval}(\mathbb{P})$ a factor of $\text{eval}(\mathbb{T})$?

The first polynomial time algorithm for fully compressed pattern matching was presented in [73], further improvements with respect to the running time were achieved in [42, 67, 81, 106].

Theorem 12 ([42, 67, 73, 81, 106]) *Fully compressed pattern matching can be solved in polynomial time.*

Basically, the algorithms from [42, 67, 73, 81, 106] can be divided into two classes: Those that are based on the periodicity lemma of Fine and Wilf [39] mentioned in the previous section [42, 73, 81, 106] and Jez’s solution from [67] that does not use any non-trivial combinatorial properties of words.

An important observation that is crucial in [42, 67, 73, 81, 106] is the obvious fact that if a pattern p is a factor of $\text{eval}(\mathbb{T})$ then there exists a production $X \rightarrow YZ$ in $\text{eval}(\mathbb{T})$ such that p has an occurrence in $\text{eval}_{\mathbb{T}}(X) = \text{eval}_{\mathbb{T}}(Y)\text{eval}_{\mathbb{T}}(Z)$ that “touches” the cut (or border) position between $\text{eval}_{\mathbb{T}}(Y)$ and $\text{eval}_{\mathbb{T}}(Z)$ (called the cut of X). It is a consequence of the periodicity lemma of Fine and Wilf [39] mentioned in the previous section that the set of all starting positions of occurrences of p in $\text{eval}_{\mathbb{T}}(X)$ that touch the cut of X forms an arithmetic progression, i.e., a set of the form $\{b + i \cdot p \mid 0 \leq i < \ell\}$ (see Figure 1), which can be represented by the binary encoded triple (b, p, ℓ) . Lifshitz’s [81] algorithm for instance computes for each nonterminal A of the pattern SLP \mathbb{P} and each nonterminal X of the text SLP \mathbb{T} the arithmetic progression corresponding to the occurrences of $\text{eval}_{\mathbb{P}}(A)$ in $\text{eval}_{\mathbb{T}}(X)$ that touch the cut of X . These arithmetic progressions are computed in a bottom-up fashion spending time $O(|\mathbb{P}|)$ for each progressions, which results in the overall time bound $O(|\mathbb{P}|^2|\mathbb{T}|)$.

As already mentioned above, Jez’s algorithm for fully compressed pattern matching [67] does not use non-trivial combinatorial properties of words. Moreover, it is the currently fastest algorithm; its running time is $O((|\mathbb{T}| + |\mathbb{P}|) \log(M) \log(|\mathbb{T}| + |\mathbb{P}|))$, where M is the length of $\text{eval}(\mathbb{P})$.

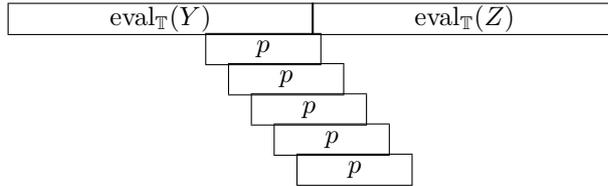


Figure 1: Occurrences of a pattern p that touch the cut of X , where $X \rightarrow YZ$

Since $\log(M) \in O(|\mathbb{P}|)$, the running time can be upper bounded by $O(n^2 \log(n))$, where n is the total input length. The basic idea of the algorithm is to rewrite the SLP for the pattern and the text simultaneously; he calls this process *recompression*. In a single round of this recompression process, maximal blocks of the form a^n (for a letter a) and length-2 words ab (for a and b different letters) become new alphabet symbols. Moreover, the combined length of the (fully decompressed) text and pattern shrinks by a constant factor in each single recompression round. Hence, the recompression process terminates after $O(\log M)$ many rounds with a pattern of length 1. Some effort is necessary in order to show that during this process the intermediate SLPs are small. Jez’s recompression technique uses some of the ideas from the equality-checking algorithms from [5, 105]. In particular the idea of breaking up a text into maximal blocks of the form a^n appears in [5, 105] too. One might say that Jez extends the techniques from [5, 105] so that pattern matching (instead of only equality checking) of SLP-compressed strings can be done efficiently.

Randomized algorithms for checking equality of compressed strings and fully compressed pattern matching are studied in [43, 129]. These algorithms are based on arithmetic modulo small prime numbers. The algorithm from [129] has a quadratic running time under the RAM model with *logarithmic cost measure*, which means that arithmetic operations on n -bit numbers need time $O(n)$. If $\text{eval}(\mathbb{A}) = \text{eval}(\mathbb{B})$ then the algorithm will correctly output “yes”; if $\text{eval}(\mathbb{A}) \neq \text{eval}(\mathbb{B})$ then the algorithm may incorrectly output “yes” with a small error probability.

In the fully compressed pattern matching problem, the term “fully” refers to the fact that both the pattern and the text are compressed. In practical applications, the pattern is usually short in comparison to the text. In such a situation, it makes sense to consider the weaker variant, where only the text is compressed, but the pattern is given explicitly. This leads to the *semi-compressed pattern matching problem* (sometimes just called *compressed pattern matching problem*), where it is asked whether an uncompressed pattern p is a factor of a compressed text $\text{eval}(\mathbb{T})$. Semi-compressed pattern matching has been studied intensively for various compression schemes, e.g. LZW (Lempel–Ziv–Welch) [7, 44] and LZ1 (a variant of LZ77 as defined in Section 3) [38, 45]. For SLP-compressed texts it is easy to show that compressed pattern matching can be solved in time $O(|p| \cdot |\mathbb{T}|)$. Assume that the text SLP \mathbb{T} is in Chomsky normal form. As remarked earlier, it suffices to check for each nonterminal X of \mathbb{T} , whether there exists an occurrence of p in $\text{eval}_{\mathbb{T}}(X)$ that touches the cut of X . For this, one computes for every nonterminal X the prefix pref_X of $\text{eval}_{\mathbb{T}}(X)$ of length $\min\{|p|, |\text{eval}_{\mathbb{T}}(X)|\}$ as well as the suffix suff_X of $\text{eval}_{\mathbb{T}}(X)$ of length $\min\{|p|, |\text{eval}_{\mathbb{T}}(X)|\}$. This is possible in time $O(|p| \cdot |\mathbb{T}|)$ by a straightforward bottom-up computation. Then, one only has to check for every production $X \rightarrow YZ$ of \mathbb{T} , whether p is a factor of $\text{suff}_Y \text{pref}_Z$, which is a string of length at most $2p$. This is again possible in time $O(|p| \cdot |\mathbb{T}|)$ using any linear time pattern matching algorithm for uncompressed strings. It should be also remarked that the currently best algorithm for semi-compressed pattern matching for LZ1 [45] first transforms the LZ1-encoding of the text into an α -balanced SLP using the method from [23] (see Theorem 6) and then runs an algorithm for semi-compressed pattern matching for α -balanced SLPs.

Several other algorithmic problems that are related to compressed pattern matching were studied. In [25, 26, 40], self-indexes for SLP-compressed strings are constructed. The goal of such a self-index is to store an SLP \mathbb{A} with small space overhead in such a form that for a given uncompressed pattern p one can quickly list all occurrences of p in $\text{eval}(\mathbb{A})$. Bille et al. [15] present an efficient algorithm for approximative compressed pattern matching for SLP-compressed strings:

Here, for a given pattern p , SLP \mathbb{T} , and $k \in \mathbb{N}$ the goal is to find all occurrences of factors of $\text{eval}(\mathbb{T})$ that have distance at most k from p with respect to a certain distance measure (e.g. Hamming distance or edit distance). The problem of computing q -gram frequencies for SLP-compressed strings is studied in [49, 50]. A q -gram is just a string of length q . An algorithm that computes in time $O(q \cdot n)$ for a Chomsky normal form SLP \mathbb{A} of size n a list with the frequencies of all q -grams occurring in $\text{eval}(\mathbb{A})$ is presented in [49]. This algorithm can be seen as a refinement of the algorithm for semi-compressed pattern matching outlined above. In [103], it is shown that the length of the longest common factor of two SLP-compressed strings can be computed in polynomial time.

7 Leaf languages and string compression

In this section, we discuss a technique that turned out to be useful for deriving lower complexity bounds for algorithmic problems on SLP-compressed strings. This technique is based on the concept of leaf languages from complexity theory. A detailed survey of leaf languages can be found in [139]. Here, we will only introduce the basic definitions necessary in order to understand the relationship to SLPs.

Let M be a nondeterministic Turing-machine. We say that M is *adequate* if it satisfies the following three properties:

- On every input M does not have an infinite computation path.
- A linear order $<$ is fixed on the set of transition tuples of M .
- To every configuration of M at most two transition tuples of M apply, i.e., every configuration has at most two successor configurations.

For an adequate Turing-machine M and an input w , the computation tree $T_M(w)$ of M on input w becomes a finite rooted ordered binary tree, where nodes are labelled with configurations of M . The root is labelled with the initial configuration on input w . Now, assume that a node v of $T_M(w)$ is labelled with the configuration c , and c has exactly two successor configurations c_1 and c_2 (the case that c has at most one successor configuration is analogous). Assume moreover that c_i results from c by applying the transition tuple t_i to c and $t_1 < t_2$. Then, v has two children v_1 and v_2 in this order and v_i is labelled with the configuration c_i . The leaf string $\text{leaf}(M, w)$ is defined as follows: Let v_1, \dots, v_m be the leaves of the ordered tree $T_M(w)$ in the natural left-to-right order, and set $b_i = 1$ (resp., $b_i = 0$) if the configuration that labels node v_i is accepting (resp., rejecting). Then, we set $\text{leaf}(M, w) = b_1 \cdots b_m$.

With a language $L \subseteq \{0, 1\}^*$, one can associate a complexity class $\text{LEAF}^{\text{P}}(L)$ as follows: A language $K \subseteq \Sigma^*$ belongs to the class $\text{LEAF}^{\text{P}}(L)$ if there exists an adequate nondeterministic polynomial time Turing-machine M with input alphabet Σ such that for every word $w \in \Sigma^*$ we have: $w \in K$ if and only if $\text{leaf}(M, w) \in L$. In this way, many complexity classes can be defined in a uniform way. Here are a few examples:

- $\text{NP} = \text{LEAF}^{\text{P}}(\{0, 1\}^*1\{0, 1\}^*)$
- $\text{coNP} = \text{LEAF}^{\text{P}}(1^*)$
- $\text{PP} = \text{LEAF}^{\text{P}}(\{x \in \{0, 1\}^* \mid |x|_1 > |x|_0\})$ (PP stands for probabilistic polynomial time.²)

Of course, the leaf language concept is not restricted to nondeterministic polynomial time Turing-machines. In the context of SLP-compression, logspace leaf language classes [65] turned out to be useful: A language $K \subseteq \Sigma^*$ belongs to the class $\text{LEAF}^{\text{log}}(L)$ if there exists an adequate

²If we view a nondeterministic Turing machine as a probabilistic machine that chooses successor configurations with uniform distribution, then PP can be defined as the class of all languages L for which there exists a probabilistic machine M such that for all inputs x : $x \in L$ if and only if M accepts x with probability $> 1/2$.

nondeterministic logspace Turing-machine M with input alphabet Σ such that for every word $w \in \Sigma^*$ we have: $w \in K$ if and only if $\text{leaf}(M, w) \in L$.

Note that the machine M together with an input w can be seen as a compressed representation of the string $\text{leaf}(M, w)$. In a certain sense, if the machine M is logspace-bounded, then this form of compression is equivalent to SLP-compression:

Theorem 13 ([88]) *The following holds:*

- *Let M be an adequate nondeterministic logspace Turing-machine. From a given input $w \in \Sigma^*$ for M we can construct in logspace an SLP \mathbb{A} over $\{0, 1\}$ such that $\text{eval}(\mathbb{A}) = \text{leaf}(M, w)$.*
- *There exists an adequate nondeterministic logspace Turing-machine M that takes as input an SLP \mathbb{A} in Chomsky normal form over the alphabet $\{0, 1\}$ and produces the leaf string $\text{leaf}(M, \mathbb{A}) = \text{eval}(\mathbb{A})$.*

The proof of this result is quite simple. For the first statement, let n be the length of the input w and assume that M works in space $\log(n)$ on an input of length n . Then the nonterminals of the SLP \mathbb{A} are all configurations of M , where the input tape contains the word w and the work tape is restricted to length $\log(n)$. There is only a polynomial number of such configurations. If the successor configurations of the configuration c are c_1, \dots, c_m and $m \geq 1$ (we must have $m \leq 2$ since M is adequate) then we add the production $c \rightarrow c_1 \cdots c_m$ to the SLP. If c is a configuration without successor configurations, then we add the production $c \rightarrow 1$ (resp., $c \rightarrow 0$) if c is an accepting (resp., rejecting) configuration. Finally, the initial nonterminal of the SLP \mathbb{A} is the initial configuration corresponding to input w . Clearly, this SLP \mathbb{A} produces the leaf string $\text{leaf}(M, w)$.

For the second statement of Theorem 13 we take the adequate nondeterministic logspace Turing-machine M that stores on the work tape a nonterminal of the input SLP \mathbb{A} ; logarithmic space suffices for this. If in a configuration c the work tape stores the nonterminal A , then M searches (deterministically) the unique production with left hand A . If this production is of the form $A \rightarrow 1$ (resp., $A \rightarrow 0$), then M terminates and accepts (resp., rejects). Otherwise, assume that the production for A has the form $A \rightarrow BC$. Then M branches nondeterministically into two successor configurations c_1 and c_2 . From c_1 (resp., c_2), M updates the content of the work tape to the nonterminal B (resp., C).

The main result of [88] shows that also the leaf strings produced by nondeterministic polynomial time Turing-machines can be produced by SLPs in a certain way. We need a few more definitions.

It is often useful to restrict adequate Turing-machines further. We say that an adequate Turing-machine M is *fully balanced* if for every input w , the computation tree $T_M(w)$ has the same number of branching nodes (i.e., nodes with exactly two children) along every maximal path. In other words, if we contract edges (v, v') such that v' is the unique child of v , then $T_M(w)$ becomes a full binary tree. Figure 2 shows a typical shape for the computation tree of a fully balanced adequate Turing-machine. For two binary strings $u = a_1 a_2 \cdots a_n$ and $v = b_1 b_2 \cdots b_n$ ($a_1, b_1, \dots, a_n, b_n \in \{0, 1\}$) we define the convolution $u \otimes v = (a_1, b_1)(a_2, b_2) \cdots (a_n, b_n)$. Finally, let $\rho : (\{0, 1\} \times \{0, 1\})^* \rightarrow \{0, 1\}^*$ be the morphism with $\rho(0, 0) = \rho(0, 1) = \varepsilon$, $\rho(1, 0) = 0$, and $\rho(1, 1) = 1$. The following result was shown in [88]:

Theorem 14 ([88]) *Let M be a fully balanced adequate polynomial time Turing-machine such that for some polynomial $p(n)$ and for every input w , every maximal path in the computation tree $T_M(w)$ has exactly $p(|w|)$ many branching nodes. From a given input $w \in \Sigma^*$ for M we can construct in logspace two SLPs \mathbb{A} and \mathbb{B} such that $|\text{eval}(\mathbb{A})| = |\text{eval}(\mathbb{B})|$ and $\text{leaf}(M, w) = \rho(\text{eval}(\mathbb{A}) \otimes \text{eval}(\mathbb{B}))$.*

The proof of Theorem 14 can be seen as a refinement of the construction from Example 3. It is based on an encoding of the subsetsum problem. In the following sections, we will see two applications of Theorem 14.

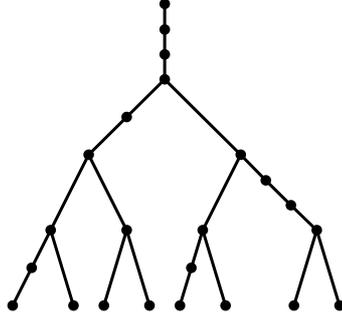


Figure 2: A computation tree of a fully balanced adequate Turing-machine

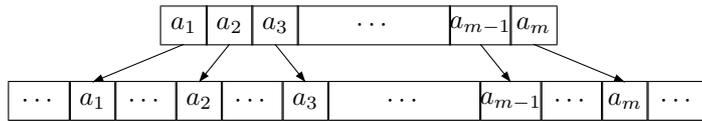


Figure 3: Subsequences

8 Compressed problems for subsequences

In many applications in computational biology, approximate occurrences of a pattern in a text are more relevant than exact matches, see e.g. [53]. A very simple formalization of an approximate occurrence is that of a *subsequence*. The string $p = a_1 \cdots a_n$ is a subsequence of $t \in \Sigma^*$, if $t \in \Sigma^* a_1 \Sigma^* \cdots a_n \Sigma^*$, see Figure 3. Note the difference between the notions of a factor and a subsequence. The *fully compressed subsequence matching problem* is defined as follows:

input: Two SLPs \mathbb{P} and \mathbb{T}

question: Is $\text{eval}(\mathbb{P})$ a subsequence of $\text{eval}(\mathbb{T})$?

When comparing fully compressed subsequence matching with fully compressed pattern matching, it is the more liberal notion of a pattern occurrence that makes the application of periodicity properties of words (in particular, the periodicity lemma of Fine and Wilf) impossible. Indeed it turned out in [82, 88] that fully compressed subsequence matching is much harder than fully compressed pattern matching.

Recall the definition of the complexity class PP (probabilistic polynomial time) from the previous section. This class contains computationally very difficult problems. A famous result of Toda [135] states that every language from the polynomial time hierarchy is polynomial time Turing-reducible to a language from PP, briefly: $\text{PH} \subseteq \text{P}^{\text{PP}}$. Moreover, PP contains the class $\text{P}_{\parallel}^{\text{NP}}$ (*parallel access to NP*), which consists of all problems that can be accepted by a deterministic polynomial time machine with access to an oracle from NP and such that furthermore all questions to the oracle are asked in parallel [110, 138].

Theorem 15 ([88]) *The fully compressed subsequence matching problem belongs to PSPACE and is PP-hard.*

Theorem 15 improves a result from [82] stating that the fully compressed subsequence matching problem is hard for the class $\text{P}_{\parallel}^{\text{NP}}$.

Membership of fully compressed subsequence matching in PSPACE is easy to prove. The straight-forward greedy linear time algorithm for uncompressed strings becomes a polynomial space bounded algorithm for SLP-compressed strings. In polynomial space, we store a position $p_{\mathbb{P}}$ in the string $\text{eval}(\mathbb{P})$ and a position $p_{\mathbb{T}}$ in the string $\text{eval}(\mathbb{T})$ (initially, $p_{\mathbb{T}} = p_{\mathbb{P}} = 1$). In each step, we search for the smallest position $p \geq p_{\mathbb{T}}$ such that $\text{eval}(\mathbb{T})[p] = \text{eval}(\mathbb{P})[p_{\mathbb{P}}]$. If we find such a position, we set $p_{\mathbb{P}} := p_{\mathbb{P}} + 1$ and $p_{\mathbb{T}} := p + 1$.

Hardness of fully compressed subsequence matching for the class PP is derived in [88] from Theorem 14. It is easy to see that there exists a fixed fully balanced adequate polynomial time Turing-machine M for which it is PP-complete to check whether the number of accepting computations of M on a given input w is at least m (a given binary coded number). Let us fix an input w for M and a binary coded number m . By Theorem 14, we can compute in logspace two SLPs \mathbb{A} and \mathbb{B} such that $|\text{eval}(\mathbb{A})| = |\text{eval}(\mathbb{B})|$ and $\text{leaf}(M, w) = \rho(\text{eval}(\mathbb{A}) \otimes \text{eval}(\mathbb{B}))$. Let $n = |\text{eval}(\mathbb{A})| = |\text{eval}(\mathbb{B})| \geq 1$; the binary encoding of this number can be computed in logspace in this particular case. For the number m we can w.l.o.g. assume that $m \leq n$ (otherwise, there cannot be m accepting computations on input w).

We define two morphisms ϕ_1 and ϕ_2 as follows:

$$\begin{aligned} \phi_1(0) &= 0^{n+1} & \phi_1(1) &= (10)^n \\ \phi_2(0) &= 0(10)^n & \phi_2(1) &= (10)^{n+1} \end{aligned}$$

It is straightforward to compute in logspace SLPs \mathbb{C} and \mathbb{D} such that

$$\text{eval}(\mathbb{C}) = \phi_1(\text{eval}(\mathbb{A}))(10)^m \quad \text{and} \quad \text{eval}(\mathbb{D}) = \phi_2(\text{eval}(\mathbb{B})).$$

Moreover, it is not too hard to show that $\text{eval}(\mathbb{C})$ is a subsequence of $\text{eval}(\mathbb{D})$ if and only if $\rho(\text{eval}(\mathbb{A}) \otimes \text{eval}(\mathbb{B}))$ contains at least m many 1's. The latter holds if and only if the number of accepting computations of M on input w is at least m . This concludes our proof sketch for the PP-hardness of fully compressed subsequence matching.

A corollary of the PP-hardness of fully compressed subsequence matching is PP-hardness of the *longest common subsequence problem* and the *shortest common supersequence problem* on SLP-compressed strings, even when both problems are restricted to two input strings. These problems have many applications e.g. in computational biology [53]. Another interesting property of the fully compressed subsequence matching problem is that it can be reduced to its own complement [82]. Hence, the fully compressed subsequence matching problem and its complementary problem have the same complexity.

As for pattern matching, also the length of the pattern in subsequence matching is quite often small in comparison to the length of the text. In such a situation it makes sense to consider the complexity of checking whether an uncompressed pattern $p \in \Sigma^*$ is a subsequence of a compressed string $\text{eval}(\mathbb{T})$. Let us call this problem the *semi-compressed subsequence matching problem*. There is a straightforward algorithm for semi-compressed subsequence matching that works in time $O(|p| \cdot |\mathbb{T}|)$, see [133, 142]. Assume that the SLP \mathbb{T} is in Chomsky normal form. Let $p = a_1 \cdots a_m$. The idea is to compute for every position $1 \leq i \leq m$ and every variable X of \mathbb{T} the length $\ell(i, X)$ of the longest prefix of $a_i \cdots a_m$ that is a subsequence of $\text{eval}_{\mathbb{T}}(X)$. The values $\ell(i, X)$ satisfy a simple recursion, which allows to compute all these values in time $O(|p| \cdot |\mathbb{T}|)$. More complex problems related to semi-compressed subsequence matching (e.g. counting the number of shortest factors of $\text{eval}(\mathbb{T})$, which contain p as a subsequence) are considered in [22, 134, 142].

The structural complexity of a generalization of semi-compressed subsequence matching was studied by Markey and Schnoebelen [100]. LOGCFL is the class of all languages that are logspace reducible to a context-free language [132]; it coincides with the class of all languages that can be recognized by a uniform boolean circuit family of polynomial size and logarithmic depth, where all \wedge -gates have bounded fan-in and all \vee -gates have unbounded fan-in [137]. It is conjectured that LOGCFL is a proper subclass of P.

Theorem 16 ([100]) *The following problem belongs to LOGCFL:*

input: Strings $p_0, p_1, \dots, p_n \in \Sigma^*$ and an SLP \mathbb{T} over Σ
question: Does $\text{eval}(\mathbb{T}) \in p_0 \Sigma^* p_1 \Sigma^* \cdots p_{n-1} \Sigma^* p_n$ hold?

9 Compressed membership problems

The complexity of membership problems for various classes of formal languages is a classical topic at the borderline between formal language theory and complexity theory. It is therefore natural to

consider membership problems in a compressed setting. For a language $L \subseteq \Sigma^*$, the *compressed membership problem* is the following decision problem:

input: An SLP \mathbb{A} over the alphabet Σ

question: Does $\text{eval}(\mathbb{A}) \in L$ hold?

An immediate corollary of Theorem 13 is the following result:

Corollary 17 *For every language $L \subseteq \{0,1\}^*$, the compressed membership problem for the language L is complete for the logspace leaf language class $\text{LEAF}^{\log}(L)$.*

If we fix some formalism for describing languages, we may also consider a uniform version of the compressed membership problem, where the language L is part of the input. In this section, we will survey result on the complexity of this problem for various kinds of automata, regular expression, and grammars. Let us start with regular languages.

9.1 Regular languages

Let us start with regular languages, the most basic languages in formal language theory.

Theorem 18 ([10, 100, 115]) *The following holds:*

- For a given nondeterministic finite automaton \mathcal{A} with n states and an SLP \mathbb{A} it can be checked in time $O(|\mathbb{A}| \cdot n^3)$ whether $\text{eval}(\mathbb{A}) \in L(\mathcal{A})$.
- There exists a fixed regular language with a P-complete compressed membership problem.

The first statement is easy to show. Let \mathbb{A} be an SLP in Chomsky normal form over the alphabet Σ and let \mathcal{A} be a nondeterministic finite automaton with state set Q and input alphabet Σ . W.l.o.g. assume that $Q = \{1, \dots, n\}$. The set of transition tuples of \mathcal{A} can be represented by a bunch of boolean matrices $B_a \in \{0,1\}^{n \times n}$, one for each input symbol $a \in \Sigma$. Entry (i, j) of B_a is 1 if and only if there is an a -labelled transition from state i to state j . Clearly, there is a path from state i to state j that is labelled with the word $a_1 \cdots a_m$ ($a_1, \dots, a_m \in \Sigma$) if and only if entry (i, j) of the product matrix $B_{a_1} B_{a_2} \cdots B_{a_m}$ is 1. This simple observation allows us to compute for each nonterminal C of the SLP \mathbb{A} a boolean matrix B_C such that entry (i, j) of B_C is 1 if and only if there is an $\text{eval}_{\mathbb{A}}(C)$ -labelled path from state i to state j : If $C \rightarrow a \in \Sigma$ is a production of \mathbb{A} , we set $B_C = B_a$. If $C \rightarrow DE$ is a production of \mathbb{A} , then we set $B_C = B_D B_E$. If S is the initial nonterminal of \mathbb{A} , then $\text{eval}(\mathbb{A}) \in L(\mathcal{A})$ if and only if in the matrix B_S there is a 1-entry at a position (i, j) , where i is an initial state of \mathcal{A} and j is a final state of \mathcal{A} . In this way, we can check whether $\text{eval}(\mathbb{A}) \in L(\mathcal{A})$ using $|\mathbb{A}|$ many multiplications of boolean $(n \times n)$ -matrices, which leads to the time bound $O(|\mathbb{A}| \cdot n^3)$. Actually, by using a better algorithm for boolean matrix multiplication than the standard n^3 school method, the time bound $O(|\mathbb{A}| \cdot n^3)$ can be improved. For about 20 years, the Coppersmith–Winograd algorithm was the asymptotically best algorithm for matrix multiplication with a complexity of $O(n^{2.3755})$. This bound was recently improved by Williams [29] to $O(n^{2.3727})$.

For a given *deterministic* finite automaton \mathcal{A} with n states and an SLP \mathbb{A} , one can verify $\text{eval}(\mathbb{A}) \in L(\mathcal{A})$ in time $O(|\mathbb{A}| \cdot n)$ [115]. Instead of $|\mathbb{A}|$ many boolean matrix multiplications, only $|\mathbb{A}|$ compositions of functions from $\{1, \dots, n\}$ to $\{1, \dots, n\}$ are necessary in the deterministic case. In [12] the following generalization has been shown: From a given SLP \mathbb{A} over an alphabet Σ and a deterministic finite state transducer τ with input alphabet Σ and n states, one can compute in time $O(|\mathbb{A}| \cdot n)$ an SLP \mathbb{B} of size $O(|\mathbb{A}| \cdot n)$ for the output string $\tau(\text{eval}(\mathbb{A}))$.

The second statement of Theorem 18 is stated in [100]. The proof is based on a result from [10], which states that there exists a fixed finite monoid M (in fact, any finite non-solvable group can be taken for M) such that the circuit evaluation problem for M is P-complete. The latter problem asks whether a given circuit over M (i.e., a circuit, where all input gates are labelled with constants from M and all internal gates are labelled with the multiplication operation of M) evaluates to a given element of M . A circuit over M can be viewed as an SLP that generates

a word over the alphabet M . The second statement of Theorem 18 follows, since the set of all words over the alphabet M that evaluate to a particular element of M is a regular language. P-hardness of the circuit evaluation problem for a finite non-solvable group is shown in [10] using a technique of Barrington [9]. In this paper, Barrington proved that the word problem for every finite non-solvable group is complete for the circuit complexity class NC^1 .

9.2 Compressed membership for regular expressions

When the regular language is not given by an automaton but by a regular expression, the complexity of the compressed membership problem depends on the operators that we allow in expressions. The same phenomenon is already well-known for uncompressed strings. For instance, whereas the membership problem for ordinary regular expressions (where only the Kleene-operators $\cdot, \cup, *$ are allowed) is NL-complete [69], membership for semi-extended regular expressions (where in addition also intersection is allowed) is LOGCFL-complete [111].

Let us be a bit more formal: For a set \mathcal{C} of language operations, \mathcal{C} -expressions over the alphabet Σ are built up from constants in $\Sigma \cup \{\varepsilon\}$ using the operations from \mathcal{C} . Thus, ordinary regular expressions are $\{\cdot, \cup, *\}$ -expressions. The language defined by ρ is $L(\rho)$. The length $|\rho|$ of an expression is defined as follows: For $\rho \in \Sigma \cup \{\varepsilon\}$ set $|\rho| = 1$. If $\rho = \text{op}(\rho_1, \dots, \rho_n)$, where op is an n -ary operator, we set $|\rho| = 1 + |\rho_1| + \dots + |\rho_n|$. For a set \mathcal{C} of language operations, the *compressed membership problem for \mathcal{C}* is the following computational problem:

input: A \mathcal{C} -expression ρ and an SLP \mathbb{A}

question: Does $\text{eval}(\mathbb{A}) \in L(\rho)$ hold?

Beside the Kleene-operators $\cdot, \cup, *$, we also consider intersection (\cap), complement (\neg), squaring (2 , where $L^2 = \{uv \mid u, v \in L\}$), and shuffle (\parallel). The latter operator is defined as follows: For words $u, v \in \Sigma^*$ let

$$u \parallel v = \{u_0 v_0 u_1 v_1 \cdots u_n v_n \mid n \geq 0, u = u_0 \cdots u_n, v = v_0 \cdots v_n\}.$$

For $L, K \subseteq \Sigma^*$ let $L \parallel K = \bigcup_{u \in L, v \in K} u \parallel v$. It is well known that the class of regular languages is closed under $\cap, \neg, ^2$, and \parallel , but each of these operators leads to more succinct representations of regular languages. Operator sets that received particular interest in the literature are $\{\cdot, \cup, \neg\}$ (star-free expressions), $\{\cdot, \cup, *, \cap\}$ (semi-extended regular expressions), $\{\cdot, \cup, *, \neg\}$ (extended regular expressions), and $\{\cdot, \cup, *, \parallel\}$ (shuffle-expressions).

The next theorem gives a rather complete picture on the complexity of compressed membership problems for regular expressions. It characterizes the complexity for all operator sets

$$\{\cdot, \cup\} \subseteq \mathcal{C} \subseteq \{\cdot, \cup, *, \cap, \neg, \parallel, ^2\}.$$

The class $\text{ATIME}(\exp(n), O(n))$ denotes the class of all problems that can be solved on an alternating Turing-machine in exponential time, but the number of alternations (i.e., transitions between an existential and a universal state or vice versa) has to be bounded by $O(n)$ on every computation path. Completeness results for $\text{ATIME}(\exp(n), O(n))$ are typical for logical theories [28], but we are not aware of any other completeness results for this class in formal language theory.

Theorem 19 ([86]) *Let $\{\cdot, \cup\} \subseteq \mathcal{C} \subseteq \{\cdot, \cup, *, \cap, \neg, \parallel, ^2\}$. The compressed membership problem for \mathcal{C} is*

- (1) in P if $\mathcal{C} \subseteq \{\cdot, \cup, \cap\}$ or $\mathcal{C} \subseteq \{\cdot, \cup, *\}$,
- (2) NP-complete if $\{\cdot, \cup, \parallel\} \subseteq \mathcal{C} \subseteq \{\cdot, \cup, \cap, \parallel\}$,
- (3) PSPACE-complete if (i) $\neg \notin \mathcal{C}$ or $\parallel \notin \mathcal{C}$ and (ii) \mathcal{C} contains $\{\cdot, \cup, \neg\}$ or $\{\cdot, \cup, ^2\}$ or $\{\cdot, \cup, *, \cap\}$ or $\{\cdot, \cup, *, \parallel\}$, and
- (4) $\text{ATIME}(\exp(n), O(n))$ -complete if $\{\cdot, \cup, \neg, \parallel\} \subseteq \mathcal{C}$.

The first statement (membership in P for $\mathcal{C} \subseteq \{\cdot, \cup, \cap\}$ or $\mathcal{C} \subseteq \{\cdot, \cup, *\}$) is clear: The case $\mathcal{C} \subseteq \{\cdot, \cup, *\}$ reduces to the uniform compressed membership problem for nondeterministic finite automata (since $\{\cdot, \cup, *\}$ -expressions can be transformed in polynomial time into nondeterministic finite automata), see Theorem 18. The case $\mathcal{C} \subseteq \{\cdot, \cup, \cap\}$ is also clear, since for a $\{\cdot, \cup, \cap\}$ -expression ρ the length of every word in $L(\rho)$ is bounded by $|\rho|$. Hence, the compressed membership problem for $\{\cdot, \cup, \cap\}$ reduces to the (uncompressed) membership problem for $\{\cdot, \cup, \cap\}$ -expressions, which can be solved in polynomial time. By the same argument, the compressed membership problem for $\{\cdot, \cup, \cap, \parallel\}$ reduces to the (uncompressed) membership problem for $\{\cdot, \cup, \cap, \parallel\}$ -expressions. But the latter problem is known to be NP-complete; in fact this already holds for $\{\cdot, \parallel\}$ -expressions [104, 140], which yields statement (2) of Theorem 19. The upper complexity bounds in statement (3) and (4) of Theorem 19 can be shown by rather straightforward algorithms (using for (3) the well-known correspondance between PSPACE and alternating polynomial time).

Let us now turn to the lower bounds in (3) and (4) from Theorem 19. PSPACE-hardness for $\{\cdot, \cup, \neg\}$ is shown in [86] by a reduction from the quantified boolean satisfiability problem; the reduction is mainly taken from [99], where it was shown that model checking the temporal logic LTL on SLP-compressed strings is PSPACE-complete. PSPACE-hardness for $\{\cdot, \cup, \neg\}$ and $\{\cdot, \cup, *, \cap\}$ is shown by generic reductions from the acceptance problem for a polynomial space bounded machine. Finally, the PSPACE lower bound for $\{\cdot, \cup, *, \parallel\}$ is proved by a reduction from the intersection nonemptiness problem for regular expressions (over the standard operator set $\{\cdot, \cup, *\}$), which is a well-known PSPACE-complete problem [77]. The ATIME($\exp(n), O(n)$)-hardness part from (4) uses a corresponding result from [89] on the model-checking problem for monadic second-order logic over SLP-compressed words over a unary alphabet.

9.3 Compressed finite automata

A *compressed nondeterministic finite automaton* (CNFA for short) is a tuple $\mathcal{A} = (Q, \Sigma, \delta, q_0, F)$, 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 final states, and δ is a finite set of transitions of the form (p, \mathbb{A}, q) , where p and q are states and \mathbb{A} is an SLP over Σ . The size of \mathcal{A} is $|\mathcal{A}| = |Q| + \sum_{(p, \mathbb{A}, q) \in \delta} |\mathbb{A}|$. We say that a word w labels a path from state p to state q in \mathcal{A} if there exists a sequence of transitions

$$(p_0, \mathbb{A}_0, p_1), (p_1, \mathbb{A}_1, p_2), \dots, (p_{n-1}, \mathbb{A}_{n-1}, p_n) \in \delta$$

($n \geq 0$) such that $p_0 = p$, $p_n = q$, and $w = \text{eval}(\mathbb{A}_0) \cdots \text{eval}(\mathbb{A}_{n-1})$. The language $L(\mathcal{A}) \subseteq \Sigma^*$ is the set of all words that label a path from the initial state q_0 to some final state $q_f \in F$. A *compressed deterministic finite automaton* (CDFA for short) is a CNFA $\mathcal{A} = (Q, \Sigma, \delta, q_0, F)$ such that for every state $p \in Q$, and all transitions $(p, \mathbb{A}_1, q_1), (p, \mathbb{A}_2, q_2) \in \delta$, neither $\text{eval}(\mathbb{A}_1)$ is a prefix of $\text{eval}(\mathbb{A}_2)$, nor $\text{eval}(\mathbb{A}_2)$ is a prefix of $\text{eval}(\mathbb{A}_1)$. The *compressed membership problem for CNFAs* (resp., *CDFAs*) is the following decision problem:

input: A CNFA (resp., CDFA) \mathcal{A} and an SLP \mathbb{A}

question: Does $\text{eval}(\mathbb{A}) \in L$ hold?

Theorem 20 ([66]) *The compressed membership problem for CNFAs is NP-complete, whereas the compressed membership problem for CDFAs is P-complete.*

In his proof of Theorem 20, Jez developed his recompression technique that he latter applied in [67] to get his $O((|\mathbb{T}| + |\mathbb{P}|) \log(M) \log(|\mathbb{T}| + |\mathbb{P}|))$ -algorithm for fully compressed pattern matching, see Section 6.

The *order* of a word u is the largest number n such that u can be written as $u = v^n w$, where w is a prefix of v . If $\mathbb{A}_1, \dots, \mathbb{A}_n$ is a list of all SLPs that occur as labels in the CNFA \mathcal{A} , then we set

$$\text{ord}(\mathcal{A}) = \max\{\text{ord}(\text{eval}(\mathbb{A}_i)) \mid 1 \leq i \leq n\}.$$

Note that $\text{ord}(\mathcal{A})$ is in general exponential in the size of \mathcal{A} . By the following result, the compressed membership problem for CNFAs \mathcal{A} with small $\text{ord}(\mathcal{A})$ can be solved efficiently.

Theorem 21 ([93]) *Given a CNFA \mathcal{A} and an SLP \mathbb{A} , we can check $\text{eval}(\mathbb{A}) \in L(\mathcal{A})$ in time polynomial in $|\mathbb{A}|$, $|\mathcal{A}|$, and $\text{ord}(\mathcal{A})$.*

9.4 Context-free languages

In this section, we consider the compressed membership problem for context-free languages. It is easy to see that for every context-free language the compressed membership problem can be solved in polynomial space. This result even holds in a uniform setting, where the input consists of a context-free grammar G and an SLP \mathbb{A} , and it is asked whether $\text{eval}(\mathbb{A}) \in L(G)$:

Theorem 22 ([85]) *For a given context-free grammar G and an SLP \mathbb{A} it can be checked in polynomial space whether $\text{eval}(\mathbb{A}) \in L(G)$.*

First, one has to transform the grammar G into Chomsky normal form; this is possible in polynomial time. By [48] the uniform (uncompressed) membership problem for context-free grammars in Chomsky normal form can be solved in $\text{DSPACE}(\log^2(|G| + |w|))$, where $|G|$ is the size of the input grammar (in Chomsky normal form) and w is the word that has to be tested for membership. This algorithm can be used to check for a given context-free grammar G in Chomsky normal form and an SLP \mathbb{A} in $\text{DSPACE}(\log^2(|G| + 2^{|\mathbb{A}|}))$, i.e., in polynomial space, whether $\text{eval}(\mathbb{A}) \in L(G)$.

In [85], a fixed deterministic context-free language with a PSPACE -complete compressed membership problem was constructed using a generic reduction from the acceptance problem for polynomial space machines. Alternatively, the existence of such a language can also be deduced from Corollary 17 and the following result from [65]: There is a deterministic context-free language L such that $\text{PSPACE} = \text{LEAF}^{\log}(L)$. In [21], this result was sharpened: There even exists a 1-turn deterministic context-free language L such that $\text{PSPACE} = \text{LEAF}^{\log}(L)$. A deterministic pushdown automaton P is 1-turn, if for every input, the unique computation of P can be divided into two phases: In a first phase, the height of the pushdown does not decrease in every step, whereas in the second phase, the height of the pushdown does not increase in every step. Moreover, in [21] also a deterministic 1-counter language L (i.e., a context-free language that is accepted by a deterministic pushdown automaton with a unary pushdown alphabet) with $\text{PSPACE} = \text{LEAF}^{\log}(L)$ is constructed. Hence, there also exist fixed 1-turn deterministic context-free languages and deterministic 1-counter languages with a PSPACE -complete compressed membership problem.

In [88] it was shown that even the important subclass of visibly pushdown languages [6] contains languages with a PSPACE -complete compressed membership problem. Let us first define visibly pushdown automata and the associated languages. Let Σ_c and Σ_r be two disjoint finite alphabets (call symbols and return symbols) and let $\Sigma = \Sigma_c \cup \Sigma_r$. A *visibly pushdown automaton* (VPA) [6] over (Σ_c, Σ_r) is a tuple $V = (Q, q_0, \Gamma, \perp, \Delta, F)$, where Q is a finite set of states, $q_0 \in Q$ is the initial state, $F \subseteq Q$ is the set of final states, Γ is the finite set of stack symbols, $\perp \in \Gamma$ is the initial stack symbol, and

$$\Delta \subseteq (Q \times \Sigma_c \times Q \times (\Gamma \setminus \{\perp\})) \cup (Q \times \Sigma_r \times \Gamma \times Q)$$

is the set of transitions.³ A configuration of V is a triple from $Q \times \Sigma^* \times (\Gamma \setminus \{\perp\})^* \perp$. For two configurations (p, au, v) and (q, u, w) (with $a \in \Sigma, u \in \Sigma^*$) we write $(p, au, v) \Rightarrow_V (q, u, w)$ if one of the following three cases holds:

- $a \in \Sigma_c$ and $w = \gamma v$ for some $\gamma \in \Gamma$ with $(p, a, q, \gamma) \in \Delta$
- $a \in \Sigma_r$ and $v = \gamma w$ for some $\gamma \in \Gamma$ with $(p, a, \gamma, q) \in \Delta$
- $a \in \Sigma_r, u = v = \perp$, and $(p, a, \perp, q) \in \Delta$.

The language $L(V)$ is defined as

$$L(V) = \{w \in \Sigma^* \mid \exists f \in F, u \in (\Gamma \setminus \{\perp\})^* \perp : (q_0, w, \perp) \Rightarrow_V^* (f, \varepsilon, u)\}.$$

The VPA V is deterministic if for every $p \in Q$ and $a \in \Sigma$ the following hold:

³In [6], the input alphabet may also contain internal symbols, on which the automaton does not touch the stack at all. For our lower bound, we will not need internal symbols.

- If $a \in \Sigma_c$, then there is at most one pair $(q, \gamma) \in Q \times \Gamma$ with $(p, a, q, \gamma) \in \Delta$.
- If $a \in \Sigma_r$, then for every $\gamma \in \Gamma$ there is at most one $q \in Q$ with $(p, a, \gamma, q) \in \Delta$.

A VPA V is called a 1-turn VPA, if $L(V) \subseteq \Sigma_c^* \Sigma_r^*$. In this case $L(V)$ is called a *linear visibly pushdown language*.

Visibly pushdown automata and their languages have many decidability results and closure results that do not hold for general context-free languages. For instance, for every VPA V there exists a deterministic VPA V' with $L(V) = L(V')$, the class of visibly pushdown languages is effectively closed under boolean operations, and equivalence, inclusion, and intersection non-emptiness are decidable for VPAs, see [6] where also precise complexity results are shown. Moreover, also the membership problem seems to be easier for visibly pushdown languages than for general context-free languages. By a classical result from [51], there exists a context-free language with a LOGCFL-complete membership problem. For visibly pushdown languages the complexity of the membership problem decreases to the circuit complexity class NC^1 [35] and is therefore of the same complexity as for regular languages [9] (complexity theorists conjecture that NC^1 is a proper subclass of LOGCFL). In contrast to this, by the following theorem, compressed membership is in general PSPACE-complete even for linear visibly pushdown languages, whereas it is P-complete for regular languages (Theorem 18):

Theorem 23 ([88]) *There is a fixed 1-turn visibly pushdown language with a PSPACE-complete compressed membership problem.*

This result can be deduced from Theorem 14. By the main result of [60] there exists a regular language $K \subseteq \{0, 1\}^*$ such that for every language L in PSPACE there exists a fully balanced adequate polynomial time Turing-machine M with the following properties:

- There is a polynomial $p(n)$ such that for every input w , every maximal path in the computation tree $T_M(w)$ has exactly $p(|w|)$ many branching nodes.
- For every input w , $w \in L$ if and only if $\text{leaf}(M, w) \in K$.⁴

Hence, by Theorem 14, the following problem is PSPACE-hard:

input: Two SLPs \mathbb{A} and \mathbb{B} over the terminal alphabet $\{0, 1\}$ such that $|\text{eval}(\mathbb{A})| = |\text{eval}(\mathbb{B})|$.

question: Does $\text{eval}(\mathbb{A}) \otimes \text{eval}(\mathbb{B}) \in \rho^{-1}(K)$ hold?

Since the class of regular languages is closed under inverse morphisms, the language $\rho^{-1}(K) \subseteq (\{0, 1\} \times \{0, 1\})^*$ is regular too. Fix a deterministic finite automaton D for $\rho^{-1}(K)$. Let $\bar{0}, \bar{1}$ be two new symbols. For a word $w = a_1 \cdots a_n$ with $a_1, \dots, a_n \in \{0, 1\}$ let $\bar{w} = \bar{a}_n \cdots \bar{a}_1$. It remains to construct a 1-turn visibly pushdown automaton P such that

$$L(P) = \{\bar{u}v \mid u, v \in \{0, 1\}^*, |u| = |v|, u \otimes v \in L(D)\}$$

(then, we get $\text{eval}(\mathbb{A}) \otimes \text{eval}(\mathbb{B}) \in \rho^{-1}(K)$ if and only if $\overline{\text{eval}(\mathbb{A})} \text{eval}(\mathbb{B}) \in L(P)$). This is straightforward: While reading a symbol from $\{\bar{0}, \bar{1}\}$, the visibly pushdown automaton stores the symbol on the stack. When reading a symbol $a \in \{0, 1\}$, the automaton checks if the topmost stack symbol is \bar{a} . If not, the automaton blocks, otherwise it pops \bar{a} from the stack.

10 Querying compressed strings

One of the simplest questions one can ask about a given string is whether a certain position in the string carries a certain symbol. Let us formally define the *compressed querying problem* as follows:

input: An SLP \mathbb{A} over an alphabet Σ , a binary-coded number $1 \leq i \leq |\text{eval}(\mathbb{A})|$, and a symbol $a \in \Sigma$.

question: Does $\text{eval}(\mathbb{A})[i] = a$ hold?

⁴The proof of this result uses again Barrington's technique mentioned at the end of Section 9.1.

Theorem 24 ([82]) *Compressed querying is P-complete.*

We already argued in Section 2.2 that compressed querying can be solved in polynomial time. P-completeness of compressed querying is shown in [82] by a reduction from the P-complete *super increasing subsetsum problem* [72], which is the following problem:

input: Binary coded natural numbers t, w_1, \dots, w_n such that $w_k < \sum_{i=1}^{k-1} w_i$ for all $1 \leq k \leq n$.
question: Do there exist $b_1, \dots, b_n \in \{0, 1\}$ such that $t = \sum_{i=1}^n b_i w_i$?

Assume that w_1, \dots, w_n are binary coded natural numbers such that $w_k < \sum_{i=1}^{k-1} w_i$ for all $1 \leq k \leq n$. In [82], it is shown how to construct in logspace an SLP \mathbb{A} over the alphabet $\{a, b\}$ such that $|\text{eval}(\mathbb{A})| = 1 + \sum_{i=1}^n w_i$ and for all $0 \leq m \leq \sum_{i=1}^n w_i$ the following holds: $\text{eval}(\mathbb{A})[m+1] = b$ if and only if there exist $b_1, \dots, b_n \in \{0, 1\}$ such that $m = \sum_{i=1}^n b_i w_i$. The construction is very similar to those from Example 3.

P-completeness of compressed querying does not imply that there do not exist highly efficient polynomial time algorithms for this problem. When considering efficient algorithms for compressed querying it makes sense to distinguish between *preprocessing time* and *query time*: During preprocessing time, an auxiliary data structure is constructed from the input SLP \mathbb{A} that allows to compute in a second phase for any given binary coded number $1 \leq i \leq |\text{eval}(\mathbb{A})|$ the symbol $\text{eval}(\mathbb{A})[i]$ within the query time bound. Hence, the preprocessing phase is executed only once.

The straightforward algorithm (see Section 2.2) that first computes the lengths of all strings generated by nonterminals of \mathbb{A} and then walks down the derivation tree of \mathbb{A} needs

- preprocessing time $\Theta(|\mathbb{A}|)$ in order to compute the lengths of all strings generated by nonterminals of \mathbb{A} , followed by
- query time $\Theta(h)$, where h is the height of \mathbb{A} , for walking down the derivation tree of \mathbb{A} .

As usual we assume here the RAM model with uniform cost measure. Note that the extra space in order to store the lengths of all strings generated by nonterminals of \mathbb{A} is $\Theta(|\mathbb{A}|)$.

The query time of this algorithm can be reduced from $\Theta(h)$ to $\Theta(\log(n))$ by making the SLP \mathbb{A} balanced using the technique from [119]. As mentioned in Section 4, this needs (preprocessing) time $O(|\mathbb{A}| \cdot \log(n))$; moreover, the best bound for the size of the resulting balanced SLP is $O(|\mathbb{A}| \cdot \log(n))$. Hence, the extra space increases from $\Theta(|\mathbb{A}|)$ to $O(|\mathbb{A}| \cdot \log(n))$. The following result from [15] combines a query time of $O(\log(n))$ with a preprocessing time and extra space of $O(|\mathbb{A}|)$:

Theorem 25 ([15]) *From a given SLP \mathbb{A} such that $n = |\text{eval}(\mathbb{A})|$ one can compute on a RAM in (preprocessing) time $O(|\mathbb{A}|)$ a data structure of size $O(|\mathbb{A}|)$ that allows to compute in (query) time $O(\log(n))$ the symbol $\text{eval}(\mathbb{A})[i]$ for a given binary coded number $1 \leq i \leq n$.*

Some lower bounds on the preprocessing and query time for the compressed querying problem can be found in [24].

11 Compressed word problems for groups

In recent years, techniques for SLP-compressed strings were successfully applied in order to get more efficient algorithms for problems in combinatorial group theory. In this section, we briefly survey some of these results. We will restrict our consideration to the compressed word problem and its application to the word problem for automorphism groups. Background in combinatorial group theory can be found for instance in [96, 131].

Let G be a finitely generated group, and let Σ be a finite generating set for G . This means that every element of G can be written as a product of elements from $\Sigma \cup \Sigma^{-1}$. The *word problem* of G with respect to Σ is the following decision problem:

input: A word $w \in (\Sigma \cup \Sigma^{-1})^*$

question: Does $w = 1$ hold in the group G ?

It is well known and easy to see that if Γ is another finite generating set for G , then the word problem for G with respect to Σ is logspace many-one reducible to the word problem for G with respect to Γ . This justifies to speak just of the word problem for the group G .

The word problem was introduced in the pioneering work of Dehn from 1910 [32] in relation with topological questions. Novikov [109] and independently Boone [16] constructed examples of finitely presented groups (i.e., finitely generated groups with only finitely many defining relations) with an undecidable word problem. Despite this negative result, many natural classes of groups with decidable word problems were found. Prominent examples are for instance finitely generated linear groups, automatic groups [36], and one-relator groups. With the advent of computational complexity theory, also the complexity of word problems became an active research area. For instance, it was shown that for a finitely generated linear group the word problem can be solved in logarithmic space [84, 130], that automatic groups have polynomial time solvable (in fact quadratic) word problems [36], and that the word problem for a one-relator group is primitive recursive [19].

For a group G let $\text{Aut}(G)$ be the automorphism group of G . In general, $\text{Aut}(G)$ is not necessarily finitely generated even if G is finitely generated. For the investigation of the complexity of the word problem of a finitely generated subgroup of $\text{Aut}(G)$, a compressed variant of the word problem for G turned out to be useful. Assume again that G is finitely generated by Σ . The *compressed word problem* of G with respect to Σ is the following computational problem:

input: An SLP \mathbb{A} over the alphabet $\Sigma \cup \Sigma^{-1}$

question: Does $\text{eval}(\mathbb{A}) = 1$ hold in G ?

It is easy to see that also for the compressed word problem the complexity does not depend on the chosen generating set, which allows to speak of the compressed word problem for the group G . The compressed word problem for G can be also viewed as the following question: Does a given circuit over G evaluate to the identity element of G . One of the main applications of the compressed word problem to “non-compressed” decision problems is the following result:

Theorem 26 ([95, 125]) *Let H be a finitely generated subgroup of $\text{Aut}(G)$. Then the word problem for H is logspace many-one reducible to the compressed word problem for G .*

The proof of Theorem 26 is quite simple. Given a sequence of H -generators (and hence automorphisms of G) $\varphi_1, \varphi_2, \dots, \varphi_n$, it has to be checked whether the composition $\varphi = \varphi_1 \circ \varphi_2 \circ \dots \circ \varphi_n$ is the identity mapping on G . But this is equivalent to $\varphi(a) = a$ for each generator $a \in \Sigma$ of G . Now, one can easily construct in logspace from the sequence $\varphi_1, \varphi_2, \dots, \varphi_n$ and a generator $a \in \Sigma$ an SLP \mathbb{A}_a over the alphabet $\Sigma \cup \Sigma^{-1}$ such that $\text{eval}(\mathbb{A}_a)$ represents the group element $\varphi(a)$. Hence, one has to check for all $a \in \Sigma$, whether $\text{eval}(\mathbb{A}_a) = a$ in the group G , which is an instance of the compressed word problem for G .

Clearly, for a finite group the compressed word problem can be solved in polynomial time. The following theorem is an immediate corollary of the results from [10].

Theorem 27 ([10]) *Let G be a finite group.*

- *If G is not solvable, then the compressed word problem for G is P-complete.*
- *If G is solvable, then the compressed word problem for G belongs to the class $\text{DET} \subseteq \text{NC}^2$.⁵*

So, for instance, the compressed word problem for the symmetric group on 5 elements (a non-solvable group) is P-complete.

Compressed word problems for infinite groups were considered for the first time in [85]. For free groups, the following result was shown:

Theorem 28 ([85]) *The compressed word problem for a finitely generated free group F of rank r can be solved in polynomial time. If $r > 1$ then this problem is P-complete.*

⁵DET is the class of all problems that are NC^1 -reducible to the problem of computing the determinant of a given integer matrix.

The proof for P-hardness for $r > 1$ uses a construction of Robinson [117], where it is shown that the (uncompressed) word problem for the free group of rank 2 is hard for the circuit complexity class NC^1 . Let us sketch a polynomial time algorithm for the compressed word problem for a finitely generated free group F . Let Σ be a free set of generators for F and let \mathbb{A} be an SLP over the terminal alphabet $\Sigma \cup \Sigma^{-1}$. By Theorem 4 it suffices to construct a composition system \mathbb{B} over the same alphabet that generates the reduced normal form of $\text{eval}(\mathbb{A})$. This normal form is obtained by iteratively deleting occurrences of subwords aa^{-1} with $a \in \Sigma \cup \Sigma^{-1}$ from $\text{eval}(\mathbb{A})$ (the order in which cancelation is done is not relevant). The composition system \mathbb{B} contains for each variable X of \mathbb{A} a variable X' such that $\text{eval}_{\mathbb{B}}(X')$ is the reduced normal form of $\text{eval}_{\mathbb{A}}(X)$. Let $X \rightarrow YZ$ be a production of \mathbb{A} and assume that \mathbb{B} already contains enough productions such that $\text{eval}_{\mathbb{B}}(Y')$ (resp., $\text{eval}_{\mathbb{B}}(Z')$) is the reduced normal form of $\text{eval}_{\mathbb{A}}(Y)$ (resp., $\text{eval}_{\mathbb{A}}(Z)$). Let $\text{eval}_{\mathbb{B}}(Y') = a_1 a_2 \cdots a_m$ and $\text{eval}_{\mathbb{B}}(Z') = b_1 b_2 \cdots b_k$. In the concatenation $\text{eval}_{\mathbb{B}}(Y')\text{eval}_{\mathbb{B}}(Z')$ cancellation may only occur at the border between $\text{eval}_{\mathbb{B}}(Y')$ and $\text{eval}_{\mathbb{B}}(Z')$. We have to compute the length n of the longest suffix of $\text{eval}_{\mathbb{B}}(Y')$ that cancels against the length- n prefix of $\text{eval}_{\mathbb{B}}(Z')$, i.e., we compute the maximal number n such that

$$(\text{eval}_{\mathbb{B}}(Y')[m-n+1, m])^{-1} = a_m^{-1} a_{m-1}^{-1} \cdots a_{m-n+1}^{-1} = b_1 b_2 \cdots b_n = (\text{eval}_{\mathbb{B}}(Z'))[1, n] \quad (1)$$

For a particular n we can check equation (1) in polynomial time by the extension of Theorem 9 to composition systems (which holds by Theorem 4). Now, the maximal n satisfying (1) can be computed in polynomial time as well using binary search. Finally, we add to the composition system \mathbb{B} the production $X' \rightarrow Y'[1, m-n]Z'[n+1, k]$.

A direct corollary of Theorem 26 and 28 is the following result, which solves an open problem from [71].

Corollary 29 ([125]) *Let F be a finitely generated free group. The word problem for $\text{Aut}(F)$ (the automorphism group of F , which is finitely generated) can be solved in polynomial time.*

Several closure properties for the complexity of the compressed word problem were shown:

- (1) If H is a finitely generated subgroup of G and the compressed word problem for G can be solved in polynomial time, then the same holds for H [95].
- (2) If G is a finite-index subgroup of H and the compressed word problem for G can be solved in polynomial time, then the same holds for H [95].
- (3) Let A and B be finite isomorphic subgroups of the group G , let $\varphi : A \rightarrow B$ be an isomorphism, and assume that the compressed word problem for G can be solved in polynomial time. Then, also the compressed word problem for the HNN-extension $H = \langle G, t; t^{-1}at = \varphi(a)(a \in A) \rangle$ can be solved in polynomial time [57].
- (4) Let G and H be groups, which both contain the finite subgroup A . If the compressed word problems of G and H can be solved in polynomial time, then also the compressed word problem for the amalgamated free product $G *_A H$ can be solved in polynomial time [57].
- (5) Let G_1, G_2, \dots, G_n be groups and let I be a symmetric and irreflexive binary relation on the set $\{1, \dots, n\}$. The *graph product* $\text{GP}(G_1, G_2, \dots, G_n, I)$ results from the free product $G_1 * G_2 * \cdots * G_n$ by taking the quotient with respect to the commutation relations $ab = ba$ for all $a \in G_i$ and $b \in G_j$ with $(i, j) \in I$. Assume that the compressed word problems for the groups G_1, \dots, G_n can be solved in polynomial time. Then, also the compressed word problem for the graph product $\text{GP}(G_1, G_2, \dots, G_n, I)$ can be solved in polynomial time [58].

In all these results, polynomial time can be replaced by any complexity class that is closed under polynomial time Turing-reductions.

Results (1)–(5) yield polynomial time algorithms for a quite large class of groups. First of all, (5) implies by taking $G_1 = G_2 = \cdots = G_n = \mathbb{Z}$ that the compressed word problems for right-angled Artin groups (which are also known as graph groups or free partially commutative groups)

can be solved in polynomial time. Due to their rich subgroup structure, right-angled Artin groups received a lot of attention in group theory during the last few years [14, 30, 46]. With (1) and (5) it follows that for every group G that virtually embeds into a right-angled Artin group (i.e., G has a finite index subgroup that embeds into a right-angled Artin group) the compressed word problem can be solved in polynomial time. Groups that virtually embed into a right-angled Artin group are also known as *virtually special*. Recently, it has been shown that virtually special groups form a quite large class of groups. For instance, every Coxeter group is virtually special [56] and every fully residually free group is virtually special [141].⁶ It follows that for every Coxeter group as well as for every fully residually free group the compressed word problem can be solved in polynomial time. For fully residually free groups, this result was directly shown by Macdonald [97] using a generalization of the technique for free groups.

For finitely generated linear groups, the word problem can be solved in logarithmic space [84, 130]. We currently do not know whether also for these groups the compressed word problem can be solved in polynomial time. The best we can achieve is a randomized polynomial time algorithm for the complement of the compressed word problem. More precisely, a language L belongs to the complexity class RP (randomized polynomial time) if there exists a randomized polynomial time algorithm A such that:

- if $x \notin L$ then $\text{Prob}[A \text{ accepts } x] = 0$
- if $x \in L$ then $\text{Prob}[A \text{ accepts } x] \geq 1/2$

The choice of the failure probability $1/2$ in case $x \in L$ is arbitrary: By repeating the algorithm c times (where c is some constant), we can reduce the failure probability to $(1/2)^c$ and still have a randomized polynomial time algorithm. A language L belongs to the class coRP, if the complement of L belongs to RP. This means that there exists a randomized polynomial time algorithm A such that:

- if $x \notin L$ then $\text{Prob}[A \text{ accepts } x] \leq 1/2$
- if $x \in L$ then $\text{Prob}[A \text{ accepts } x] = 1$

Theorem 30 ([95]) *If G is a finitely generated linear group, then the compressed word problem for G belongs to coRP.*

The idea for the proof of Theorem 30 is to reduce (by using results from [84, 130]) the compressed word problem for a finitely generated linear group G to the problem, whether a circuit over a polynomial ring $\mathbb{Z}[x_1, \dots, x_n]$ or $\mathbb{F}_p[x_1, \dots, x_n]$ (for p a prime) evaluates to the zero polynomial. This problem is known as *algebraic identity testing*. It belongs to coRP by [63]. Whether algebraic identity testing belongs to P is a major open problem in complexity theory, see [1] for a survey.

It also makes sense to consider the compressed word problem for a finitely generated monoid M . In that case, the input consists of two SLPs \mathbb{A} and \mathbb{B} over a generating set of M , and it is asked, whether $\text{eval}(\mathbb{A}) = \text{eval}(\mathbb{B})$ in M . For instance, Theorem 9 says that the compressed word problem for a finitely generated free monoid can be solved in polynomial time. In [87], it was shown that the compressed word problem for a finitely generated free inverse monoid of rank $r \geq 2$ is complete for Π_2^p , which is the universal second level of the polynomial time hierarchy. It can be defined as the class of all languages $L \subseteq \Sigma^*$ such that there exists a polynomial time set $P \subseteq \Sigma^* \times \{0, 1\}^* \times \{0, 1\}^*$ and a polynomial $p(x)$ such that for all $u \in \Sigma^*$ we have $(\{0, 1\}^{\leq m})$ denotes the set of binary words of length at most m):

$$u \in L \iff \forall v \in \{0, 1\}^{\leq p(|u|)} \exists w \in \{0, 1\}^{\leq p(|u|)} : (u, v, w) \in P$$

The class Π_2^p contains both NP and coNP.

Recently, power circuits [107], which are a compression scheme for representing huge integers (of multiply exponential size), have been used for solving the word problem for the one-relator Baumslag group and Higman's group [34].

⁶A group G is fully residually free, if for every tuple (g_1, \dots, g_n) of elements from $G \setminus \{1\}$ there exists a free group F and a homomorphism $\varphi : G \rightarrow F$ such that $\varphi(g_i) \neq 1$ for all $1 \leq i \leq n$.

12 Word equations

Let Σ be a finite alphabet and \mathcal{X} be a finite set of variables. A word equation over (Σ, \mathcal{X}) is a pair (u, v) with $u, v \in (\Sigma \cup \mathcal{X})^*$. For a mapping $\sigma : \mathcal{X} \rightarrow \Sigma^*$ let $\widehat{\sigma} : (\Sigma \cup \mathcal{X})^* \rightarrow \Sigma^*$ be the unique homomorphism with $\widehat{\sigma}(a) = a$ for $a \in \Sigma$ and $\widehat{\sigma}(x) = \sigma(x)$. A solution of the word equation (u, v) is a mapping $\sigma : \mathcal{X} \rightarrow \Sigma^*$ such that $\widehat{\sigma}(u) = \widehat{\sigma}(v)$. In the following, we write a word equation (u, v) as $u = v$. A system of word equations over (Σ, \mathcal{X}) is a conjunction $\mathcal{S} = \bigwedge_{i=1}^n u_i = v_i$ of word equations over (Σ, \mathcal{X}) . If σ is a solution of every word equation $u_i = v_i$, then σ is called a solution of \mathcal{S} . Standard arguments show that from a given system of word equations \mathcal{S} one can compute (in logspace) a single word equation $u = v$ such that the set of solutions of $u = v$ is equal to the set of solutions of the system \mathcal{S} , see e.g. [33].

Deciding whether a given word equation has a solution turned out to be very difficult. In the 1950's, Markov conjectured that it is undecidable whether a given word equation has a solution. He knew that solvability of word equations can be reduced to Hilbert's 10th problem, and his goal was to prove the undecidability of Hilbert's 10th problem by proving undecidability of solvability of word equations. Whereas Hilbert's 10th problem was finally shown to be undecidable by the seminal work of Matiyasevich from 1971, solvability of word equations was shown to be decidable by Makanin [98] in 1977. Makanin's algorithm is very complicated, both with respect to its computational complexity and its technical difficulty. Over the years, the estimate on the time and space complexity of Makanin's algorithm was gradually improved. Currently, the best known bound is EXPSPACE [54]. In 1999, Plandowski came up with an alternative algorithm for checking solvability of word equations, which put the problem into PSPACE [113]. This is still the best upper complexity bound. The best known lower bound is NP, and it was repeatedly conjectured that the precise complexity is NP too.

Now, consider the situation, where for every variable $x \in \mathcal{X}$ a length-constraint $\ell(x) \in \mathbb{N}$ is specified in binary representation. Consider the question whether a given word equation $u = v$ has a solution that respects the length constraints. Clearly, this problem is decidable. In fact, the naive algorithm that guesses for every variable $x \in \mathcal{X}$ a word of length $\ell(x)$ over the alphabet Σ and then checks whether the guess is indeed a solution leads to a NEXPTIME-algorithm (recall that we represent the numbers for the length constraints in binary). This trivial complexity bound can drastically be improved by the following result of Plandowski and Rytter [114]:

Theorem 31 ([114]) *The following problem can be solved in polynomial time:*

input: *A system of word equations \mathcal{S} over (Σ, \mathcal{X}) and a mapping $\ell : \mathcal{X} \rightarrow \mathbb{N}$, where every value $\ell(x)$ with $x \in \mathcal{X}$ is binary encoded.*

question: *Is there a solution $\sigma : \mathcal{X} \rightarrow \Sigma^*$ of \mathcal{S} such that $|\sigma(x)| = \ell(x)$ for all $x \in \mathcal{X}$?*

Moreover, in case the answer is affirmative, one can compute in polynomial time for every variable $x \in \mathcal{X}$ an SLP \mathbb{A}_x such that the mapping $x \mapsto \text{eval}(\mathbb{A}_x)$ is a solution of \mathcal{S} and $|\text{eval}(\mathbb{A}_x)| = \ell(x)$ for every $x \in \mathcal{X}$.

A system of word equations \mathcal{S} is *quadratic* if every variable $x \in \mathcal{X}$ occurs at most twice in \mathcal{S} . For quadratic systems of word equations, Diekert and Robson [118] improved the polynomial time upper bound in Theorem 31 to linear time:

Theorem 32 ([118]) *The following problem can be solved in linear time:*

input: *A quadratic system \mathcal{S} of word equations over (Σ, \mathcal{X}) and a mapping $\ell : \mathcal{X} \rightarrow \mathbb{N}$, where every value $\ell(x)$ with $x \in \mathcal{X}$ is binary encoded.*

question: *Is there a solution $\sigma : \mathcal{X} \rightarrow \Sigma^*$ of \mathcal{S} such that $|\sigma(x)| = \ell(x)$ for all $x \in \mathcal{X}$?*

Moreover, in case the answer is affirmative, one can compute in linear time for every variable $x \in \mathcal{X}$ an SLP \mathbb{A}_x such that the mapping $x \mapsto \text{eval}(\mathbb{A}_x)$ is a solution of \mathcal{S} and $|\text{eval}(\mathbb{A}_x)| = \ell(x)$ for every $x \in \mathcal{X}$.

The second statement of Theorem 32 is not stated explicitly in [118] but follows easily from the proof.

Let σ be a solution for a word equation $u = v$. We say that σ is *minimal*, if for every solution σ' of $u = v$ we have $|\sigma(u)| \leq |\sigma'(u)|$. We say that $|\sigma(u)|$ is the length of a minimal solution of $u = v$. By the following result of Plandowski and Rytter [114], minimal solutions for word equations are highly compressible:

Theorem 33 ([114]) *Let $u = v$ be a word equation over (Σ, \mathcal{X}) and let $n = |uv|$. Assume that $u = v$ has a solution and let N be the length of a minimal solution of $u = v$. Then, for every minimal solution σ of $u = v$, the word $\sigma(u)$ can be generated by an SLP of size $O(n^2 \log^2(N)(\log n + \log \log N))$.*

In [68], Jez applied his recompression technique from [67, 66] (see Section 6) to word equations and obtained an alternative PSPACE-algorithm for solving word equations. Moreover, his technique yields a $O(\text{poly}(n) \log N)$ bound in Theorem 33 (instead of the $O(n^2 \log^2(N)(\log n + \log \log N))$ bound); this result is not explicitly stated in [68].

13 Computational topology

In [37, 122, 124] efficient algorithms for SLP-compressed strings are used to solve various problems in computational topology efficiently. Fix a *connected, compact, and orientable surface* M , possibly with a boundary ∂M , see for instance [131] for more details. Typical examples are the 2-dimensional sphere and the torus (both have an empty boundary). A *simple path* in M is the image of a continuous injective mapping f from the unit interval $[0, 1] \subseteq \mathbb{R}$ into M . If moreover $f([0, 1]) \cap \partial M = \{f(0), f(1)\}$, then the simple path is called a *simple arc*. A *simple cycle* in M is the image of a continuous injective mapping f from the unit circle S_1 into M such that $f(S_1) \cap \partial M = \emptyset$. A *simple curve* in M is either a simple arc or a simple cycle. A *simple multi-curve* in M is the union of finitely many pairwise disjoint simple curves. Two simple multi-curves α and β are *isotopic*, if there exists a continuous deformation from α into β that leaves the boundary ∂M fixed. For instance, any two simple cycles on the 2-dimensional sphere are isotopic.

A *triangulation* T of M consists of a finite set V of points in M and a set E of simple paths in M (the edges of the triangulation) such that (i) each endpoint of a path $\alpha \in E$ belongs to V , (ii) if two paths from E intersect, then the intersection is a point of V , and (iii) every connected component of $M \setminus E$ is an open disc that is bounded by three paths from E (a triangle of T). A simple multi-curve α in M is normal w.r.t. T if (i) for every $\beta \in E$ the intersection $\alpha \cap \beta$ consists of finitely many points that are different from the endpoints of β , and (ii) the intersection of α with a triangle of T consists of finitely many paths and for each of these paths the endpoints are located on different sides of the triangle. The latter means that if α enters a triangle at the side β , then it has to leave the triangle at a side different from β .

For each edge e let e^{-1} be a formal inverse of e and choose an orientation \vec{e} of e . A normal simple curve α can be represented by a finite word over the alphabet $E \cup E^{-1}$ as follows: Choose an orientation of α and follow the simple curve in this orientation (if α is a simple cycle then the starting point is chosen arbitrarily). Each time α intersects \vec{e} from left to right (resp. from right to left) we write down e (resp. e^{-1}). The resulting word is called an *intersection sequence* of α . Clearly, in case α is closed this word is not unique, since a cyclic rotation of an intersection sequence is again an intersection sequence. A normal simple multi-curve can be represented by the set of intersection sequences of its connected components. More compact representations of normal simple (multi-)curves are known. First of all, one can use SLPs in order to describe long intersection sequences. Another better known representation are normal coordinates. The vector of normal coordinates of the normal simple multi-curve α is the tuple $(|\alpha \cap \beta|)_{\beta \in E}$ of natural numbers. One obtains a succinct description of the curve by encoding these natural numbers in binary. Two normal simple multi-curves with the same normal coordinates must be isotopic. The number of connected components of α can grow exponentially with the total number of bits of the normal coordinates of α . Hence, there is no chance to compute from the binary encoded normal coordinates of α in polynomial time a set of SLPs for the intersection sequences of the connected components of α ; simply because the number of these SLPs is too large. On the other hand, if

α is connected then such an SLP can be computed by point (1) of Theorem 34 below. Moreover, the number of non-isotopic connected components of α is at most $6t$, where t is the number of triangles of the triangulation T (Lemma 3 in [122]).

Theorem 34 *Let M be a connected, compact, and orientable surface M and let $T = (V, E)$ be a triangulation for M . The following problems can be solved in polynomial time, where the input consists of binary encoded normal coordinates for normal simple multi-curves α and β .*

1. *Compute an SLP for an intersection sequence of α , assuming that α is connected and hence a normal simple curve [124].*
2. *Compute the binary encoded normal coordinates of the connected component of α that contains a given intersection point of α and T [122].*
3. *Compute the number of connected components of α [122].*
4. *Compute a list of binary encoded normal coordinates of connected components $\alpha_1, \dots, \alpha_n$ of α , such that every connected component of α is isotopic to exactly one α_i [122].*
5. *Check, whether α and β are isotopic (assuming that $\partial M \neq \emptyset$) [122].*
6. *Compute the binary encoded normal coordinates of a Dehn twist of α along β assuming that β is a simple cycle (see [131] for the definition of a Dehn-twist) [124].*
7. *Compute the geometric intersection number of α and β [124].⁷*

Note that the converse of point (1) is straightforward: One can compute the binary encoded normal coordinates of a normal simple curve α from an SLP for an intersection sequence of α by simply counting the number of occurrences of each symbol $e \in E \cup E^{-1}$ in the generated intersection sequence.

The main tool for the proof of Theorem 34 are word equations and their connection to SLPs, see Section 12. For instance, for point (1) Schaefer et al. construct a quadratic system of word equations with length constraints such that the solutions of the system are exactly the intersection sequences of α . Then, Theorem 32 can be used in order to construct in linear time an SLP for a solution of the system. This technique was first used by Schaefer et al. in [123] in order to show that string graphs can be recognized in NP. Let us also remark that the algorithm from [124] for computing Dehn twists in polynomial time (see point (6) in Theorem 34) uses the fact that the unique freely reduced normal form for a given SLP-compressed word over $\Sigma \cup \Sigma^{-1}$ can be computed in polynomial time, see the paragraph after Theorem 28.

In [124] also Dehn-Thurston coordinates, which are another succinct representation of simple multi-curves, are discussed. It is shown that Theorem 34 is also true with Dehn-Thurston coordinates instead of normal coordinates.

14 Other applications

Let us briefly discuss some other applications of algorithms for SLPs, in particular Theorem 9 (equality testing for SLP-compressed strings is in P). For precise definitions concerning the discussed problems we refer the interested reader to the original literature.

One of the earliest applications of Theorem 9 concerns normed context-free processes [61, 79]. In [61], the first algorithm for checking bisimulation equivalence of normed context-free processes was developed; in [79] a more efficient algorithm was presented. In both papers, checking bisimulation equivalence of normed context-free processes is reduced to checking equality of SLP-compressed words. In [52], Theorem 9 is applied in the area of program analysis. It is shown that

⁷The geometric intersection number of α and β is the minimum of $|\alpha' \cap \beta'|$, where α' (resp., β') is isotopic to α (resp., β).

the so called interprocedural global value numbering problem for unary uninterpreted function symbols can be solved in polynomial time.

Applications of Theorem 9 for general context-free languages can be found in [13, 136]. In [136] it is shown that one can decide in polynomial time whether a context-free grammar with terminal alphabet $\Sigma \cup \bar{\Sigma}$ (where $\bar{\Sigma} = \{\bar{a} \mid a \in \Sigma\}$ is a disjoint copy of Σ) only generates well-bracketed strings. Here we view \bar{a} as the closing bracket of a , so for instance $ab\bar{b}c\bar{c}a$ is well-bracketed but $a\bar{b}$ is not. In [13], this result is generalized as follows. Assume that R is a terminating and confluent string rewriting system over an alphabet Σ such that the presented monoid Σ^*/R is cancellative. Then by [13] one can decide in polynomial time for a given context-free grammar G whether $L(G)$ is contained in the set of all strings over Σ that can be reduced with R to the empty word.

In [94], Theorem 9 is used in order to solve the isomorphism problem for regular words (a particular class of colored linear orders) in polynomial time. A regular word is given by a set of equations $\{X_i := w_i \mid 1 \leq i \leq n\}$, where every right-hand side w_i is a word over the variables X_i and terminal letters. This is an SLP, where we do not require condition (2) (acyclicity) from Definition 1. Starting from the distinguished variable X_1 , we obtain an infinite derivation tree, where every leaf is labelled with a terminal symbol. The generated regular word is obtained by taking the natural left-to-right order on the leaves of the derivation tree, together with the labels on the leaves. For instance, the recursive equation $X_1 := abX_1$ defines the regular word, where the underlying linear order is (\mathbb{N}, \leq) , and every even (resp., odd) position is labelled with a (resp., b). Another example is $X_1 := X_1aX_1bX_1$. It is not hard to see that this equation defines the regular word consisting of the rational numbers with the natural order, which are densely colored with a and b (between every two different rationals, there is an a -colored point as well as a b -colored point). In [94], it is shown that one can decide for two given sets of equations in polynomial time, whether they define isomorphic regular words. The proof is based on a reduction to equality for SLP-compressed strings.

15 Extensions to trees and pictures

The idea of grammar-based compression can be extended from strings to more complicated structures. In fact, all one needs is a set of operations on a domain D of objects, that construct “larger” D -objects from smaller ones. Using straight-line programs over D with these operations allows to construct “very large” D -objects.

15.1 Trees

For background on trees, tree grammars, and tree automata see [27]. The concept of an SLP can easily be generalized from strings to trees. Here, trees are rooted node-labelled trees. Every node has a label from an alphabet Σ . Moreover, with every symbol $a \in \Sigma$ a natural number (the rank of a) is associated. If a tree node v is labelled with a symbol of rank n , then v has exactly n children, which are linearly ordered. Such trees can conveniently be represented as terms. The size $|t|$ of a tree t is the number of nodes of t . Here is an example:

Example 35 *Let f be a symbol of rank 2, h a symbol of rank 1, and a a symbol of rank 0 (a constant). Then the term $h(f(h(f(h(h(a)), a)), h(f(h(h(a)), a)))$ corresponds to the node-labelled tree of size 14, shown in Figure 4.*

A tree straight-line program (TSLP for short and also called SLCF tree grammar in [90, 92]) over the terminal alphabet Σ (which is a ranked alphabet in the above sense) is a tuple $\mathbb{A} = (N, \Sigma, S, P)$, such that N is a finite set of ranked symbols (the nonterminals), $S \in N$ has rank 0 (the initial nonterminal) and P is a finite set of productions of the form $A(y_1, \dots, y_n) \rightarrow t$ where A is a nonterminal of rank n and t is a tree built up from the ranked symbols in $\Sigma \cup N$ and the parameters y_1, \dots, y_n which are considered as symbols of rank 0 (i.e., constants). Moreover, as for ordinary SLPs, it is required that every nonterminal occurs on the left-hand side of exactly one production, and the relation $\{(A, B) \in N \times N \mid (A(y_1, \dots, y_n) \rightarrow t) \in P, B \text{ occurs in } t\}$ is

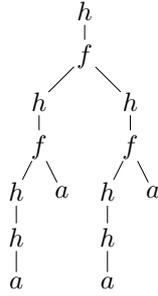


Figure 4: A node-labelled tree

acyclic. A TSLP \mathbb{A} generates a tree $\text{eval}(\mathbb{A})$ in the natural way. During the derivation process, the parameters y_1, \dots, y_n are instantiated with concrete trees. Instead of giving a formal definition, let us consider an example.

Example 36 Let S, A, B, C be nonterminals, let S be the start nonterminal and let the TSLP \mathbb{A} consist of the following productions:

$$\begin{aligned}
 S &\rightarrow A(B(a), B(a)) \\
 A(x_1, x_2) &\rightarrow C(C(x_1, a), C(x_2, a)) \\
 C(x_1, x_2) &\rightarrow h(f(x_1, x_2)) \\
 B(x) &\rightarrow h(h(x))
 \end{aligned}$$

Then $\text{eval}(G)$ is the tree from Example 35. It can be derived as follows:

$$\begin{aligned}
 S &\rightarrow A(B(a), B(a)) \\
 &\rightarrow C(C(B(a), a), C(B(a), a)) \\
 &\rightarrow C(C(h(h(a)), a), C(h(h(a)), a)) \\
 &\rightarrow h(f(C(h(h(a)), a), C(h(h(a)), a))) \\
 &\rightarrow h(f(h(f(h(h(a)), a)), h(f(h(h(a)), a))))
 \end{aligned}$$

The size of a TSLP $\mathbb{A} = (N, \Sigma, S, P)$ is defined as the sum of the sizes of all right-hand sides of P . Due to multiple occurrences of parameters in right-hand sides, TSLPs are able to generate trees of doubly exponential size.

Example 37 Let the TSLP \mathbb{A}_n consist of the following productions:

$$\begin{aligned}
 S &\rightarrow A_0(a) \\
 A_i(y_1) &\rightarrow A_{i+1}(A_{i+1}(y_1)) \quad \text{for } 0 \leq i < n \\
 A_n(y_1) &\rightarrow f(y_1, y_1)
 \end{aligned}$$

Then $\text{eval}(\mathbb{A}_n)$ is a complete binary tree of height 2^n . Thus, $|\text{eval}(\mathbb{A}_n)| = 2 \cdot 2^{2^n} - 1$.

Example 37 motivates the definition of linear TSLPs. A TSLP $\mathbb{A} = (N, S, P)$ is linear if for every rule $(A(y_1, \dots, y_n) \rightarrow t) \in P$, each of the parameters y_1, \dots, y_n occurs exactly once in the tree t . The TSLP from Example 36 is linear. In contrast to non-linear TSLPs, for a linear TSLP \mathbb{A} the size of tree $\text{eval}(\mathbb{A})$ is bounded by $2^{O(|\mathbb{A}|)}$. Efficient algorithms that generate for a given input tree t a linear TSLP \mathbb{A} with $\text{eval}(\mathbb{A}) = t$ are described in [18, 91] and [2] (for a slightly different type of grammars). A TSLP is nothing else than a context-free tree grammar that generates a single tree. An extension of TSLPs to higher order tree grammars was recently proposed in [76].

Let us now consider algorithmic problems for TSLP-compressed trees. Concerning equality checking, we have the following:

Theorem 38 ([18, 126]) *For two given TSLPs \mathbb{A} and \mathbb{B} it can be checked in PSPACE, whether $\text{eval}(\mathbb{A}) = \text{eval}(\mathbb{B})$. For two given linear TSLPs \mathbb{A} and \mathbb{B} it can be checked in P, whether $\text{eval}(\mathbb{A}) = \text{eval}(\mathbb{B})$.*

For the PSPACE-bound in the first statement, one has to notice that a TSLP \mathbb{A} can also be viewed as a graph generating grammar that produces a node-labelled directed acyclic graph $\text{dag}(\mathbb{A})$, whose unfolding is $\text{eval}(\mathbb{A})$. This graph grammar is obtained by merging in a right-hand side $t(y_1, \dots, y_n)$ of a production $A(y_1, \dots, y_n) \rightarrow t$ all occurrences of a variable y_i into a single y_i -labelled node. The number of nodes of $\text{dag}(\mathbb{A})$ is bounded singly exponential in the size of \mathbb{A} , hence a node of $\text{dag}(\mathbb{A})$ can be represented in polynomial space. A path in $\text{dag}(\mathbb{A})$ that starts in the root can be represented by the sequence of natural number n_1, \dots, n_k , such that in the i^{th} step the path moves from the current node to its n_i^{th} child node. A PSPACE-algorithm for checking $\text{eval}(\mathbb{A}) \neq \text{eval}(\mathbb{B})$ (this suffices, since PSPACE is closed under complement) simply guesses such a path p step by step and thereby stores the nodes of $\text{dag}(\mathbb{A})$ and $\text{dag}(\mathbb{B})$, respectively, that are reached by the path p . The algorithm rejects as soon as these two nodes are labelled by different symbols. It remains open, whether equality of TSLP-compressed trees can be checked more efficiently.

For the second statement of Theorem 38 one constructs in polynomial time from a linear TSLP \mathbb{A} a (string generating) SLP \mathbb{A}' such that $\text{eval}(\mathbb{A}')$ represents a depth-first left-to-right transversal of the tree $\text{eval}(\mathbb{A})$. For two TSLPs \mathbb{A} and \mathbb{B} we have $\text{eval}(\mathbb{A}) = \text{eval}(\mathbb{B})$ if and only if $\text{eval}(\mathbb{A}') = \text{eval}(\mathbb{B}')$. Hence, equality of trees that are represented by linear TSLPs can be reduced to checking equality of SLP-compressed strings, which can be checked in polynomial time by Theorem 9. In [126], the second statement of Theorem 38 is shown by a direct extension of Plandowski's algorithm for string SLPs.

In [90, 92], the problem of evaluating tree automata over TSLP-compressed input trees is considered. A tree automaton runs on an node-labelled input tree bottom-up and thereby assigns states to tree nodes. Transitions are of the form (q_1, \dots, q_n, f, q) , where f is a symbol (node label) of rank n and q_1, \dots, q_n, q are states of the tree automaton. Then a *run* of the tree automaton is a mapping ρ from the tree nodes to states that is consistent with the set of transitions in the following sense: If a tree node is labelled with the symbol f (of rank n) and v_1, \dots, v_n are the children of v in that order, then $(\rho(v_1), \dots, \rho(v_n), f, \rho(v))$ must be a transition of the tree automaton. A tree automaton accepts a tree if there is a run that assigns a final state to the root of the tree (every tree automaton has a distinguished set of final states).

Theorem 39 *The following complexity results hold:*

- *It is PSPACE-complete to check for a given TSLP \mathbb{A} and a given tree automaton A , whether A accepts $\text{eval}(\mathbb{A})$ [90].*
- *It is P-complete to check for a given linear TSLP \mathbb{A} and a given tree automaton A , whether A accepts $\text{eval}(\mathbb{A})$ [92].*

The polynomial time algorithm in the linear case works in two steps:

- (1) Transform in polynomial time the input TSLP \mathbb{A} into a TSLP \mathbb{B} such that $\text{eval}(\mathbb{A}) = \text{eval}(\mathbb{B})$ and every nonterminal of \mathbb{B} has rank at most one. This is the difficult step in [92].
- (2) Evaluate the input tree automaton A on $\text{eval}(\mathbb{B})$ in polynomial time, using the fact that every nonterminal of \mathbb{B} has rank at most one, see [90].

Several other algorithmic problems for TSLP-compressed input trees are studied in [20, 41, 80, 127, 128].

15.2 Pictures

A picture over the alphabet Σ is a mapping $p : \{1, \dots, n\} \times \{1, \dots, m\} \rightarrow \Sigma$ for some $n, m \geq 1$. We say that (n, m) is the dimension of the picture p . For more information on pictures see [47].

A picture can be viewed as a 2-dimensional word. We define two partial concatenation operations for pictures. Given two pictures p_1 and p_2 , where p_1 has dimension (n, m) and p_2 has dimension (n, k) , we define the horizontal concatenation $(p_1, p_2) : \{1, \dots, n\} \times \{1, \dots, m+k\} \rightarrow \Sigma$ (a picture of dimension $(n, m+k)$) as follows, where $x \in \{1, \dots, n\}$ and $y \in \{1, \dots, m+k\}$:

$$(p_1, p_2)(x, y) = \begin{cases} p_1(x, y) & \text{if } y \leq m \\ p_2(x, y - m) & \text{if } y > m \end{cases}$$

Given two pictures p_1 and p_2 , where p_1 has dimension (m, n) and p_2 has dimension (k, n) , we define the vertical concatenation $\begin{pmatrix} p_2 \\ p_1 \end{pmatrix} : \{1, \dots, m+k\} \times \{1, \dots, n\} \rightarrow \Sigma$ (a picture of dimension $(m+k, n)$) as follows, where $x \in \{1, \dots, m+k\}$ and $y \in \{1, \dots, n\}$:

$$\begin{pmatrix} p_2 \\ p_1 \end{pmatrix}(x, y) = \begin{cases} p_1(x, y) & \text{if } x \leq m \\ p_2(x - m, y) & \text{if } x > m \end{cases}$$

It is now straightforward to define 2-dimensional straight-line programs (briefly 2SLPs). The productions of a 2SLP \mathbb{A} have one of the forms

$$X \rightarrow a, \quad X \rightarrow (Y, Z), \quad X \rightarrow \begin{pmatrix} Z \\ Y \end{pmatrix},$$

where X, Y, Z are nonterminals and a is a terminal symbol. In order to generate a picture, we have to require that the dimensions match. Define the dimension of a nonterminal X inductively as follows. If the unique production for X has the form $X \rightarrow a$, then the dimension of X is $(1, 1)$. If the unique production for X has the form $X \rightarrow (Y, Z)$, then, inductively the dimension (n, m) (resp., (k, ℓ)) of Y (resp., Z) is already defined. We require that $n = k$ and define the dimension of X as $(n, m + \ell)$. The constraint for a production $X \rightarrow \begin{pmatrix} Z \\ Y \end{pmatrix}$ is analogous. If the dimension of the initial nonterminal S is (n, m) , then the 2SLP \mathbb{A} defines a picture $\text{eval}(\mathbb{A})$ of dimension (n, m) .

Algorithmic problems for 2SLPs were studied in [11]. For equality checking, the following was shown; recall the definition of the complexity class coRP from Section 11.

Theorem 40 ([11]) *The problem of checking $\text{eval}(\mathbb{A}) = \text{eval}(\mathbb{B})$ for two given 2SLPs belongs to the complexity class coRP .*

For the proof, one encodes a picture p of dimension (m, n) over the alphabet $\{1, \dots, k\}$ as the polynomial

$$\text{poly}_p(x, y) = \sum_{1 \leq i \leq m} \sum_{1 \leq j \leq n} p(i, j) x^{i-1} y^{j-1}.$$

For two 2SLPs \mathbb{A} and \mathbb{B} , we have $\text{eval}(\mathbb{A}) = \text{eval}(\mathbb{B})$ if and only if $\text{poly}_{\text{eval}(\mathbb{A})}(x, y) - \text{poly}_{\text{eval}(\mathbb{B})}(x, y)$ is the zero polynomial. Now, from a given 2SLP \mathbb{A} one can construct in polynomial time a circuit over the polynomial ring $\mathbb{Z}[x, y]$ that defines the polynomial $\text{poly}_{\text{eval}(\mathbb{A})}(x, y)$. Hence, the initial question whether $\text{eval}(\mathbb{A}) = \text{eval}(\mathbb{B})$ can be reduced to the problem whether a circuit over the polynomial ring $\mathbb{Z}[x, y]$ evaluates to the zero polynomial. This is an instance of the algebraic identity testing that we already encountered after Theorem 30 and that belongs to coRP . It remains open, whether equality of 2SLP-compressed pictures can be checked in polynomial time.

Also pattern matching problems can naturally be defined for pictures. A picture p of dimension (k, ℓ) occurs in the picture q of dimension (m, n) if there exist $1 \leq i \leq m$ and $1 \leq j \leq n$ such that (i) $i + k - 1 \leq m$, (ii) $j + \ell - 1 \leq n$, and (iii) for all $0 \leq x \leq k - 1$ and $0 \leq y \leq \ell - 1$ we have $q(i + x, j + y) = p(x + 1, y + 1)$. The following two results from [11] pinpoint the complexity of 2-dimensional pattern matching:

Theorem 41 ([11]) *It is coNP -complete to check for an explicitly given picture p and a 2SLP \mathbb{A} , whether p occurs in $\text{eval}(\mathbb{A})$.*

The class Σ_2^p consists of all complements of Π_2^p -sets, see the end of Section 11 for the definition of Π_2^p . Hence, $L \subseteq \Sigma^*$ belongs to Σ_2^p if there exists a polynomial time set $P \subseteq \Sigma^* \times \{0, 1\}^* \times \{0, 1\}^*$ and a polynomial $p(x)$ such that for all $u \in \Sigma^*$ we have:

$$u \in L \iff \exists v \in \{0, 1\}^{\leq p(|u|)} \forall w \in \{0, 1\}^{\leq p(|u|)} : (u, v, w) \in P \quad (2)$$

Theorem 42 ([11]) *It is Π_2^p -complete to check for two given 2SLP \mathbb{A} and \mathbb{B} , whether $\text{eval}(\mathbb{A})$ occurs in $\text{eval}(\mathbb{B})$.*

The upper complexity bounds in Theorem 41 and 42 are straightforward. For instance, for two given 2SLPs \mathbb{A} and \mathbb{B} the definition of “ $\text{eval}(\mathbb{A})$ occurs in $\text{eval}(\mathbb{B})$ ” directly leads to a formula of the general form (2). The lower complexity bound in Theorem 41 is shown by a reduction from 3SAT. Assume that $F(x_1, \dots, x_n) = \bigwedge_{i=1}^n C_i$ is a boolean formula, where every C_i is a conjunction of three literals (a literal is a boolean variable or a negated boolean variable) and x_1, \dots, x_m are the variables that occur in F . Hence there are 2^m truth assignments. For a given clause $C_i = (A \vee B \vee C)$, where A , B , and C are literals, one can construct an SLP \mathbb{A}_i over the alphabet $\{0, 1\}$ such that $|\text{eval}(\mathbb{A}_i)| = 2^m$ and for all $1 \leq j \leq 2^m$, $\text{eval}(\mathbb{A}_i)[j]$ is the truth value of C_i under the j^{th} truth assignment in lexicographic order. Now, one can easily construct a 2SLP \mathbb{A} such that $\text{eval}(\mathbb{A})$ is a picture of dimension $(n, 2^m)$ such that the i^{th} row of $\text{eval}(\mathbb{A})$ is the word $\text{eval}(\mathbb{A}_i)$ for every $1 \leq i \leq n$. Hence, the formula F is satisfiable if and only if the picture $\text{eval}(\mathbb{A})$ contains a column only consisting of 1’s. The lower bound in Theorem 42 is shown by a reduction from a Π_2^p -complete variant of the subsetsum problem.

16 Open problems

Let us state some open problems that in the opinion of the author deserve further investigations.

- Is the problem of checking $\text{eval}(\mathbb{A}) = \text{eval}(\mathbb{B})$ for two given SLPs \mathbb{A} and \mathbb{B} P-complete, or does this problem belong to the parallel complexity class NC (see Section 5)? The latter would mean that the problem can be solved in polylogarithmic time using polynomially many processors. The same question can be considered also for the fully compressed pattern matching problem (see Section 6).
- What is the precise complexity of the fully compressed subsequence matching problem (see Section 8)? The author conjectures that this problem is PSPACE-complete.
- Is the compressed word problem (see Section 11) for a finitely generated linear group solvable in polynomial time? What about subclasses of linear groups, e.g., braid groups and polycyclic groups.
- The *compressed subgroup membership problem* for a free group $F(\Sigma)$ is the following problem: Given SLPs $\mathbb{A}, \mathbb{B}_1, \dots, \mathbb{B}_n$ (over $\Sigma \cup \Sigma^{-1}$), does the word $\text{eval}(\mathbb{A})$ represent a group element from the subgroup of $F(\Sigma)$ generated by $\text{eval}(\mathbb{B}_1), \dots, \text{eval}(\mathbb{B}_n)$? We are only aware of an exponential time algorithm for this problem.
- Can equality of 2SLP-compressed pictures (see Section 15.2) be checked in polynomial time?

Acknowledgment

The author is very grateful to Pawel Gawrychowski, Artur Jez, Sebastian Maneth, Wojciech Rytter, Marcus Schäfer, and Manfred Schmidt-Schauß for many valuable comments and reading a preliminary version of this paper.

References

- [1] M. Agrawal and R. Saptharishi. Classifying polynomials and identity testing. *Current Trends in Science – Platinum Jubilee Special*, pages 149–162, 2009.
- [2] T. Akutsu. A bisection algorithm for grammar-based compression of ordered trees. *Inf. Process. Lett.*, 110(18-19):815–820, 2010.
- [3] E. Allender, P. Bürgisser, J. Kjeldgaard-Pedersen, and P. B. Miltersen. On the complexity of numerical analysis. *SIAM J. Comput.*, 38(5):1987–2006, 2009.
- [4] S. Alstrup, G. S. Brodal, and T. Rauhe. Dynamic pattern matching. Technical Report DIKU Report 98/27, Department of Computer Science, University of Copenhagen, 1998.
- [5] S. Alstrup, G. S. Brodal, and T. Rauhe. Pattern matching in dynamic texts. In *Proceedings of the Eleventh Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2000*, pages 819–828. ACM/SIAM, 2000.
- [6] R. Alur and P. Madhusudan. Visibly pushdown languages. In *Proceedings of the 36th Annual ACM Symposium on Theory of Computing, STOC 2004*, pages 202–211. ACM Press, 2004.
- [7] A. Amir, G. Benson, and M. Farach. Let sleeping files lie: Pattern matching in Z-compressed files. *Journal of Computer and System Sciences*, 52(2):299–307, 1996.
- [8] J. Arpe and R. Reischuk. On the complexity of optimal grammar-based compression. In *Proceedings of the Data Compression Conference, DCC 2006*, pages 173–182. IEEE Computer Society, 2006.
- [9] D. A. M. Barrington. Bounded-width polynomial-size branching programs recognize exactly those languages in NC^1 . *Journal of Computer and System Sciences*, 38:150–164, 1989.
- [10] M. Beaudry, P. McKenzie, P. Péladeau, and D. Thérien. Finite monoids: From word to circuit evaluation. *SIAM Journal on Computing*, 26(1):138–152, 1997.
- [11] P. Berman, M. Karpinski, L. L. Larmore, W. Plandowski, and W. Rytter. On the complexity of pattern matching for highly compressed two-dimensional texts. *Journal of Computer and System Sciences*, 65(2):332–350, 2002.
- [12] A. Bertoni, C. Choffrut, and R. Radicioni. Literal shuffle of compressed words. In *Proceeding of the 5th IFIP International Conference on Theoretical Computer Science, IFIP TCS 2008*, pages 87–100. Springer, 2008.
- [13] A. Bertoni, C. Choffrut, and R. Radicioni. The inclusion problem of context-free languages: Some tractable cases. *Int. J. Found. Comput. Sci.*, 22(2):289–299, 2011.
- [14] M. Bestvina and N. Brady. Morse theory and finiteness properties of groups. *Inventiones Mathematicae*, 129(3):445–470, 1997.
- [15] P. Bille, G. M. Landau, R. Raman, K. Sadakane, S. R. Satti, and O. Weimann. Random access to grammar-compressed strings. In *Proceedings of the 22nd Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2011*, pages 373–389. SIAM, 2011.
- [16] W. W. Boone. The word problem. *Annals of Mathematics (2)*, 70:207–265, 1959.
- [17] I. Burmistrov and L. Khvorost. Straight-line programs: A practical test. In *Proceedings of the 1st International Conference on Data Compression, Communications and Processing, CCP 2011*, pages 76–81. IEEE, 2011.
- [18] G. Busatto, M. Lohrey, and S. Maneth. Efficient memory representation of XML document trees. *Information Systems*, 33(4-5):456–474, 2008.

- [19] F. B. Cannonito and R. W. Gatterdam. The word problem and power problem in 1-relator groups are primitive recursive. *Pacific Journal of Mathematics*, 61(2):351–359, 1975.
- [20] A. G. Carles Creus and G. Godoy. One-context unification with STG-compressed terms is in NP. In *Proceedings of the 23rd International Conference on Rewriting Techniques and Applications, RTA 2012*, LIPIcs. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2012.
- [21] H. Caussinus, P. McKenzie, D. Thérien, and H. Vollmer. Nondeterministic NC^1 computation. *Journal of Computer and System Sciences*, 57(2):200–212, 1998.
- [22] P. Cégielski, I. Guessarian, Y. Lifshits, and Y. Matiyasevich. Window subsequence problems for compressed texts. In *Proceedings of the International Computer Science Symposium in Russia, CSR 2006*, number 3967 in Lecture Notes in Computer Science, pages 127–136. Springer, 2006.
- [23] M. Charikar, E. Lehman, A. Lehman, D. Liu, R. Panigrahy, M. Prabhakaran, A. Sahai, and A. Shelat. The smallest grammar problem. *IEEE Transactions on Information Theory*, 51(7):2554–2576, 2005.
- [24] S. Chen, E. Verbin, and W. Yu. Data structure lower bounds on random access to grammar-compressed strings. Technical report, arXiv.org, 2012. <http://arxiv.org/abs/1203.1080>.
- [25] F. Claude and G. Navarro. Improved grammar-based compressed indexes. Technical report, arXiv.org, 2011. <http://arxiv.org/abs/1110.4493>.
- [26] F. Claude and G. Navarro. Self-indexed grammar-based compression. *Fundam. Inform.*, 111(3):313–337, 2011.
- [27] H. Comon, M. Dauchet, R. Gilleron, F. Jacquemard, D. Lugiez, C. Löding, S. Tison, and M. Tommasi. Tree automata techniques and applications. Available on: <http://tata.gforge.inria.fr/>, 2007.
- [28] K. J. Compton and C. W. Henson. A uniform method for proving lower bounds on the computational complexity of logical theories. *Annals of Pure and Applied Logic*, 48:1–79, 1990.
- [29] M. M. F. T. Coppersmith-Winograd. Virginia vassilevska williams. In *Proceedings of the 44th ACM Symposium on Theory of Computing, STOC 2012*. ACM, 2012.
- [30] J. Crisp and B. Wiest. Embeddings of graph braid and surface groups in right-angled artin groups and braid groups. *Algebraic & Geometric Topology*, 4:439–472, 2004.
- [31] M. Crochemore, G. M. Landau, and M. Ziv-Ukelson. A subquadratic sequence alignment algorithm for unrestricted scoring matrices. *SIAM J. Comput.*, 32(6):1654–1673, 2003.
- [32] M. Dehn. Über die Topologie des dreidimensionalen Raumes. *Mathematische Annalen*, 69:137–168, 1910. In German.
- [33] V. Diekert. Makanin’s algorithm. In *Algebraic Combinatorics on Words*, pages 342–390. Cambridge University Press, 2001.
- [34] V. Diekert, J. Laun, and A. Ushakov. Efficient algorithms for highly compressed data: The word problem in Higman’s group is in P. In *Proceedings of the 29th International Symposium on Theoretical Aspects of Computer Science, STACS 2012*, number 14 of LIPIcs, pages 218–229. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2012.
- [35] P. W. Dymond. Input-driven languages are in $\log n$ depth. *Information Processing Letters*, 26(5):247–250, 1988.

- [36] D. B. A. Epstein, J. W. Cannon, D. F. Holt, S. V. F. Levy, M. S. Paterson, and W. P. Thurston. *Word processing in groups*. Jones and Bartlett, Boston, 1992.
- [37] J. Erickson and A. Nayyeri. Tracing compressed curves in triangulated surfaces. In *Proceedings of the 28th Symposium on Computational Geometry, SOCG 2012*, 2012.
- [38] M. Farach and M. Thorup. String matching in Lempel-Ziv compressed strings. *Algorithmica*, 20(4):388–404, 1998.
- [39] N. J. Fine and H. S. Wilf. Uniqueness theorems for periodic functions. *Proceedings of the American Mathematical Society*, 16:109–114, 1965.
- [40] T. Gagie, P. Gawrychowski, J. Kärkkäinen, Y. Nekrich, and S. J. Puglisi. A faster grammar-based self-index. In *Proceedings of the 6th International Conference on Language and Automata Theory and Applications, LATA 2012*, number 7183 of *Lecture Notes in Computer Science*, pages 240–251. Springer, 2012.
- [41] A. Gascón, G. Godoy, and M. Schmidt-Schauß. Unification and matching on compressed terms. *ACM Trans. Comput. Log.*, 12(4):26, 2011.
- [42] L. Gasieniec, M. Karpinski, W. Plandowski, and W. Rytter. Efficient algorithms for Lempel-Ziv encoding (extended abstract). In *Proceedings of the 5th Scandinavian Workshop on Algorithm Theory, SWAT 1996*, number 1097 in *Lecture Notes in Computer Science*, pages 392–403. Springer, 1996.
- [43] L. Gasieniec, M. Karpinski, W. Plandowski, and W. Rytter. Randomized efficient algorithms for compressed strings: The finger-print approach (extended abstract). In *Proceedings of the 7th Annual Symposium on Combinatorial Pattern Matching, CPM 1996*, number 1075 of *Lecture Notes in Computer Science*, pages 39–49. Springer, 1996.
- [44] P. Gawrychowski. Optimal pattern matching in LZW compressed strings. In *Proceedings of the 22nd Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2011*, pages 362–372. SIAM, 2011.
- [45] P. Gawrychowski. Pattern matching in Lempel-Ziv compressed strings: Fast, simple, and deterministic. In *Proceedings of the 19th Annual European Symposium on Algorithms, ESA 2011*, number 6942 of *Lecture Notes in Computer Science*, pages 421–432. Springer, 2011.
- [46] R. Ghrist and V. Peterson. The geometry and topology of reconfiguration. *Advances in Applied Mathematics*, 38(3):302–323, 2007.
- [47] D. Giammarresi and A. Restivo. Two-dimensional languages. In G. Rozenberg and A. Salomaa, editors, *Handbook of Formal Languages*, volume 3, pages 216–267. Springer, 1997.
- [48] L. M. Goldschlager. ε -productions in context-free grammars. *Acta Informatica*, 16:303–308, 1981.
- [49] K. Goto, H. Bannai, S. Inenaga, and M. Takeda. Fast q-gram mining on SLP compressed strings. In *Proceedings of the 18th International Symposium on String Processing and Information Retrieval, SPIRE 2011*, number 7024 of *Lecture Notes in Computer Science*, pages 278–289. Springer, 2011.
- [50] K. Goto, H. Bannai, S. Inenaga, and M. Takeda. Speeding-up q-gram mining on grammar-based compressed texts. Technical report, arXiv.org, 2012. <http://arxiv.org/abs/1202.3311>.
- [51] S. Greibach. The hardest context-free language. *SIAM Journal on Computing*, 2(4):304–310, 1973.

- [52] S. Gulwani and A. Tiwari. Computing procedure summaries for interprocedural analysis. In *Proceedings of the 16th European Symposium on Programming, ESOP 2007*, number 4421 of *Lecture Notes in Computer Science*, pages 253–267. Springer, 2007.
- [53] D. Gusfield. *Algorithms on Strings, Trees, and Sequences*. Cambridge University Press, 1999.
- [54] C. Gutiérrez. Satisfiability of word equations with constants is in exponential space. In *Proceedings of the 39th Annual Symposium on Foundations of Computer Science, FOCS 1998*, pages 112–119. IEEE Computer Society Press, 1998.
- [55] C. Hagenah. *Gleichungen mit regulären Randbedingungen über freien Gruppen*. PhD thesis, University of Stuttgart, Institut für Informatik, 2000.
- [56] F. Haglund and D. T. Wise. Coxeter groups are virtually special. *Advances in Mathematics*, 224(5):1890–1903, 2010.
- [57] N. Haubold and M. Lohrey. Compressed word problems in HNN-extensions and amalgamated products. *Theory of Computing Systems*, 49(2):283–305, 2011.
- [58] N. Haubold, M. Lohrey, and C. Mathissen. Compressed conjugacy and the word problem for outer automorphism groups of graph groups. In *Proceedings of Developments in Language Theory, DLT 2010*, number 6224 in *Lecture Notes in Computer Science*, pages 218–230. Springer, 2010. long version appears in *International Journal of Algebra and Computation*.
- [59] D. Hermelin, G. M. Landau, S. Landau, and O. Weimann. Unified compression-based acceleration of edit-distance computation. Technical report, arXiv.org, 2010. <http://arxiv.org/abs/1004.1194>.
- [60] U. Hertrampf, C. Lautemann, T. Schwentick, H. Vollmer, and K. W. Wagner. On the power of polynomial time bit-reductions. In *Proceedings of the 8th Annual Structure in Complexity Theory Conference*, pages 200–207. IEEE Computer Society Press, 1993.
- [61] Y. Hirshfeld, M. Jerrum, and F. Moller. A polynomial algorithm for deciding bisimilarity of normed context-free processes. *Theoretical Computer Science*, 158(1&2):143–159, 1996.
- [62] J. E. Hopcroft and J. D. Ullman. *Introduction to automata theory, languages and computation*. Addison–Wesley, Reading, MA, 1979.
- [63] O. H. Ibarra and S. Moran. Probabilistic algorithms for deciding equivalence of straight-line programs. *Journal of the ACM*, 30(1):217–228, 1983.
- [64] H. B. H. III, D. J. Rosenkrantz, and T. G. Szymanski. On the equivalence, containment, and covering problems for the regular and context-free languages. *J. Comput. Syst. Sci.*, 12(2):222–268, 1976.
- [65] B. Jenner, P. McKenzie, and D. Thérien. Logspace and logtime leaf languages. *Information and Computation*, 129(1):21–33, 1996.
- [66] A. Jez. Compressed membership for NFA (DFA) with compressed labels is in NP (P). In *Proceedings of the 29th International Symposium on Theoretical Aspects of Computer Science, STACS 2012*, number 14 of *LIPICs*, pages 136–147. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2012.
- [67] A. Jez. Faster fully compressed pattern matching by recompression. In *Proceedings of the 39th International Colloquium on Automata, Languages, and Programming, ICALP 2012*, number 7391 of *Lecture Notes in Computer Science*, pages 533–544. Springer, 2012.
- [68] A. Jez. Recompression: a simple and powerful technique for word equations. Technical report, arXiv.org, 2012. <http://arxiv.org/abs/1203.3705>.

- [69] T. Jiang and B. Ravikumar. A note on the space complexity of some decision problems for finite automata. *Information Processing Letters*, 40(1):25–31, 1991.
- [70] E. Kaltofen. Greatest common divisors of polynomials given by straight-line programs. *Journal of the ACM*, 35(1):231–264, 1988.
- [71] I. Kapovich, A. Miasnikov, P. Schupp, and V. Shpilrain. Generic-case complexity, decision problems in group theory, and random walks. *Journal of Algebra*, 264(2):665–694, 2003.
- [72] H. J. Karloff and W. L. Ruzzo. The iterated mod problem. *Information and Computation*, 80(3):193–204, 1989.
- [73] M. Karpinski, W. Rytter, and A. Shinohara. Pattern-matching for strings with short descriptions. In *Proceedings of the 6th Annual Symposium on Combinatorial Pattern Matching, CPM 1995*, number 937 in Lecture Notes in Computer Science, pages 205–214. Springer, 1995.
- [74] T. Kida, T. Matsumoto, Y. Shibata, M. Takeda, A. Shinohara, and S. Arikawa. Collage system: a unifying framework for compressed pattern matching. *Theor. Comput. Sci.*, 1(298):253–272, 2003.
- [75] D. E. Knuth, J. H. Morris, Jr., and V. R. Pratt. Fast pattern matching in strings. *SIAM Journal on Computing*, 6(2):323–350, 1977.
- [76] N. Kobayashi, K. Matsuda, and A. Shinohara. Functional programs as compressed data. In *Proceedings of the ACM SIGPLAN 2012 Workshop on Partial Evaluation and Program Manipulation, PEPM 2012*, pages 121–130. ACM, 2012.
- [77] D. Kozen. Lower bounds for natural proof systems. In *Proceedings of the 18th Annual Symposium on Foundations of Computer Science, FOCS 1977*, pages 254–266. IEEE Computer Society Press, 1977.
- [78] N. J. Larsson and A. Moffat. Offline dictionary-based compression. In *Proceedings of the Data Compression Conference, DCC 1999*, pages 296–305. IEEE Computer Society Press, 1999.
- [79] S. Lasota and W. Rytter. Faster algorithm for bisimulation equivalence of normed context-free processes. In *Proceedings of the 31th International Symposium on Mathematical Foundations of Computer Science, MFCS 2006*, number 4162 in Lecture Notes in Computer Science, pages 646–657. Springer, 2006.
- [80] J. Levy, M. Schmidt-Schauß, and M. Villaret. The complexity of monadic second-order unification. *SIAM J. Comput.*, 38(3):1113–1140, 2008.
- [81] Y. Lifshits. Processing compressed texts: A tractability border. In B. Ma and K. Zhang, editors, *Proceedings of the 18th Annual Symposium on Combinatorial Pattern Matching, CPM 2007*, number 4580 in Lecture Notes in Computer Science, pages 228–240. Springer, 2007.
- [82] Y. Lifshits and M. Lohrey. Querying and embedding compressed texts. In *Proceedings of the 31th International Symposium on Mathematical Foundations of Computer Science, MFCS 2006*, number 4162 in Lecture Notes in Computer Science, pages 681–692. Springer, 2006.
- [83] Y. Lifshits, S. Mozes, O. Weimann, and M. Ziv-Ukelson. Speeding up HMM decoding and training by exploiting sequence repetitions. *Algorithmica*, 54(3):379–399, 2009.
- [84] R. J. Lipton and Y. Zalcstein. Word problems solvable in logspace. *Journal of the ACM*, 24(3):522–526, 1977.

- [85] M. Lohrey. Word problems and membership problems on compressed words. *SIAM Journal on Computing*, 35(5):1210 – 1240, 2006.
- [86] M. Lohrey. Compressed membership problems for regular expressions and hierarchical automata. *International Journal of Foundations of Computer Science*, 21(5):817–841, 2010.
- [87] M. Lohrey. Compressed word problems for inverse monoids. In *Proceedings of the 36th International Symposium on Mathematical Foundations of Computer Science, MFCS 2011*, number 6907 of *Lecture Notes in Computer Science*, pages 448–459. Springer, 2011.
- [88] M. Lohrey. Leaf languages and string compression. *Information and Computation*, 209(6):951–965, 2011.
- [89] M. Lohrey. Model-checking hierarchical structures. *Journal of Computer and System Sciences*, 78(2):461–490, 2012.
- [90] M. Lohrey and S. Maneth. The complexity of tree automata and XPath on grammar-compressed trees. *Theoretical Computer Science*, 363(2):196–210, 2006.
- [91] M. Lohrey, S. Maneth, and R. Mennicke. Tree structure compression with repair. In *Proceedings of the Data Compression Conference, DCC 2011*, pages 353–362. IEEE Computer Society, 2011.
- [92] M. Lohrey, S. Maneth, and M. Schmidt-Schauß. Parameter reduction and automata evaluation for grammar-compressed trees. *Journal of Computer and Systems Sciences*, 2012. to appear.
- [93] M. Lohrey and C. Mathissen. Compressed membership in automata with compressed labels. In *Proceedings of the 6th International Computer Science Symposium in Russia, CSR 2011*, number 6651 in *Lecture Notes in Computer Science*, pages 275–288. Springer, 2011.
- [94] M. Lohrey and C. Mathissen. Isomorphism of regular trees and words. In *Proceeding of the 38th International Colloquium on , Automata, Languages and Programming, ICALP 2011*, number 6756 in *Lecture Notes in Computer Science*, pages 210–221. Springer, 2011.
- [95] M. Lohrey and S. Schleimer. Efficient computation in groups via compression. In *Proceedings of Computer Science in Russia, CSR 2007*, number 4649 in *Lecture Notes in Computer Science*, pages 249–258. Springer, 2007.
- [96] R. C. Lyndon and P. E. Schupp. *Combinatorial Group Theory*. Springer, 1977.
- [97] J. Macdonald. Compressed words and automorphisms in fully residually free groups. *International Journal of Algebra and Computation*, 20(3):343–355, 2010.
- [98] G. S. Makanin. The problem of solvability of equations in a free semigroup. *Math. Sbornik*, 103:147–236, 1977. In Russian; English translation in *Math. USSR Sbornik 32, 1977*.
- [99] N. Markey and P. Schnoebelen. Model checking a path. In *Proceedings of the 14th International Conference on Concurrency Theory, CONCUR 2003*, number 2761 in *Lecture Notes in Computer Science*, pages 248–262, 2003.
- [100] N. Markey and P. Schnoebelen. A PTIME-complete matching problem for SLP-compressed words. *Information Processing Letters*, 90(1):3–6, 2004.
- [101] S. Maruyama, M. Takeda, M. Nakahara, and H. Sakamoto. An online algorithm for lightweight grammar-based compression. In *Proceedings of the 1st International Conference on Data Compression, Communications and Processing, CCP 2011*, pages 19–28, 2011.

- [102] S. Maruyama, Y. Tanaka, H. Sakamoto, and M. Takeda. Context-sensitive grammar transform: Compression and pattern matching. In *Proceedings of the 15th International Symposium on String Processing and Information Retrieval, SPIRE 2008*, number 5280 of *Lecture Notes in Computer Science*, pages 27–38. Springer, 2008.
- [103] W. Matsubara, S. Inenaga, A. Ishino, A. Shinohara, T. Nakamura, and K. Hashimoto. Efficient algorithms to compute compressed longest common substrings and compressed palindromes. *Theor. Comput. Sci.*, 410(8-10):900–913, 2009.
- [104] A. J. Mayer and L. J. Stockmeyer. The complexity of word problems—this time with interleaving. *Information and Computation*, 115(2):293–311, 1994.
- [105] K. Mehlhorn, R. Sundar, and C. Uhrig. Maintaining dynamic sequences under equality tests in polylogarithmic time. *Algorithmica*, 17(2):183–198, 1997.
- [106] M. Miyazaki, A. Shinohara, and M. Takeda. An improved pattern matching algorithm for strings in terms of straight-line programs. In *Proceedings of the 8th Annual Symposium on Combinatorial Pattern Matching, CPM 1997*, number 1264 in *Lecture Notes in Computer Science*, pages 1–11. Springer, 1997.
- [107] A. G. Myasnikov, A. Ushakov, and D. W. Won. Power circuits, exponential algebra, and time complexity. *International Journal of Algebra and Computation*, 22(6), 2012.
- [108] M.-J. Nederhof and G. Satta. The language intersection problem for non-recursive context-free grammars. *Information and Computation*, 192(2):172–184, 2004.
- [109] P. S. Novikov. On the algorithmic unsolvability of the word problem in group theory. *American Mathematical Society, Translations, II. Series*, 9:1–122, 1958.
- [110] C. H. Papadimitriou. *Computational Complexity*. Addison Wesley, 1994.
- [111] H. Petersen. The membership problem for regular expressions with intersection is complete in LOGCFL. In *Proceedings of the 19th Annual Symposium on Theoretical Aspects of Computer Science, STACS 2002*, number 2285 in *Lecture Notes in Computer Science*, pages 513–522. Springer, 2002.
- [112] W. Plandowski. Testing equivalence of morphisms on context-free languages. In *Proceedings of the 2nd Annual European Symposium on Algorithms, ESA 1994*, number 855 in *Lecture Notes in Computer Science*, pages 460–470. Springer, 1994.
- [113] W. Plandowski. Satisfiability of word equations with constants is in PSPACE. *Journal of the ACM*, 51(3):483–496, 2004.
- [114] W. Plandowski and W. Rytter. Application of Lempel-Ziv encodings to the solution of word equations. In *Proceedings of the 25th International Colloquium on Automata, Languages and Programming, ICALP 1998*, number 1443 in *Lecture Notes in Computer Science*, pages 731–742. Springer, 1998.
- [115] W. Plandowski and W. Rytter. Complexity of language recognition problems for compressed words. In J. Karhumäki, H. A. Maurer, G. Paun, and G. Rozenberg, editors, *Jewels are Forever, Contributions on Theoretical Computer Science in Honor of Arto Salomaa*, pages 262–272. Springer, 1999.
- [116] R. Radicioni and A. Bertoni. Grammatical compression: compressed equivalence and other problems. *Discrete Mathematics & Theoretical Computer Science*, 12(4):109–126, 2010.
- [117] D. Robinson. *Parallel Algorithms for Group Word Problems*. PhD thesis, University of California, San Diego, 1993.

- [118] J. M. Robson and V. Diekert. On quadratic word equations. In *Proceedings of the 16th Annual Symposium on Theoretical Aspects of Computer Science, STACS 1999*, number 1563 in Lecture Notes in Computer Science, pages 217–226. Springer, 1999.
- [119] W. Rytter. Application of Lempel-Ziv factorization to the approximation of grammar-based compression. *Theoretical Computer Science*, 302(1–3):211–222, 2003.
- [120] W. Rytter. Grammar compression, LZ-encodings, and string algorithms with implicit input. In *Proceedings of the 31st International Colloquium on Automata, Languages and Programming, ICALP 2004*, number 3142 in Lecture Notes in Computer Science, pages 15–27. Springer, 2004.
- [121] H. Sakamoto. A fully linear-time approximation algorithm for grammar-based compression. *J. Discrete Algorithms*, 3(2-4):416–430, 2005.
- [122] M. Schaefer, E. Sedgwick, and D. Štefankovič. Algorithms for normal curves and surfaces. In *Proceedings of the 8th Annual International Conference on Computing and Combinatorics, COCOON 2002*, number 2387 in Lecture Notes in Computer Science, pages 370–380. Springer, 2002.
- [123] M. Schaefer, E. Sedgwick, and D. Štefankovič. Recognizing string graphs in NP. *Journal of Computer and System Sciences*, 67(2):365–380, 2003.
- [124] M. Schaefer, E. Sedgwick, and D. Štefankovič. Computing Dehn twists and geometric intersection numbers in polynomial time. In *Proceedings of the 20th Annual Canadian Conference on Computational Geometry, CCCG 2008*, 2008.
- [125] S. Schleimer. Polynomial-time word problems. *Commentarii Mathematici Helvetici*, 83(4):741–765, 2008.
- [126] M. Schmidt-Schauß. Polynomial equality testing for terms with shared substructures. Technical Report Report 21, Institut für Informatik, J. W. Goethe-Universität Frankfurt am Main, 2005.
- [127] M. Schmidt-Schauß. Matching of compressed patterns with character-variables. In *Proceedings of the 23rd International Conference on Rewriting Techniques and Applications, RTA 2012*, number 15 of LIPIcs, pages 272–287. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2012.
- [128] M. Schmidt-Schauß, D. Sabel, and A. Anis. Congruence closure of compressed terms in polynomial time. In *Proceedings of the 8th International Symposium on Frontiers of Combining Systems, FroCos 2011*, number 6989 of *Lecture Notes in Computer Science*, pages 227–242. Springer, 2011.
- [129] M. Schmidt-Schauß and G. Schnitger. Fast equality test for straight-line compressed strings. *Inf. Process. Lett.*, 112(8-9):341–345, 2012.
- [130] H.-U. Simon. Word problems for groups and contextfree recognition. In *Proceedings of Fundamentals of Computation Theory, FCT 1979*, pages 417–422. Akademie-Verlag, 1979.
- [131] J. Stillwell. *Classical Topology and Combinatorial Group Theory (2nd edition)*. Springer, 1995.
- [132] I. H. Sudborough. On the tape complexity of deterministic context-free languages. *Journal of the ACM*, 25(3):405–414, 1978.
- [133] A. Tiskin. Faster subsequence recognition in compressed strings. *Journal of Mathematical Sciences*, 158(5):759–769, 2009.

- [134] A. Tiskin. Towards approximate matching in compressed strings: Local subsequence recognition. In *Proceedings of the 6th International Computer Science Symposium in Russia, CSR 2011*, number 6651 of *Lecture Notes in Computer Science*, pages 401–414. Springer, 2011.
- [135] S. Toda. PP is as hard as the polynomial-time hierarchy. *SIAM J. Comput.*, 20(5):865–877, 1991.
- [136] A. Tozawa and Y. Minamide. Complexity results on balanced context-free languages. In *Proceedings of the 10th International Conference on Foundations of Software Science and Computational Structures, FOSSACS 2007*, number 4423 of *Lecture Notes in Computer Science*, pages 346–360. Springer, 2007.
- [137] H. Venkateswaran. Properties that characterize LOGCFL. *Journal of Computer and System Sciences*, 43:380–404, 1991.
- [138] K. W. Wagner. The complexity of combinatorial problems with succinct input representation. *Acta Informatica*, 23(3):325–356, 1986.
- [139] K. W. Wagner. Leaf language classes. In *Proceedings of the 4th International Conference on Machines, Computations, and Universality, MCU 2004*, number 3354 of *Lecture Notes in Computer Science*, pages 60–81. Springer, 2005.
- [140] M. K. Warmuth and D. Haussler. On the complexity of iterated shuffle. *Journal of Computer and System Sciences*, 28(3):345–358, 1984.
- [141] D. T. Wise. Research announcement: the structure of groups with a quasiconvex hierarchy. *Electronic Research Announcements in Mathematical Sciences*, 16:44–55, 2009.
- [142] T. Yamamoto, H. Bannai, S. Inenaga, and M. Takeda. Faster subsequence and don’t-care pattern matching on compressed texts. In *Proceedings of the 22nd Annual Symposium on Combinatorial Pattern Matching, CPM 2011*, number 6661 of *Lecture Notes in Computer Science*, pages 309–322. Springer, 2011.
- [143] A. C.-C. Yao. On the evaluation of powers. *SIAM J. Comput.*, 5(1):100–103, 1976.
- [144] J. Ziv and A. Lempel. A universal algorithm for sequential data compression. *IEEE Transactions on Information Theory*, 23(3):337–343, 1977.