THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Normalization of Sequential Top-Down Tree-to-Word Transducers

Citation for published version:

Laurence, G, Lemay, A, Niehren, J, Staworko, S & Tommasi, M 2011, Normalization of Sequential Top-Down Tree-to-Word Transducers. in *Language and Automata Theory and Applications: 5th International Conference, LATA 2011, Tarragona, Spain, May 26-31, 2011. Proceedings*. vol. 6638, Springer Berlin Heidelberg, pp. 354-365. https://doi.org/10.1007/978-3-642-21254-3_28

Digital Object Identifier (DOI):

[10.1007/978-3-642-21254-3_28](https://doi.org/10.1007/978-3-642-21254-3_28)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Peer reviewed version

Published In:

Language and Automata Theory and Applications

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

Normalization of Sequential Top-Down Tree-to-Word Transducers

Grégoire Laurence^{1,2}, Aurélien Lemay^{1,2}, Joachim Niehren^{1,3},
Sławek Staworko^{1,2}, and Marc Tommasi^{1,2}

¹ Mostrare project, INRIA & LIFL (CNRS UMR8022)

² University of Lille

³ INRIA, Lille

Abstract. We study normalization of deterministic sequential top-down tree-to-word transducers (STWs), that capture the class of deterministic top-down nested-word to word transducers. We identify the subclass of *earliest* STWs (eSTWs) that yield unique normal forms when minimized. The main result of this paper is an effective normalization procedure for STWs. It consists of two stages: we first convert a given STW to an equivalent eSTW, and then, we minimize the eSTW.

1 Introduction

The classical problems on transducers are equivalence, minimization, learning, type checking, and functionality [2, 13, 14, 6]. Except for the latter two questions, one usually studies deterministic transducers because non-determinism quickly leads to fundamental limitations. For instance, equivalence of non-deterministic string transducers is known to be undecidable [8]. We thus follow the tradition to study classes of deterministic transducers. The problems of equivalence, minimization, and learning are often solved using unique normal representation of transformations definable with a transducer from a given class [9, 7, 5, 11]. Normalization i.e., constructing the normal form of a given transducer, has been studied independently for various classes, including string transducers [4, 3], top-down tree transducers [5], and bottom-up tree transducers [7].

In this paper, we study the normalization problem for the class of deterministic sequential top-down tree-to-word transducers (STWs). STWs are finite state machines that traverse the input tree in top-down fashion and at every node produce words obtained by the concatenation of constant words and the results from processing the child nodes. The main motivation to study this model is because tree-to-word transformations are better suited to model general XML transformations as opposed to tree-to-tree transducers [5, 11, 14]. This follows from the observation that general purpose XML transformation languages, like XSLT, allow to define transformations from XML documents to arbitrary, not necessarily structured, formats. Also, STWs capture a large subclass of deterministic nested-word to word transducers (dn2w), which have recently been the object of an enlivened interest [6, 15, 16].

Expressiveness of STWs suffers from two limitations: 1) every node is visited exactly once, and 2) the nodes are visited in the fix left-to-right preorder traversal of the input tree. Consequently, STWs cannot express transformations that reorder the nodes of the input tree or make multiple copies of a part of the input document. STWs remain, however, very powerful and are capable of: concatenation in the output, producing arbitrary context-free languages, deleting inner nodes, and verifying that the input tree belongs to the domain even when deleting parts of it. These features are often missing in tree-to-tree transducers, and for instance, make STWs incomparable with the class of top-down tree-to-tree transducers [5, 11].

Normal forms of transducers are typically obtained in two steps: output normalization followed by machine minimization. A natural way of output normalization is (re)arranging output words among the transitions rules so that the output is produced as soon as possible when reading the input, and thus transducers producing output in this fashion are called *earliest*. Our method subscribes to this approach but we note that it is a challenging direction that is not always feasible in the context of tree transformations. For instance, it fails for bottom-up tree-to-tree transducers, where *ad-hoc* solutions need to be employed [7].

We propose a natural normal form for STWs based on the notion of being earliest for STWs and define the corresponding class of *earliest* STWs (eSTWs) using easy to verify structural requirements. We present an effective procedure to convert an STW to an equivalent eSTW. This process is very challenging and requires novel tools on word languages. We point out that while this procedure works in time polynomial in the size of the output eSTW, we only know a doubly-exponential upper-bound and a single-exponential lower bound of the size of the output eSTW. This high complexity springs from the fact that the output language of an STW may be an arbitrary context-free language. We also show that minimization of earliest STWs is in PTIME thanks to a fundamental property: two equivalent eSTWs have rules of the same form and allow bisimulation. General STWs are unlikely to enjoy a similar property because their minimization is NP-complete.

Overall, we obtain an effective normalization procedure for STWs. Our results also offer an important step towards a better understanding of the same problem for dn2Ws because STWs capture a large class of top-down dn2Ws modulo the standard first-child next-sibling encoding and the conversion from one model to another can be done efficiently [16]. It is a significant result because there exist arguments suggesting that arbitrary dn2Ws are unlikely to have natural normal forms [1].

Organization. In Section 2 we present basic notions and introduce STWs and eSTWs. Section 3 introduces important tools on word languages and presents an STW to eSTW conversion algorithm. In Section 4 we deal with minimization of STWs and eSTWs. Section 5 summarizes our work and outlines future directions. Because of space restrictions we omit the proofs, which can be found in the full version at <http://hal.inria.fr/inria-00566291/en/>.

2 Sequential Top-down Tree-to-word Transducers

A *ranked alphabet* is a finite set of ranked symbols $\Sigma = \bigcup_{k \geq 0} \Sigma^{(k)}$ where $\Sigma^{(k)}$ is the set of k -ary symbols. We assume that every symbol has a unique arity i.e., $\Sigma^{(i)} \cap \Sigma^{(j)} = \emptyset$ for $i \neq j$. We use $f, g, \dots$ to range over symbols of non negative arity and $a, b, \dots$ to range over *constants* i.e., symbols of arity 0. We write $f^{(k)}$ to indicate that $f \in \Sigma^{(k)}$ if the arity of f is not known from the context. A *tree* is a ranked ordered term over Σ . We use $t, t_0, t_1, \dots$ to range over trees. For instance, $t_0 = f(a, g(b))$ is a tree over $\Sigma = \{f^{(2)}, g^{(1)}, a^{(0)}, b^{(0)}\}$.

For a finite set Δ of symbols by Δ^* we denote the free monoid on Δ . We write $u \cdot v$ for the concatenation of two words u and v and ε for the empty word. We use $a, b, \dots$ to range over Δ and $u, v, w, \dots$ to range over Δ^* . For a word w by $|w|$ we denote its length. Given a word $u = u_p \cdot u_f \cdot u_s$, u_p is a *prefix* of u , u_f a *factor* of u , and u_s a *suffix* of u . The *longest common prefix* of a nonempty set of words W , denoted $lcp(W)$, is the longest word u that is a prefix of every word in W . Analogously, we define the *longest common suffix* $lcs(W)$.

Definition 1. A deterministic sequential top-down tree-to-word transducer (STW) is a tuple $M = (\Sigma, \Delta, Q, \text{init}, \delta)$, where Σ is a ranked alphabet of input trees, Δ is a finite alphabet of output words, Q is a finite set of states, $\text{init} \in \Delta^* \cdot Q \cdot \Delta^*$ is the initial rule, δ is a partial transition function from $Q \times \Sigma$ to $(\Delta \cup Q)^*$ such that if $\delta(q, f^{(k)})$ is defined, then it has exactly k occurrences of elements from Q . By STWs we denote the class of deterministic sequential top-down tree-to-word transducers.

In the sequel, if $u_0 \cdot q_0 \cdot u_1$ is the initial rule, then we call q_0 the initial state. Also, we often view δ as a set of *transition rules* i.e., a subset of $Q \times \Sigma \times (\Delta \cup Q)^*$, which allows us to quantify over δ . The *size* of the STW M is the number of its states and the lengths of its rules, including the lengths of words used in the rules. The semantics of the STW M is defined with the help of auxiliary partial functions T_q (for $q \in Q$), recursively defined on the structure of trees as follows:

$$T_q(f(t_1, \dots, t_k)) = \begin{cases} u_0 \cdot T_{q_1}(t_1) \cdot u_1 \cdot \dots \cdot T_{q_k}(t_k) \cdot u_k, & \text{if } \delta(q, f) = u_0 \cdot q_1 \cdot u_1 \cdot \dots \cdot q_k \cdot u_k, \\ \text{undefined}, & \text{if } \delta(q, f) \text{ is undefined.} \end{cases}$$

The *transformation* T_M defined by M is a partial function mapping trees over Σ to words over Δ defined by $T_M(t) = u_0 \cdot T_{q_0}(t) \cdot u_1$, where $\text{init} = u_0 \cdot q_0 \cdot u_1$. Two transducers are *equivalent* iff they define the same transformation.

Example 1. We fix the input alphabet $\Sigma = \{f^{(2)}, g^{(1)}, a^{(0)}\}$ and the output alphabet $\Delta = \{a, b, c\}$. The STW M_1 has the initial rule q_0 and the following transition rules:

$$\delta(q_0, f) = q_1 \cdot ac \cdot q_1, \quad \delta(q_1, g) = q_1 \cdot abc, \quad \delta(q_1, a) = \varepsilon.$$

It defines the transformation $T_{M_1}(f(g^m(a), g^n(a))) = (abc)^m ac(abc)^n$, where $m, n \geq 0$, and T_{M_1} is undefined on all other input trees. The STW M_2 has the initial rule p_0 and these transition rules:

$$\begin{aligned} \delta(p_0, f) &= p_1 \cdot p_3 \cdot ab, & \delta(p_1, g) &= a \cdot p_2, & \delta(p_2, g) &= ab \cdot p_3, & \delta(p_3, g) &= p_3, \\ \delta(p_0, a) &= ba, & \delta(p_1, a) &= \varepsilon, & \delta(p_2, a) &= \varepsilon, & \delta(p_3, a) &= \varepsilon. \end{aligned}$$

Now, $T_{M_2}(a) = ba$ and for $n \geq 0$, the result of $T_{M_2}(f(g^m(a), g^n(a)))$ is ab for $m = 0$, aab for $m = 1$, and $aabab$ for $m \geq 2$; T_{M_2} is undefined for all other input trees. Note that p_3 is a *deleting* state: it does not produce any output but allows to check that the input tree belongs to the domain of the transducer. $\square$

In the sequel, to simplify notation we assume every state belongs to exactly one transducer, and so T_q above is defined in unambiguous manner. We consider only trimmed STWs i.e., transducers where all states define a nonempty transformation and are accessible from the initial rule. Also, by dom_q we denote the set $dom(T_q)$, the domain of T_q i.e., the set of trees on which T_q is defined, and by L_q the range of T_q i.e., the set of words returned by T_q . For instance, $dom_{q_0} = \{f(g^m(a), g^n(a)) \mid m, n \geq 0\}$ and $L_{q_0} = (abc)^* ac(abc)^*$. We observe that dom_q is a regular tree language and L_q is a context-free word language (CFL).

Next, we introduce the notion of being *earliest* that allows us to identify normal forms of transformations definable with STWs. It is a challenging task because the notion of being earliest needs to be carefully crafted so that every transducer can be made earliest. Take, for instance, the transformation *turn* that takes a tree over $\Sigma = \{a^{(1)}, b^{(1)}, \perp^{(0)}\}$ and returns the sequence of its labels in the reverse order e.g., $turn(a(b(b(\perp)))) = bba$. It is definable with a simple STW.

$$\delta(q_{turn}, a) = q_{turn} \cdot a, \quad \delta(q_{turn}, b) = q_{turn} \cdot b, \quad \delta(q_{turn}, \perp) = \varepsilon.$$

One way to view the transformation is a preorder traversal of the input tree that produces one output word upon entering the node and another word prior to leaving the node. When analyzing *turn* from this perspective, the earliest moment to produce any output is when the control reaches $\perp$, and in fact, the whole output can be produced at that point because all labels have been seen. This requires storing the label sequence in memory, which is beyond the capabilities of a finite state machine, and thus, *turn* cannot be captured with a transducer satisfying this notion of being earliest.

We propose a notion of being *earliest* that is also based on preorder traversal but with the difference that both output words are specified on entering the node and the output of a node is constructed right before leaving the node. Intuitively, we wish to *push up* all possible factors in the rules. Clearly, the STW above satisfies the condition. We remark that in some cases the output words in the rule can be placed in several positions, e.g. the rule $\delta(q_1, g) = q_1 \cdot abc$ in M_1 (Examples 1) can be replaced by $\delta(q_1, g) = abc \cdot q_1$ without changing the semantics of M_1 . Consequently, we need an additional requirement that resolves this ambiguity: intuitively, we wish to *push left* the words in a rule as much as possible.

Definition 2. An STW $M = (\Sigma, \Delta, Q, \text{init}, \delta)$ is earliest (eSTW) iff the following two conditions are satisfied:

- (E₁) $\text{lcp}(L_q) = \varepsilon$ and $\text{lcs}(L_q) = \varepsilon$ for every state q ,
- (E₂) $\text{lcp}(L_{q_0} \cdot u_1) = \varepsilon$ for the initial rule $u_0 \cdot q_0 \cdot u_1$ and for every transition $\delta(q, f) = u_0 \cdot q_1 \dots q_k \cdot u_k$ and $1 \leq i \leq k$ we have $\text{lcp}(L_{q_i} \cdot u_i \dots L_{q_k} \cdot u_k) = \varepsilon$.

Intuitively, the condition (E₁) ensures that no factor can be pushed up in the traversal and (E₂) ensures that no factor can be pushed left. We note that (E₁) and (E₂) can be efficiently checked in an STW because we need only to check that the results of lcp and lcs are ε . The main contribution of this paper follows.

Theorem 1. For every STW there exists a unique minimal equivalent eSTW.

The proof consists of an effective procedure that works in two stages: In the first stage we normalize the outputs i.e., from the input STW we construct an equivalent eSTW, and in the second stage we minimize the obtained eSTW. The first stage is, however, quite complex as illustrated in the following example.

Example 2 (contd. Example 1). M_1 is not earliest because (E₁) is not satisfied at q_0 : every word of $L_{q_0} = (abc)^*ac(abc)^*$ begins with a i.e., $\text{lcp}(L_{q_0}) = a$, and ends with c i.e., $\text{lcs}(L_{q_0}) = c$. Consequently, we need to *push up* these two symbols to the new initial rule $a \cdot q'_0 \cdot c$, but we also need to retract them from the rule $\delta(q_0, f) = q_1 \cdot ac \cdot q_1$ producing a new state q'_0 and new rules for this state. Essentially, we need to push the symbol a to the left through the first occurrence of q_1 and push the symbol c to the right through the second occurrence of q_1 . Pushing symbols through states produces again new states with rules obtained by reorganizing the output words. Finally, we obtain

$$\delta'(q'_0, f) = q'_1 \cdot q''_1, \quad \delta'(q'_1, g) = bca \cdot q'_1, \quad \delta'(q''_1, g) = cab \cdot q''_1, \quad \delta'(q'_1, a) = \delta'(q''_1, a) = \varepsilon.$$

M_2 is not earliest because (E₂) is not satisfied by $\delta(p_0, f) = p_1 \cdot p_3 \cdot ab$: every word produced by this rule starts with a . First, we push the word ab through the state p_3 , and then we push the symbol a through the state p_1 . Pushing through p_3 is easy because it is a deleting state and the rules do not change. Pushing through p_1 requires a recursive push through the states of the rules of p_1 and this process affects the rules of p_2 . Finally, we obtain an eSTW with the initial rule p'_0 and the transition rules

$$\begin{aligned} \delta'(p'_0, f) &= a \cdot p'_1 \cdot b \cdot p'_3, & \delta'(p'_1, g) &= a \cdot p'_2, & \delta'(p'_2, g) &= ba \cdot p'_3, & \delta'(p'_3, g) &= p'_3, \\ \delta'(p'_0, a) &= ba, & \delta'(p'_1, a) &= \varepsilon, & \delta'(p'_2, a) &= \varepsilon, & \delta'(p'_3, a) &= \varepsilon. \quad \square \end{aligned}$$

3 Output Normalization

The first phase of normalization of an STW consists of constructing an equivalent eSTW, which involves changing the placement of the factors in the rules of the transducer and deals mainly with the output. Consequently, we begin with several notions and constructions inspired by the conditions (E₁) and (E₂) but set in a simpler setting of word languages. We consider only nonempty languages because in trimmed STWs the ranges of the states are always nonempty.

3.1 Reducing Languages

Enforcement of **(E₁)** corresponds to what we call constructing the *reduced decomposition* of a language. A nonempty language L is *reduced* iff $lcp(L) = \varepsilon$ and $lcs(L) = \varepsilon$. Note that the assumption that we work with a nonempty language is essential here. Now, take a nonempty language L , that is not necessarily reduced. We decompose it into its *reduced core* $Core(L)$ and two words $Left(L)$ and $Right(L)$ such that $Core(L)$ is reduced and

$$L = Left(L) \cdot Core(L) \cdot Right(L). \quad (1)$$

We observe that different decompositions are possible. For instance, $L = \{a, aba\}$ has two decompositions $L = a \cdot \{\varepsilon, ba\} \cdot \varepsilon$ and $L = \varepsilon \cdot \{\varepsilon, ab\} \cdot a$. We resolve the ambiguity by choosing the former decomposition because it is consistent with **(E₁)** and **(E₂)** which indicate to *push to the left*. Formally, $Left(L) = lcp(L)$ and $Right(L) = lcs(L')$, where $L = Left(L) \cdot L'$. The reduced core $Core(L)$ is obtained from (1). As an example, the reduced decomposition of $L_{q_0} = (abc)^*ac(abc)^*$ from Example 1 is $Left(L_{q_0}) = a$, $Right(L_{q_0}) = c$, and $Core(L_{q_0}) = (bca)^*(cba)^*$.

3.2 Pushing Words Through Languages

In this subsection, we work with *nonempty* and *reduced* languages only. Condition **(E₂)** introduces the problem that we call pushing words through languages. To illustrate it, suppose we have a language $L = \{\varepsilon, a, aa, aaab\}$ and a word $w = aab$, which together give $L \cdot w = \{aab, aaab, aaaab, aaabaab\}$. The goal is to find the longest prefix v of w such that $L \cdot w = v \cdot L' \cdot u$, where $w = v \cdot u$ and L' is some derivative of L . Intuitively speaking, we wish to push (a part of) the word w forward i.e., from right to left, through the language L . In the example above, the solution is $v = aa$, $L' = \{\varepsilon, a, aa, abaa\}$, and $u = b$ (note that L' is different from L). In this section, we show that this process is always feasible and for CFLs it is constructive.

The result of pushing a word w through a language L will consist of three words: $push(L, w)$ the longest part of w that can be pushed through L , $rest(L, w)$ the part that cannot be pushed through, and $offset(L, w)$ a special word that allows to identify the corresponding derivative of L . There are three classes of languages that need to be considered, which we present next together with an outline of how the pushing is done.

The first class contains only the *trivial* language $L = \{\varepsilon\}$ e.g., the range of the state p_3 of M_2 in Example 1. This language allows every word to be pushed through and it never changes in the process. For instance, if $w_0 = ab$, then $push(L_{p_3}, w_0) = ab$, $rest(L_{p_3}, w_0) = \varepsilon$, and $offset(L_{p_3}, w_0) = \varepsilon$.

The second class consists of non-trivial periodic languages, essentially languages contained in the Kleene closure of some period word. An example is $L_{q_1} = (abc)^* = \{\varepsilon, abc, abcabc, \dots\}$ whose period is abc . Periodic languages allow to push multiplicities of the period and then some prefix of the period e.g., if we take $w_1 = abcabcaba$, then $push(L_{q_1}, w_1) = abcabcab$ and $rest(L_{q_1}, w_1) = a$. The offset here is the corresponding prefix of the period: $offset(L_{q_1}, w_1) = ab$.

The third class contains all remaining languages i.e., non-trivial non-periodic languages. Interestingly, we show that for a language in this class there exists a word that is the longest word that can be pushed fully through the language, and furthermore, every other word that can be pushed through is a prefix of this word. For instance, for $L_{p_1} = \{\varepsilon, a, aab\}$ from Example 1, aa is the longest word that can be pushed through. If we take $w_2 = ab$, then we get $push(L_{p_1}, w_2) = a$ and $rest(L_{p_1}, w_2) = b$. Here, the offset is the prefix of aa that has been already pushed through: $offset(L_{p_1}, w_2) = a$. Note that this class also contains the languages that do not permit any pushing through e.g., $L_{p_0} = \{ba, ab, aab\}$ does not allow pushing through because it contains two words that start with a different symbol.

We now define formally the pushing process. First, for $L \subseteq \Delta^*$ we define the set of words that can be pushed fully through L :

$$Shovel(L) = \{w \in \Delta^* \mid w \text{ is a common prefix of } L \cdot w\}.$$

For instance, $Shovel(L_{p_1}) = \{\varepsilon, a, aa\}$ and $Shovel(L_{q_0}) = (abc)^* \cdot \{\varepsilon, a, ab\}$. We note that $Shovel(\{\varepsilon\}) = \Delta^*$ and $Shovel(L)$ always contains at least one element ε because L is assumed to be nonempty. Also, as we prove in appendix, $Shovel(L)$ is prefix-closed and totally ordered by the prefix relation.

Next, we define periodic languages (cf. [12]). A language $L \subseteq \Delta^*$ is *periodic* iff there exists a nonempty word $v \in \Delta^*$, called a *period* of L , such that $L \subseteq v^*$. A word w is *primitive* if there is no v and $n \geq 0$ such that $w = v^n$. Recall from [12] that every non-trivial periodic language L has a unique primitive period, which we denote $Period(L)$. For instance, the language $\{\varepsilon, abab, abababab\}$ is periodic and its primitive period is ab ; $abab$ is also its period but not primitive. In the sequel, by $Prefix(w)$ we denote the set of prefixes of the word w .

Proposition 1. *Given a reduced and non-trivial language L , $Shovel(L)$ is infinite iff L is periodic. Furthermore, if L is periodic then $Shovel(L) = Period(L)^* \cdot Prefix(Period(L))$.*

This result and the observations beforehand lead to three relevant cases in the characterisation of $Shovel(L)$ for a language L .

- 0° $L = \{\varepsilon\}$ (trivial language), and then $Shovel(L) = \Delta^*$,
- 1° L is periodic, $L \neq \{\varepsilon\}$, and then $Shovel(L) = Period(L)^* \cdot Prefix(Period(L))$.
- 2° L is non-periodic, and $Shovel(L) = Prefix(v)$ for some $v \in Shovel(L)$.

Now, suppose we wish to *push* a word $w \in \Delta^*$ through a language $L \subseteq \Delta^*$ and let $s = \max_{\leq_{\text{prefix}}} (Prefix(w) \cap Shovel(L))$ and $w = s \cdot r$. We define $push(L, w)$, $rest(L, w)$, and $offset(L, w)$ depending on the class L belongs to:

- 0° $L = \{\varepsilon\}$: $push(L, w) = w$, $rest(L, w) = \varepsilon$, and $offset(L, w) = \varepsilon$.
- 1° L is non-trivial and periodic: $s = Period(L)^k \cdot o$ for some (maximal) proper prefix o of $Period(L)$, and we assign $push(L, w) = s$, $rest(L, w) = r$, and $offset(L, w) = o$.
- 2° L is non-periodic: $push(L, w) = s$, $rest(L, w) = r$, and $offset(L, w) = s$.

Offsets play a central role in the output normalization procedure, which is feasible thanks to the following result.

Proposition 2. *The set $\{offset(L, w) \mid w \in \Delta^*\}$ is finite for any reduced L .*

3.3 Pushing Words Backwards

Until now, we have considered the problem of pushing a word through a language from right to left. However, in Example 1 if we consider the second occurrence of q_1 in the rule $\delta(q_0, f) = q_1 \cdot ac \cdot q_1$, we realize that pushing words in the opposite direction needs to be investigated as well. These two processes are dual but before showing in what way, we present a natural extension of the free monoid Δ^* to a pregroup (or groupoid) $\mathbb{G}_\Delta$. It allows to handle pushing in two directions in a unified manner and simplifies the output normalization algorithm.

A *pregroup of words over Δ* is the set $\mathbb{G}_\Delta = \Delta^* \cup \{w^{-1} \mid w \in \Delta^+\}$, where w^{-1} is a term representing the inverse of a nonempty word w . This set comes with two operators, a unary inverse operator: $(w)^{-1} = w^{-1}$, $\varepsilon^{-1} = \varepsilon$, and $(w^{-1})^{-1} = w$ for $w \in \Delta^*$, and a partial extension of the standard concatenation that satisfies the following equations (complete definition in appendix): $w^{-1} \cdot w = \varepsilon$ and $w \cdot w^{-1} = \varepsilon$ for $w \in \Delta^*$, and $v^{-1} \cdot u^{-1} = (uv)^{-1}$ for $u, v \in \Delta^*$. We note that some expressions need to be evaluated diligently e.g., $ab \cdot (cb)^{-1} \cdot cd = ab \cdot b^{-1} \cdot c^{-1} \cdot cd = ad$, while some are undefined e.g., $ab \cdot a^{-1}$. In the sequel, we use $w, u, v, \dots$ to range over Δ^* only and $z, z_1, \dots$ to range over elements of $\mathbb{G}_\Delta$.

Now, we come back to pushing a word w backwards through L , which consists of finding $u \cdot v = w$ and L' such that $w \cdot L = u \cdot L' \cdot v$. We view this process as pushing the inverse w^{-1} through L i.e., we wish to find $u \cdot v = w$ such that $L \cdot w^{-1} = v^{-1} \cdot L' \cdot u^{-1}$ because then $L \cdot v^{-1} = v^{-1} \cdot L'$, and consequently, $w \cdot L = (u \cdot v) \cdot (v^{-1} \cdot L' \cdot v) = u \cdot L' \cdot v$.

But to define pushing backwards more properly we use another perspective based on the standard reverse operation of a word e.g., $(abc)^{\text{rev}} = cba$. Namely, pushing w backwards through L is essentially pushing w^{rev} through L^{rev} because $(w \cdot L)^{\text{rev}} = L^{\text{rev}} \cdot w^{\text{rev}}$ and if $L^{\text{rev}} \cdot w^{\text{rev}} = v_0 \cdot L_0 \cdot u_0$, then $w \cdot L = u_0^{\text{rev}} \cdot L_0^{\text{rev}} \cdot v_0^{\text{rev}}$. Thus $\text{push}(L, w^{-1}) = (\text{push}(L^{\text{rev}}, w^{\text{rev}})^{\text{rev}})^{-1}$, $\text{rest}(L, w^{-1}) = (\text{rest}(L^{\text{rev}}, w^{\text{rev}})^{\text{rev}})^{-1}$, and $\text{offset}(L, w^{-1}) = (\text{offset}(L^{\text{rev}}, w^{\text{rev}})^{\text{rev}})^{-1}$.

Now, the main condition of pushing words through languages is: for every L and $z \in \mathbb{G}_\Delta$ we have $L \cdot z = \text{push}(L, z) \cdot (\text{offset}(L, z)^{-1} \cdot L \cdot \text{offset}(L, z)) \cdot \text{rest}(L, z)$. Because the output normalization procedure works on STWs and not languages, to prove its correctness we need a stronger statement treating independently every word of the language.

Proposition 3. *Given a reduced and nonempty language $L \subseteq \Delta^*$ and $z \in \mathbb{G}_\Delta$, for any word $u \in L$*

$$u \cdot z = \text{push}(L, z) \cdot (\text{offset}(L, z)^{-1} \cdot u \cdot \text{offset}(L, z)) \cdot \text{rest}(L, z).$$

3.4 Output Normalization Algorithm

We fix an STW $M = (\Sigma, \Delta, Q, \text{init}, \delta)$ and introduce the following macros:

$$\begin{aligned} \hat{L}_q &= \text{Core}(L_q), & \text{Left}(q) &= \text{Left}(L_q), & \text{Right}(q) &= \text{Right}(L_q), \\ \text{push}(q, z) &= \text{push}(\hat{L}_q, z), & \text{offset}(q, z) &= \text{offset}(\hat{L}_q, z), & \text{rest}(q, z) &= \text{rest}(\hat{L}_q, z). \end{aligned}$$

Also, let $Offsets(q) = \{offset(q, z) \mid z \in \mathbb{G}_\Delta\}$ and note that by Proposition 2 it is finite. The constructed STW $M' = (\Sigma, \Delta, Q', init', \delta')$ has the following states

$$Q' = \{\langle q, w \rangle \mid q \in Q, w \in Offsets(q)\}.$$

Our construction ensures that $T_M = T_{M'}$ and for every $q \in Q$, every $z \in Offsets(q)$, and every $t \in dom_q$

$$T_{\langle q, z \rangle}(t) = z^{-1} \cdot Left(q)^{-1} \cdot T_q(t) \cdot Right(q)^{-1} \cdot z$$

If $init = u_0 \cdot q_0 \cdot u_1$, then $init' = u'_0 \cdot q'_0 \cdot u'_1$, where u'_0 , u'_1 , and q'_0 are calculated as follows:

- 1: $v := Right(q_0) \cdot u_1$
- 2: $q'_0 := \langle q_0, offset(q_0, v) \rangle$
- 3: $u'_0 := u_0 \cdot Left(q_0) \cdot push(q_0, v)$
- 4: $u'_1 := rest(q_0, v)$

For a transition rule $\delta(p, f) = u_0 \cdot p_1 \cdot u_1 \cdot \dots \cdot u_{k-1} \cdot p_k \cdot u_k$ and any $z \in Offsets(p)$ we introduce a rule $\delta'(\langle p, z \rangle, f) = u'_0 \cdot p'_1 \cdot u'_1 \cdot \dots \cdot u'_{k-1} \cdot p'_k \cdot u'_k$, where $u'_0, \dots, u'_k$ and $p'_1, \dots, p'_k$ are calculated as follows:

- 1: $z_k := Right(p_k) \cdot u_k \cdot Right(p)^{-1} \cdot z$
- 2: **for** $i := k, \dots, 1$ **do**
- 3: $u'_i := rest(p_i, z_i)$
- 4: $p'_i := \langle p_i, offset(p_i, z_i) \rangle$
- 5: $z_{i-1} := Right(p_{i-1}) \cdot u_{i-1} \cdot Left(p_i) \cdot push(p_i, z_i)$
- 6: $u'_0 := z^{-1} \cdot Left(p)^{-1} \cdot z_0$

where (for convenience of the presentation) we let $Right(p_0) = \varepsilon$. We remark that not all states in Q' are reachable from the initial rule and in fact the conversion procedure can identify the reachable states *on the fly*. This observation is the basis of a conversion algorithm that is polynomial in the size of the output.

Example 3. We normalize the STW M_1 from Example 1. The initial rule q_0 becomes $a \cdot \langle q_0, \varepsilon \rangle \cdot c$ with $Left(q_0) = a$ and $Right(q_0) = c$ being pushed up from q_0 but with nothing pushed through q_0 . The construction of the state $\langle q_0, \varepsilon \rangle$ triggers the normalization algorithm for the rule $\delta(q_0, f) = q_1 \cdot ac \cdot q_1$ with $Left(q_0) = a$ and $Right(q_0) = c$ to be retracted from left and right side resp. (and nothing pushed through since $z = \varepsilon$). This process can be viewed as a taking the left hand side of the original rule with the inverses of retracted words $a^{-1} \cdot q_1 \cdot ac \cdot q_1 \cdot c^{-1}$ and pushing words forward as much as possible, which gives $a^{-1} \cdot q_1 \cdot ac \cdot c^{-1} \cdot \langle q_1, c^{-1} \rangle$ and then $a^{-1} \cdot a \cdot \langle q_1, a \rangle \cdot \langle q_1, c^{-1} \rangle$. This gives $\delta'(\langle q_0, \varepsilon \rangle, f) = \langle q_1, a \rangle \cdot \langle q_1, c^{-1} \rangle$. Note that while $Offsets(q_1) = \{(bc)^{-1}, c^{-1}, \varepsilon, a, ab\}$, only two states are constructed.

Next, we need to construct rules for the new state $\langle q_1, a \rangle$ with $z = a$ and $Left(q_1) = Right(q_1) = \varepsilon$. We start with the rule $\delta(q_1, a) = \varepsilon$ and to its left hand side we add a^{-1} at the beginning and a at its end: $a^{-1} \cdot \varepsilon \cdot a = \varepsilon$, which yields the rule $\delta'(\langle q_1, a \rangle, a) = \varepsilon$. Now, for the rule $\delta(q_1, g) = q_1 \cdot abc$ we obtain the

expression $a^{-1} \cdot q_1 \cdot abca$. Recall that $L_{q_1} = (abc)^*$ is a periodic language, and so $push(q_1, abca) = abca$, $rest(q_1, abca) = \varepsilon$, and $offset(q_1, abca) = a$. Consequently, we obtain the rule $\delta'(\langle q_1, a \rangle, g) = bca \cdot \langle q_1, a \rangle$. Here, it is essential to use the offsets to avoid introducing a redundant state $\langle q_1, abca \rangle$ and entering an infinite loop. Similarly, we obtain: $\delta'(\langle q_1, c^{-1} \rangle, g) = cab \cdot \langle q_1, c^{-1} \rangle$ and $\delta'(\langle q_1, c^{-1} \rangle, a) = \varepsilon$. $\square$

Theorem 2. *For an STW M let M' be the STW obtained with the method described above. Then, M' is equivalent to M and satisfies **(E₁)** and **(E₂)**. Furthermore, M' can be constructed in time polynomial in the size M' , which is at most doubly-exponential in the size of M .*

Because of space limitations, the details on complexity have been omitted and can be found in the full version available online.

3.5 Exponential Lower Bound

First, we show that the size of a rule may increase exponentially.

Example 4. For $n \geq 0$ define an STW M_n over the input alphabet $\Sigma = \{f^{(2)}, a^{(0)}\}$ with the initial rule q_0 , and these transition rules (with $0 \leq i < n$):

$$\delta(q_i, f) = q_{i+1} \cdot q_{i+1}, \quad \delta(q_n, a) = a.$$

The transformation defined by M_n maps a perfect binary tree of height n to a string a^{2^n} . M_n is not earliest. To make it earliest we need to replace the initial rule by $a^{2^n} \cdot q_0(x_0)$ and the last transition rule by $\delta(q_n, a) = \varepsilon$. $\square$

The next example shows that also the number of states may become exponential.

Example 5. For $n \geq 0$ and take the STW N_n with $\Sigma = \{g_1^{(1)}, g_0^{(1)}, a_1^{(0)}, a_0^{(0)}\}$, the initial rule q_0 , and these transition rules (with $0 \leq i < n$):

$$\begin{aligned} \delta(q_i, g_0) &= q_{i+1}, & \delta(q_n, a_0) &= \varepsilon, \\ \delta(q_i, g_1) &= q_{i+1} \cdot a^{2^i}, & \delta(q_n, a_1) &= a^{2^n} \cdot \#. \end{aligned}$$

While the size of this transducer is exponential in n , one can easily compress the exponential factors a^{2^i} and obtain an STW of size linear in n (cf. Example 4). M_n satisfies **(E₁)** but it violates **(E₂)**, and defines the following transformation.

$$\begin{aligned} T_{N_n} = & \{(g_{b_0}(g_{b_1}(\dots g_{b_{n-1}}(a_0)\dots)), a^{\mathbf{b}}) \mid \mathbf{b} = (b_{n-1}, \dots, b_0)_2\} \cup \\ & \{(g_{b_0}(g_{b_1}(\dots g_{b_{n-1}}(a_1)\dots)), a^{2^n} \cdot \# \cdot a^{\mathbf{b}}) \mid \mathbf{b} = (b_{n-1}, \dots, b_0)_2\}, \end{aligned}$$

where $(b_{n-1}, \dots, b_0)_2 = \sum_i b_i \cdot 2^i$. The normalized version N'_n has the initial rule $\langle q_0, \varepsilon \rangle$ and these transition rules:

$$\begin{aligned} \delta'(\langle q_i, a^j \rangle, g_0) &= \langle q_{i+1}, a^j \rangle, & \delta'(\langle q_n, a^k \rangle, a_0) &= \varepsilon, \\ \delta'(\langle q_i, a^j \rangle, g_1) &= a^{2^i} \cdot \langle q_{i+1}, a^{j+2^i} \rangle, & \delta'(\langle q_n, a^k \rangle, a_1) &= a^{2^n-k} \# a^k, \end{aligned}$$

where $0 \leq i < n$, $0 \leq j < 2^i$, and $0 \leq k < 2^n$. We also remark that N'_n is the minimal eSTW that recognises T_{N_n} . $\square$

4 Minimization

In this section we investigate the problem of minimizing the size of a transducer. Minimization of eSTWs is simple and relies on testing the equivalence of eSTWs known to be in PTIME [16]. For an eSTW M the minimization procedure constructs a binary equivalence relation $\equiv_M$ on states such that $q \equiv_M q'$ iff $T_q = T_{q'}$. The result of minimization is the quotient transducer $M/\equiv_M$ obtained by choosing in every equivalence class C of $\equiv_M$ exactly one representative state $q \in C$, and then replacing in rules of M every state of C by q .

To show that the obtained eSTW is minimal among all eSTWs defining the same transformation, we use an auxiliary result stating that all eSTWs defining the same transformation use rules with the same distribution of the output words and allow bisimulation.

A *labeled path* is a word over $\bigcup_{k \geq 0} \Sigma^{(k)} \times \{1, \dots, k\}$, which identifies a node in a tree together with the labels of its ancestors: ε is the root node and if a node π is labeled with f , then $\pi \cdot (f, i)$ is its i -th child. By $\text{paths}(t)$ we denote the set of labeled paths of a tree t . For instance, for $t_0 = f(a, g(b))$ we get $\text{paths}(t_0) = \{\varepsilon, (f, 1), (f, 2), (f, 2) \cdot (g, 1)\}$. We extend the transition function δ to identify the state reached at a path π : $\delta(q, \varepsilon) = q$ and $\delta(q, \pi \cdot (f, i)) = q_i$, where $\delta(q, \pi) = q'$ and $\delta(q', f) = u_0 \cdot q_1 \cdot u_1 \cdot \dots \cdot q_k \cdot u_k$. Now, the lemma of interest.

Lemma 1. *Take two eSTWs $M = (\Sigma, \Delta, Q, \text{init}, \delta)$ and $M' = (\Sigma, \Delta, Q', \text{init}', \delta')$ defining the same transformation $T = T_M = T_{M'}$ and let $\text{init} = u_0 \cdot q_0 \cdot u_1$ and $\text{init}' = u'_0 \cdot q'_0 \cdot u'_1$. Then, $u_0 = u'_0$ and $u_1 = u'_1$, and for every $\pi \in \text{paths}(\text{dom}(T))$, we let $q = \delta(q_0, \pi)$ and $q' = \delta'(q'_0, \pi)$, and we have*

1. $T_q = T_{q'}$,
2. $\delta(q, f)$ is defined if and only if $\delta'(q', f)$ is, for every $f \in \Sigma$, and
3. if $\delta(q, f) = u_0 \cdot q_1 \cdot u_1 \cdot \dots \cdot q_k \cdot u_k$ and $\delta'(q', f) = u'_0 \cdot q'_1 \cdot u'_1 \cdot \dots \cdot q'_k \cdot u'_k$, then $u_i = u'_i$ for $0 \leq i \leq k$.

The proof is inductive and relies on properties **(E₁)** and **(E₂)**, and the determinism of the transducers. We show the correctness of our minimization algorithm by observing that it produces an eSTW whose size is smaller than the input one, and Lemma 1 essentially states that the result of minimization of two equivalent transducers is the same transducer (modulo state renaming). This argument also proves Theorem 1. We also point out that Lemma 1 (with $M = M'$) allows to devise a simpler and more efficient minimization algorithm along the lines of the standard dFA minimization algorithm [10].

Theorem 3. *Minimization of eSTWs is in PTIME.*

In STWs the output words may be arbitrarily distributed among the rules, which is the main pitfall of minimizing general STWs. This difficulty is unlikely to be overcome as suggested by the following result.

Theorem 4. *Minimization of STWs i.e., deciding whether for an STW M and $k \geq 0$ there exists an equivalent STW M' of size at most k , is NP-complete.*

5 Conclusions and Future Work

We have presented an effective normalization procedure for STWs, a subclass of top-down tree-to-word transducers closely related to a large subclass of nested-word to word transducers. One natural continuation of this work is find whether it can be extended to a Myhill-Nerode theorem for STWs, and then, to a polynomial learning algorithm. Also, the question of exact complexity of the normalization remains open. Finally, the model of STWs can be generalized to allow arbitrary non-sequential rules and multiple passes over the input tree.

References

1. Alur, R., Kumar, V., Madhusudan, P., Viswanathan, M.: Congruences for visibly pushdown languages. In: Automata, Languages and Programming. LNCS 3580 (2005) 1102–1114
2. Berstel, J., Boasson, L.: Transductions and context-free languages. Teubner Studienbücher (1979)
3. Choffru, C.: Contribution à l'étude de quelques familles remarquables de fonctions rationnelles. PhD thesis, Université de Paris VII (1978)
4. Choffrut, C.: Minimizing subsequential transducers: a survey. Theoretical Computer Science **292**(1) (2003) 131–143
5. Engelfriet, J., Maneth, S., Seidl, H.: Deciding equivalence of top-down XML transformations in polynomial time. Journal of Computer and System Science **75**(5) (2009) 271–286
6. Filiot, E., Raskin, J.F., Reynier, P.A., Servais, F., Talbot, J.M.: Properties of visibly pushdown transducers. In: Mathematical Foundations of Computer Science (MFCS). LNCS 6281 (2010) 355–367
7. Frieze, S., Seidl, H., Maneth, S.: Minimization of deterministic bottom-up tree transducers. In: Developments in Language Theory (DLT). LNCS 6224 (2010) 185–196
8. Griffiths, T.V.: The unsolvability of the equivalence problem for lambda-free nondeterministic generalized machines. Journal of the ACM **15**(3) (1968) 409–413
9. Gurari, E.M.: The equivalence problem for deterministic two-way sequential transducers is decidable. SIAM Journal on Computing **11**(3) (1982) 448–452
10. Hopcroft, J.E., Motwani, R., Ullman, J.D.: Introduction to Automata Theory, Languages, and Computation. 2nd edn. Addison Wesley (2001)
11. Lemay, A., Maneth, S., Niehren, J.: A learning algorithm for top-down xml transformations. In: ACM Symposium on Principles of Database Systems (PODS). (2010) 285–296
12. Lothaire, M., ed.: Combinatorics on Words. 2nd edn. Cambridge Mathematical Library. Cambridge University Press (1997)
13. Maneth, S.: Models of Tree Translation. PhD thesis, Leiden University (2003)
14. Martens, W., Neven, F., Gyssens, M.: Typechecking top-down XML transformations: Fixed input or output schemas. Information and Computation **206**(7) (2008) 806–827
15. Raskin, J.F., Servais, F.: Visibly pushdown transducers. In: Automata, Languages and Programming. LNCS 5126 (2008) 386–397
16. Staworko, S., Laurence, G., Lemay, A., Niehren, J.: Equivalence of nested word to word transducers. In: Fundamentals of Computer Theory (FCT). LNCS 5699 (2009) 310–322

A Appendix

A.1 Proof of Proposition 1

Below we present an alternative characterization of periodic languages which will be useful later on. This result is a direct consequence of Proposition 1.3.2 in [12].

Proposition 4. *A language L is periodic iff any pair of its words commute i.e., $w_1 \cdot w_2 = w_2 \cdot w_1$ for every $w_1, w_2 \in L$.*

We also present several observations about the shoveling capabilities of a reduced language.

Fact 1 *$Shovel(L)$ is nonempty for any language L .*

Proof. In fact, the empty word is a common prefix of any (even empty) set of words, and thus $\varepsilon \in Shovel(L)$ for any L . $\square$

Fact 2 *Given a reduced language L , if $Shovel(L)$ contains a nonempty word, then L contains the empty word.*

Proof. If $a \cdot w \in Shovel(L)$, then a is the first letter of every nonempty word in L , and since L is reduced, L must contain the empty word, or otherwise $lcp(L)$ would not be ε .

Fact 3 *If L is trivial, then $Shovel(L) = \Delta^*$.*

Proof. Trivial. $\square$

Proposition 5. *If L is reduced and nontrivial, then $Shovel(L)$ is prefix-closed and totally ordered by the prefix relation.*

Proof. Showing that $Shovel(L)$ is prefix-closed follows from the definition. Take any $w \in Shovel(L)$ and a prefix w' of w whose length is $|w'| = k$. Now, fix a word $v \in L$ and observe that w is a prefix of $v \cdot w$. Since w' is a prefix of w , then $v \cdot w'$ is a prefix of $v \cdot w$ and furthermore $v \cdot w'$ is of length at least k . Consequently, w' is a prefix of $v \cdot w'$.

We show that $Shovel(L)$ is ordered with a simple induction over the length of words in $Shovel(L)$. Take two words $w \cdot a, w \cdot b \in Shovel(L)$. Since L is nontrivial, there is some nonempty word $u \in L$. Now, $w \cdot a$ is a prefix of $u \cdot w \cdot a$ and $w \cdot b$ is a prefix of $u \cdot w \cdot b$. Since u has length at least 1, both $w \cdot a$ and $w \cdot b$ are prefixes of $u \cdot w$. Consequently, $a = b$. $\square$

And now the proof of Proposition 1. For the *if* part, take any nontrivial $L \subseteq w^*$ and observe that $w^k \in Shovel(L)$ for any $k \geq 0$. Furthermore, by Proposition 5 we get that $Shovel(L) = w^* \cdot Prefix(w)$.

For the *only if* part we point out that Proposition 4 characterises nontrivial periodic languages as exactly those that self-commute i.e., a nontrivial L is periodic iff $w_1 \cdot w_2 = w_2 \cdot w_1$ for any two words $w_1, w_2 \in L$.

First, we observe that by Fact 2 $\varepsilon \in L$. Next, take any $w_1, w_2 \in L$. Since $Shovel(L)$ is infinite, there exists a word $v \in Shovel(L)$ whose length is greater than $|w_1| + |w_2|$. In addition we remark that v is a prefix of v , $w_1 \cdot v$, and $w_2 \cdot v$. This allows us to infer that $v = w_1 \cdot v'$ and $v = w_2 \cdot v''$, which implies that $w_1 \cdot v = w_1 \cdot w_2 \cdot v''$ and $w_2 \cdot v = w_2 \cdot w_1 \cdot v'$. Since the length of v is greater than $|w_1| + |w_2|$, the aforementioned two words agree on the first $|w_1| + |w_2|$ letters and hence $w_1 \cdot w_2 = w_2 \cdot w_1$. $\square$

A.2 Proof of Proposition 2

Recall the definitions of $push(L, w)$, $rest(L, w)$, and $offset(L, w)$ for a language nonempty language L and a word w in three cases with $s = \max_{\leq_{\text{prefix}}} (Prefix(w) \cap Shovel(L))$:

- 0° $L = \{\varepsilon\}$. Then $push(L, w) = w$, $rest(L, w) = \varepsilon$ and $offset(L, w) = \varepsilon$.
- 1° L is periodic and $L \neq \{\varepsilon\}$. Then there exists $k \geq 0$ and $o \in Prefix(Period(L))$ such that $s = Period(L)^k \cdot o$ and $push(L, w) = s$, $rest(L, w) = r$ and $offset(L, w) = o$.
- 2° L is non-periodic. Then $push(L, w) = s$, $rest(L, w) = r$ and $offset(L, w) = s$.

The fact that the set of offsets $O = \{offset(L, w) \mid w \in \Delta^*\}$ is finite follows directly from the definitions in all three cases. We just show that this set can be constructed in time doubly-exponential in the size of a context-free grammar G defining L . The algorithms outlined below allow also to classify the case the language L belongs to.

For the case 0° we note that the condition $L = \{\varepsilon\}$ can be easily checked on G by testing that every (reachable) nonterminal of G does not produce the non-empty word. This can be done with a simple closure algorithm working in time polynomial in the size of G .

We handle the remaining two cases together and let $w_{\min}$ a shortest word in L . First, we observe that L is periodic if and only if its primitive period $Period(L)$ is also the primitive period of $w_{\min}$. Now, let v be the shortest prefix v of $w_{\min}$ such that $w_{\min} = v^k$ for some $k > 0$ (possibly equal to $w_{\min}$). One can easily see that $Period(\{w_{\min}\})$ is the shortest prefix v of $w_{\min}$ such that $w_{\min} = v^k$ for some $k > 0$. Consequently, L is periodic if and only if $L \subseteq v^*$.

We observe that $w_{\min}$ may be of length exponential in the size of G (cf. Example 4), and that may be also the length of v . Testing the inclusion $L \subseteq v^*$, when G is a CFG defining L , can be done using standard automata techniques: we construct a push-down automaton A_G defining L , a dFA A_v defining v^* , and its complement A_v^c defining $\Delta^* \setminus v^*$. Now, we take the product $P = A_G \times A_v^c$ and test it for emptiness. Clearly, $L \subseteq v^*$ iff P defines an empty language. As for complexity, we note that the size of A_G is $poly(|G|)$, the sizes of A_v and A_v^c are $poly(|v|) = exp(|G|)$, and thus the product automaton P is of size $exp(|G|)$.

If L is periodic, then $O = Prefix(Period(L)) = Prefix(v)$ and its size is single-exponential in the size of G . If L is not periodic, then the test described above fails i.e., P is nonempty and accepts a non-empty word w_0 . Note that because

P is a push-down automaton whose size is $\exp(|G|)$ the shortest word recognized by P may be of size doubly-exponential in the size of G . We claim that if a word can be pushed through L , then it cannot be longer than w_0 . The proof is combinatorial and we omit it here.

A.3 Definition of the pregroup of words over Δ

A pregroup is a set G with two operations: the unary inverse operator $\square^{-1} : G \rightarrow G$ and a partially defined binary operation $\square \cdot \square$ that takes a pair of elements of G and returns an element of G (or is undefined). Additionally, the following conditions need to be satisfied (for $z, z', z'' \in G$):

- *Associativity*: if both $z \cdot z'$ and $z' \cdot z''$ are defined, then both $(z \cdot z') \cdot z''$ and $z \cdot (z' \cdot z'')$ are defined and equal.
- *Inverse*: $z \cdot z^{-1}$ and $z^{-1} \cdot z$ are defined.
- *Annihilation*: if $z \cdot z'$ is defined, then $z \cdot z' \cdot (z')^{-1} = z$ and $z^{-1} \cdot z \cdot z' = z'$.

The *pregroup of words over Δ* is the set $\mathbb{G}_\Delta = \Delta^* \cup \{w^{-1} \mid w \in \Delta^+\}$, where w^{-1} represents the inverse of a nonempty word w . The inverse operator is defined as: $(w)^{-1} = w^{-1}$, $\varepsilon^{-1} = \varepsilon$, and $(w^{-1})^{-1} = w$, for $w \in \Delta$. The concatenation operator is an extension of the standard word concatenation to inverse elements as follows: $w^{-1} \cdot u^{-1} = (u \cdot w)^{-1}$, $w^{-1} \cdot u = v$ if $u = w \cdot v$, $w^{-1} \cdot u = v^{-1}$ if $w = u \cdot v$, $u \cdot w^{-1} = v$ if $u = v \cdot w$, and $u \cdot w^{-1} = v^{-1}$ if $w = u \cdot v$; the result is undefined in all other cases where one argument of concatenation is an inverse.

A.4 Proof of Proposition 3

To prove Proposition 3 we distinguish 3 cases: L is trivial or L is periodic and non trivial or L is non periodic.

If L is trivial, then $push(L, z) = rest(L, z) = offset(L, z) = \varepsilon$ and the proposition holds.

Let us assume that L is periodic and non trivial. Let us first consider the case where $z = w \in \Delta^*$. Let $p = Period(L)$. We have:

$$\begin{aligned} Shovel(L, w) &= p^* \cdot Prefix(p) \\ push(L, w) &= p^i \cdot o \text{ for some } i > 0 \text{ and } o \in Prefix(p) \\ offset(L, w) &= (p^i \cdot o) \mod p = o \end{aligned}$$

Let us recall that $push(L, w)$ which is also a prefix of w . Note that for any prefix u of a word v we have $u \cdot (u^{-1} \cdot v) = v$. Any word u in L is of the form p^k for some $k > 0$, and we have

$$\begin{aligned} &push(L, w) \cdot (offset(L, w)^{-1} \cdot u \cdot offset(L, w)) \cdot rest(L, w) \\ &= p^i \cdot o \cdot (o^{-1} \cdot p^k \cdot o) \cdot (p^i \cdot o)^{-1} \cdot w \\ &= p^k \cdot p^i \cdot o \cdot (p^i \cdot o)^{-1} \cdot w \\ &= p^k \cdot w = u \cdot w \end{aligned}$$

Let us now consider the case where $z = w^{-1}$ and $w \in \Delta^*$. Let $p^{\text{rev}} = \text{Period}(L^{\text{rev}})$. We have:

$$\begin{aligned} \text{Shovel}(L^{\text{rev}}, w^{-1}) &= (p^{\text{rev}})^* \cdot \text{Prefix}(p^{\text{rev}}) \\ \text{push}(L^{\text{rev}}, w^{-1}) &= (p^{\text{rev}})^i \cdot o \text{ for some } i > 0 \text{ and } o \in \text{Prefix}(p^{\text{rev}}) \\ \text{offset}(L^{\text{rev}}, w^{-1}) &= ((p^{\text{rev}})^i \cdot o) \bmod (p^{\text{rev}}) = o. \end{aligned}$$

Let $u = p^k \in L$

$$\begin{aligned} &\text{push}(L, w^{-1}) \cdot (\text{offset}(L, w^{-1})^{-1} \cdot u \cdot \text{offset}(L, w^{-1})) \cdot \text{rest}(L, w^{-1}) \\ &= (\text{push}(L^{\text{rev}}, w^{\text{rev}})^{\text{rev}})^{-1} \cdot (\text{offset}(L^{\text{rev}}, w^{\text{rev}})^{\text{rev}} \cdot u \cdot (\text{offset}(L^{\text{rev}}, w^{\text{rev}})^{\text{rev}})^{-1}) \\ &\quad \cdot (\text{rest}(L^{\text{rev}}, w^{\text{rev}})^{\text{rev}})^{-1} \\ &= (o^{\text{rev}} \cdot p^i)^{-1} \cdot (o^{\text{rev}} \cdot p^k \cdot (o^{\text{rev}})^{-1}) \cdot \left((((p^{\text{rev}})^i \cdot o)^{-1} \cdot w^{\text{rev}})^{\text{rev}} \right)^{-1} \end{aligned}$$

Now, recall that for any u and v in Δ^* , $(u^{-1} \cdot v)^{\text{rev}} = v^{\text{rev}} \cdot (u^{\text{rev}})^{-1}$ and $(u \cdot v^{-1})^{-1} = v \cdot u^{-1}$. Thus we have

$$\begin{aligned} \left((((p^{\text{rev}})^i \cdot o)^{-1} \cdot w^{\text{rev}})^{\text{rev}} \right)^{-1} &= \left((w^{\text{rev}})^{\text{rev}} \cdot (((p^{\text{rev}})^i \cdot o)^{\text{rev}})^{-1} \right)^{-1} \\ &= ((p^{\text{rev}})^i \cdot o)^{\text{rev}} \cdot w^{-1} \\ &= o^{\text{rev}} \cdot ((p^{\text{rev}})^i)^{\text{rev}} \cdot w^{-1} \\ &= o^{\text{rev}} \cdot p^i \cdot w^{-1} \end{aligned}$$

Hence,

$$\begin{aligned} &\text{push}(L, w^{-1}) \cdot (\text{offset}(L, w^{-1})^{-1} \cdot u \cdot \text{offset}(L, w^{-1})) \cdot \text{rest}(L, w^{-1}) \\ &= (o^{\text{rev}} \cdot p^i)^{-1} \cdot (o^{\text{rev}} \cdot p^k \cdot (o^{\text{rev}})^{-1}) \cdot o^{\text{rev}} \cdot p^i \cdot w^{-1} \\ &= (p^{k-i} \cdot (o^{\text{rev}})^{-1}) \cdot o^{\text{rev}} \cdot p^i \cdot w^{-1} \end{aligned}$$

Since, o^{rev} is a suffix of p , this gives $p^k \cdot w^{-1} = u \cdot w^{-1}$ and the lemma holds.

Consider L is non periodic and the case where $z = w \in \Delta^*$. Then there exists some word s such that $\text{Shovel}(L) = \text{Prefix}(s)$. In this case, $\text{push}(L, w) = \text{lcp}(\{w, s\})$ and therefore $\text{push}(L, w)$ is a prefix of s . We obtain that $\text{offset}(L, w) = \text{push}(L, w)$ and thus for all $u \in L$ we have $\text{push}(L, w) \cdot (\text{offset}(L, w))^{-1} \cdot u \cdot \text{offset}(L, w) \cdot \text{rest}(L, w) = \text{push}(L, w) \cdot (\text{push}(L, w))^{-1} \cdot u \cdot \text{push}(L, w) \cdot \text{push}(L, w)^{-1} \cdot w$. Recall that $\text{push}(L, w)$ is in $\text{Shovel}(L)$ and by the definition of Shovel we have that $\text{push}(L, w)$ is a prefix of $u \cdot \text{push}(L, w)$. Also, $\text{push}(L, w)$ is a prefix of w . Thus $\text{push}(L, w) \cdot (\text{push}(L, w))^{-1} \cdot u \cdot \text{push}(L, w) \cdot \text{push}(L, w)^{-1} \cdot w = u \cdot w$.

The case where $z = w^{-1}$ and $w \in \Delta^*$ is similar. There exists some word s such that $\text{Shovel}(L^{\text{rev}}) = \text{Prefix}(s)$, and $\text{push}(L^{\text{rev}}, w^{\text{rev}}) = \text{lcp}(\{w, s\}) = \text{offset}(L^{\text{rev}}, w^{\text{rev}})$. Thus, for all $u \in L$ we have

$$\begin{aligned}
& push(L, w^{-1}) \cdot (offset(L, w^{-1})^{-1} \cdot u \cdot offset(L, w^{-1})) \cdot rest(L, w^{-1}) \\
&= (push(L^{\text{rev}}, w^{\text{rev}})^{\text{rev}})^{-1} \cdot (push(L^{\text{rev}}, w^{\text{rev}})^{\text{rev}} \cdot u \cdot (push(L^{\text{rev}}, w^{\text{rev}})^{\text{rev}})^{-1}) \\
&\quad \cdot ((push(L^{\text{rev}}, w^{\text{rev}})^{-1} \cdot w^{\text{rev}})^{\text{rev}})^{-1} \\
&= (u \cdot (push(L^{\text{rev}}, w^{\text{rev}})^{\text{rev}})^{-1}) \cdot push(L^{\text{rev}}, w^{\text{rev}})^{\text{rev}} \cdot w^{-1} \\
&= u \cdot w^{-1}
\end{aligned}$$

□

A.5 Proof of Theorem 2

For a language L we define $Offsets(L) = \{offset(L, w) \mid w \in \Delta^*\}$.

Proposition 6. *For any reduced language L we have*

1. $\forall w \in Offsets(L)$ the language $w^{-1} \cdot L \cdot w$ is reduced,
2. $\forall w \in Offsets(L^{\text{rev}})$ the language $w \cdot L \cdot w^{-1}$ is reduced as well.

Proof. We only prove 1 since 2 is a consequence of 1. If $\varepsilon \in L$ then $\varepsilon \in w^{-1} \cdot L \cdot w$ and the proposition is trivial. Otherwise, there exists two words u and v that differ on the first letter. But in this case $Offsets(L)$ is reduced to ε and the proposition is also trivially true. □

Lemma 2. *For every $q \in Q$, every $z \in Offsets(q)$, and every $t \in dom(q)$*

$$T_{\langle q, z \rangle}(t) = z^{-1} \cdot Left(q)^{-1} \cdot T_q(t) \cdot Right(q)^{-1} \cdot z \quad (2)$$

Proof. Rules are built using the following algorithm.

- 1: $z_k := Right(q_k) \cdot u_k \cdot Right(q)^{-1} \cdot z$
- 2: **for** $i := k, \dots, 1$ **do**
- 3: $u'_i := rest(q_i, z_i)$
- 4: $q'_i := \langle q_i, offset(q_i, z_i) \rangle$
- 5: $z_{i-1} := Right(q_{i-1}) \cdot u_{i-1} \cdot Left(q_i) \cdot push(q_i, z_i)$
- 6: $u'_0 := z^{-1} \cdot Left(q)^{-1} \cdot z_0$

We prove the lemma by induction on the structure of terms.

For the base case, we consider constants and therefore rules of the form $\delta'(\langle q, z \rangle, a) = u$ and $T_q(a) = u$. The algorithm just realizes the affectation⁴:

$$u' = z^{-1} \cdot Left(q)^{-1} \cdot u \cdot Right(q)^{-1} \cdot z$$

Therefore $T_{\langle q, z \rangle}(a) = z^{-1} \cdot Left(q)^{-1} \cdot u \cdot Right(q)^{-1} \cdot z$ and the lemma holds.

---

⁴ maybe make this case more clear in section Normalisation

Let us now consider a term $t = f(t_1, \dots, t_k)$, a word z and a rule $\delta(q, f) = u_0 \cdot q_1 \cdots q_k \cdot u_k$. Let us assume as our induction hypothesis that Equation (2) holds for each subterm in $t_1, \dots, t_k$. We need to prove that:

$$u'_0 \cdot T_{q'_1}(t_1) \cdot u'_1 \cdots T_{q'_k}(t_k) \cdot u'_k = z^{-1} \cdot \text{Left}(q)^{-1} \cdot u_0 \cdot T_{q_1}(t_1) \cdots T_{q_k}(t_k) \cdot u_k \cdot \text{Right}(q)^{-1} \cdot z$$

Fact 4 $\text{push}(q_i, z_i) \cdot T_{q'_i}(t_i) \cdot \text{rest}(q_i, z_i) = \text{Left}(q_i)^{-1} \cdot T_{q_i}(t_i) \cdot \text{Right}(q_i)^{-1} \cdot z_i$

Proof. By induction hypothesis, we develop $T_{q'_i}(t_i) = \text{offset}(q_i, z_i)^{-1} \cdot \text{Left}(q_i)^{-1} \cdot T_{q_i}(t_i) \cdot \text{Right}(q_i)^{-1} \cdot \text{offset}(q_i, z_i)$. We observe that $\text{Left}(q_i)^{-1} \cdot T_{q_i}(t_i) \cdot \text{Right}(q_i)^{-1}$ is a word of $\hat{L}_{q_i}$, thus we conclude using Proposition 3. $\square$

We prove the following invariant of the algorithm for every $1 \leq i \leq k$:

$$z_{i-1} \cdot T_{q'_i}(t_i) \cdot u'_i \cdots T_{q'_k}(t_k) \cdot u'_k = \text{Right}(q_{i-1}) \cdot u_{i-1} \cdot T_{q_i}(t_i) \cdot u_i \cdots T_{q_k}(t_k) \cdot u_k \cdot \text{Right}(q)^{-1} \cdot z \quad (3)$$

We proceed by induction on i from k to 1. For the base case $i = k$ we have:

$$\begin{aligned} z_{k-1} \cdot T_{q'_k}(t_k) \cdot u'_k &= \text{Right}(q_{k-1}) \cdot u_{k-1} \cdot \text{Left}(q_k) \cdot \text{push}(q_k, z_k) \cdot T_{q'_k}(t_k) \cdot \text{rest}(q_k, z_k) \\ &= \text{Right}(q_{k-1}) \cdot u_{k-1} \cdot \text{Left}(q_k) \cdot \text{Left}(q_k)^{-1} \cdot T_{q_k}(t_k) \cdot \text{Right}(q_k)^{-1} \cdot u_k \cdot \text{Right}(q)^{-1} \cdot z \\ &= \text{Right}(q_{k-1}) \cdot u_{k-1} \cdot T_{q_k}(t_k) \cdot u_k \cdot \text{Right}(q)^{-1} \cdot z \end{aligned}$$

Thus the invariant (3) holds for the base case. Let us now consider it holds for some $1 \leq i \leq k$. For $i - 1$ we have

$$\begin{aligned} z_{i-1} \cdot T_{q'_i}(t_i) \cdot u'_i \cdots T_{q'_k}(t_k) \cdot u'_k &= \text{Right}(q_{i-1}) \cdot u_{i-1} \cdot \text{Left}(q_i) \cdot \text{push}(q_i, z_i) \cdot T_{q'_i}(t_i) \cdot \text{rest}(q_i, z_i) \\ &\quad \cdots T_{q'_k}(t_k) \cdot u'_k \\ &= \text{Right}(q_{i-1}) \cdot u_{i-1} \cdot \text{Left}(q_i) \cdot \text{Left}(q_i)^{-1} \cdot T_{q_i}(t_i) \cdot \text{Right}(q_i)^{-1} \cdot z_i \\ &\quad \cdots T_{q'_k}(t_k) \cdot u'_k \\ &= \text{Right}(q_{i-1}) \cdot u_{i-1} \cdot T_{q_i}(t_i) \cdot \text{Right}(q_i)^{-1} \cdot \text{Right}(q_i) \cdot u_i \\ &\quad \cdots T_{q'_k}(t_k) \cdot u_k \cdot \text{Right}(q)^{-1} \cdot z \\ &= \text{Right}(q_{i-1}) \cdot u_{i-1} \cdot T_{q_i}(t_i) \cdot u_i \cdots T_{q_k}(t_k) \cdot u_k \cdot \text{Right}(q)^{-1} \cdot z \end{aligned}$$

This proves the invariant. Now, since $u'_0 = z^{-1} \cdot \text{Left}(q)^{-1} \cdot z_0$, we obtain that

$$u'_0 \cdot T_{q'_1}(t_1) \cdot u'_1 \cdots T_{q'_k}(t_k) \cdot u'_k = z^{-1} \cdot \text{Left}(q)^{-1} \cdot u_0 \cdot T_{q_1}(t_1) \cdots T_{q_k}(t_k) \cdot u_k \cdot \text{Right}(q)^{-1} \cdot z$$

$\square$

Lemma 3. $(\mathbf{E}_1)$ is satisfied in M' .

Proof. To prove $(\mathbf{E}_1)$, we need to prove that $L_{\langle q, z \rangle}$ is reduced. By previous lemma, we have for every $q \in Q$, every $z \in \text{Offsets}(q)$, and every $t \in \text{dom}(q)$

$$L_{\langle q, z \rangle} = z^{-1} \cdot \text{Left}(q)^{-1} \cdot L_q \cdot \text{Right}(q)^{-1} \cdot z$$

Since $z \in \text{Offsets}(q)$, then $z = w \in \text{Shovel}(\hat{L}_q)$ or $z = w^{-1}$ with $w \in \text{Shovel}(\hat{L}_q^{\text{rev}})$. The languages $\text{Left}(q)^{-1} \cdot L_q \cdot \text{Right}(q)^{-1} = \hat{L}_q$ are reduced. Therefore using Proposition 6, we obtain that $(\mathbf{E}_1)$ is satisfied. $\square$

Lemma 4. Let L be a reduced language. For any $w \in \text{Prefix}(\text{Left}(L \cdot L'))$ we have $w \in \text{Prefix}(L')$, $\text{push}(L, w) = w$ and $\text{rest}(L, w) = \varepsilon$.

Proof. If $w = \varepsilon$ then the lemma is trivial. Otherwise, we know that $\varepsilon \in L$ otherwise L would not be reduced. Hence, w is a common prefix of all words in L' . Therefore, $L \cdot w$ has also w as a common prefix and by definition $w \in \text{Shovel}(L)$. It follows that $\text{push}(L, w) = w$ and $\text{rest}(L, w) = \varepsilon$. $\square$

Lemma 5. For any rule $\delta'(\langle q, z \rangle, f) = u'_0 \cdot q'_1 \cdot u'_1 \cdots u'_{k-1} \cdot q'_k \cdot u'_k$ of M' , for every $0 \leq i \leq k$ $u'_i \in \Delta^*$.

Proof. We consider a rule $\delta'(\langle q, z \rangle, f) = u'_0 \cdot q'_1 \cdot u'_1 \cdots u'_{k-1} \cdot q'_k \cdot u'_k$ of M' and $w \in \text{Offsets}(\hat{L}_q^{\text{rev}}) \cup \text{Offsets}(\hat{L}_q)$. We essentially have to show that for all z_i computed by the normalisation algorithm, if $z_i = w_i^{-1}$ with $w_i \in \Delta^*$, then $u_i = \text{rest}(\hat{L}_{q_i}^{\text{rev}}, w_i^{\text{rev}}) = \varepsilon$.

Let us consider L_i defined by:

$$L_i = z^{-1} \cdot \text{Left}(q)^{-1} \cdot u_0 \cdot \text{Left}(L_{q_1}) \cdot \hat{L}_{q_1} \cdot \text{Right}(L_{q_1}) \cdot u_1 \cdots u_{i-1} \cdot \text{Left}(L_{q_i})$$

Observe that if w_i is a suffix of $\text{Right}(L_i \cdot \hat{L}_{q_i})$ then w_i^{rev} is a prefix of $\text{Left}(\hat{L}_{q_i}^{\text{rev}} \cdot L_i^{\text{rev}})$. The language $\hat{L}_{q_i}^{\text{rev}}$ is reduced and by Lemma 4, $\text{rest}(L_{q_i}^{\text{rev}}, w_i^{\text{rev}}) = \varepsilon$ and $\text{push}(L_{q_i}^{\text{rev}}, w_i^{\text{rev}}) = w_i^{\text{rev}}$. So it suffices to prove that for every i , w_i is a suffix of $\text{Right}(L_i \cdot \hat{L}_{q_i})$.

We proceed by induction on i from k to 1. For the base case, $z_k = w_k^{-1} = \text{Right}(q_k) \cdot u_k \cdot \text{Right}(q)^{-1} \cdot z$. We have $L_{\langle q, z \rangle} = L_k \cdot \hat{L}_{q_k} \cdot w_k^{-1}$ and according to Lemma 3, $L_{\langle q, z \rangle}$ is reduced. Therefore $w_k = \text{Right}(L_k \cdot \hat{L}_{q_k})$.

For the induction case, we have by induction hypothesis that w_i is a suffix of $\text{Right}(L_i \cdot \hat{L}_{q_i})$, and from Lemma 4, w_i is a suffix of $\text{Right}(L_i)$ and $\text{push}(L_{q_i}^{\text{rev}}, w_i^{\text{rev}}) = w_i^{\text{rev}}$. From the algorithm $z_{i-1} = \text{Right}(q_{i-1}) \cdot u_{i-1} \cdot \text{Left}(q_i) \cdot \text{push}(q_i, z_i)$. Hence, $z_{i-1} = \text{Right}(q_{i-1}) \cdot u_{i-1} \cdot \text{Left}(q_i) \cdot w_i^{-1}$. If $z_{i-1} \in \Delta^*$ we are done. Otherwise $z_{i-1} = w_{i-1}^{-1}$ with $w_{i-1} \in \Delta^*$. Since w_i is a suffix of $\text{Right}(L_i) = \text{Right}(L_{i-1} \cdot \text{Right}(q_{i-1}) \cdot u_{i-1} \cdot \text{Left}(q_i))$, then z_{i-1} is a suffix of $\text{Right}(L_{i-1} \cdot \hat{L}_{q_{i-1}})$. $\square$

Lemma 6. Given a reduced language $L \subseteq \Delta^*$ and a word $z \in \mathbb{G}_\Delta$

$$\text{lcp}((\text{offset}(L, z)^{-1} \cdot L \cdot \text{offset}(L, z)) \cdot \text{rest}(L, z)) = \varepsilon$$

Proof. Let us consider the case where $z = w \in \Delta^*$, the other case being symmetric. Let us remark using Lemma 6, that $L' = \text{offset}(L, z)^{-1} \cdot L \cdot \text{offset}(L, z)$ is reduced. If L' does not contain ε then the Lemma is trivial. So let us consider the case where $\varepsilon \in L'$.

We prove the lemma by contradiction and consider $u \neq \varepsilon$ such that $u = \text{lcp}((\text{offset}(L, w)^{-1} \cdot L \cdot \text{offset}(L, w)) \cdot \text{rest}(L, w))$. Let $v = \text{push}(L, w) \cdot u$. We prove that v is a prefix of w . Indeed, since $\varepsilon \in L'$, u is a prefix of $\text{rest}(L, w) = \text{push}(L, w)^{-1} \cdot w$. Using Proposition 3, we have that v is a common prefix of $L \cdot w$. Therefore, v belongs to $\text{Shovel}(L)$. This contradicts the fact that is the maximum prefix of w that belongs to $\text{Shovel}(L)$. $\square$

Lemma 7. $(\mathbf{E}_2)$ is satisfied in M' .

Proof. According to Lemma 3, each $L_{\langle q, z \rangle}$ is reduced. For the init rule, $(\mathbf{E}_2)$ is direct consequence of Lemma 6.

Consider a rule $\delta'(f, q') = u'_0 \cdot q_1 \cdots q_k \cdot u'_k$. We know from Lemma 5 that each u'_i is either ε or a word in Δ^* . Moreover, according to line 3 of the algorithm, for for each $i > 0$, $u'_i = \text{rest}(q_i, z_i)$. Therefore, by Lemma 6, $\text{lcp}(L_{q'_i} \cdot u'_i) = \varepsilon$. Note that if two languages L and L' are such that $\text{lcp}(L) = \text{lcp}(L') = \varepsilon$ then $\text{lcp}(L \cdot L') = \varepsilon$. So, $\text{lcp}(L_{q'_i} \cdot u'_i \cdots L_{q'_k} \cdot u'_k) = \varepsilon$ and $(\mathbf{E}_2)$ is thus satisfied for every rule. $\square$

Lemma 8. M equivalent with M'

Proof. We prove that for all t , $u_0 \cdot T_{q_0}(t) \cdot u_1 = u'_0 \cdot T_{q'_0}(t) \cdot u'_1$. To do that, we first inspect the initial rule. Using Lemma 2, we have, with $v = \text{Right}(q_0) \cdot u_1$:

$$\begin{aligned} u'_0 \cdot T_{q'_0}(t) \cdot u'_1 &= u_0 \cdot \text{Left}(q_0) \cdot \text{push}(q_0, v) \\ &\quad \cdot \text{offset}(q_0, v)^{-1} \cdot \text{Left}(q_0)^{-1} \cdot T_{q_0}(t) \cdot \text{Right}(q_0)^{-1} \cdot \text{offset}(q_0, v) \\ &\quad \cdot \text{rest}(q_0, v) \\ &= u_0 \cdot \text{Left}(q_0) \cdot \text{Left}(q_0)^{-1} \cdot T_{q_0}(t) \cdot \text{Right}(q_0)^{-1} \cdot v \\ &\quad \text{(using Prop. 3)} \\ &= u_0 \cdot T_{q_0}(t) \cdot u_1 \end{aligned}$$

$\square$

A.6 Proof of Lemma 1

First, note that for every $t \in \text{dom}(T)$ we have that

$$T_{M_1}(t) = u_0 \cdot T_{q_0}(t) \cdot u_1 = u'_0 \cdot T_{q'_0}(t) \cdot u'_1 = T_{M_2}(t).$$

Therefore, u_0 is a prefix of u'_0 or u'_0 is a prefix of u_0 . W.l.o.g. we assume that u_0 is a prefix of u'_0 i.e., $u'_0 = u_0 \cdot v$. Thus, we obtain $u_0 \cdot T_{q_0}(t) \cdot u_1 = u_0 \cdot v \cdot T_{q'_0}(t) \cdot u'_1$, and $T_{q_0}(t) \cdot u_1 = v \cdot T_{q'_0}(t) \cdot u'_1$. Consequently, v is also a prefix of $L_{q_0} \cdot u_1$ which by $(\mathbf{E}_2)$ implies that $v = \varepsilon$.

Also, we note that u_1 is a suffix of u'_1 or u'_1 is a suffix of u_1 . Again, w.l.o.g. we assume that u_1 is a suffix of u'_1 i.e., $u'_1 = v \cdot u_1$. We get that $u_0 \cdot T_{q_0}(t) \cdot u_1 = u_0 \cdot T_{q'_0}(t) \cdot v \cdot u_1$ and $T_{q_0}(t) = T_{q'_0}(t) \cdot v$. This implies that v is a suffix of L_{q_0} but by **(E₁)** we get that v must be an empty word, and so $u_1 = u'_1$ and $T_{q_0} = T_{q'_0}$.

Next, we prove the following inductive argument. If for a path $\pi \in \text{paths}(\text{dom}(T))$, $q = \delta(q_0, \pi)$, and $q' = \delta'(q'_0, \pi)$ we have $T_q = T_{q'}$, then

1. $\delta(q, f)$ is defined if and only if $\delta'(q', f)$ is, for every $f \in \Sigma$, and
2. if $\delta(q, f) = u_0 \cdot q_1 \cdot u_1 \cdot \dots \cdot q_k \cdot u_k$ and $\delta'(q', f) = u'_0 \cdot q'_1 \cdot u'_1 \cdot \dots \cdot q'_k \cdot u'_k$, then $u_i = u'_i$ for $0 \leq i \leq k$ and $T_{q_i} = T_{q'_i}$ for every $1 \leq i \leq k$.

The first condition follows trivially from equality of the domains of T_q and $T_{q'}$. We prove the second condition with induction over i : we assume that $u_j = u'_j$ and $T_{q_j} = T_{q'_j}$ for every $0 \leq j < i \leq k$ and show that $u_i = u'_i$ and then $T_{q_{i+1}} = T_{q'_{i+1}}$. The arguments we use are analogous to those used to show equalities for the initial rules.

Since $T_q = T_{q'}$, for every $t = f(t_1, \dots, t_k) \in \text{dom}(T_q)$

$$T_q(t) = u_0 \cdot T_{q_1}(t_1) \cdot u_1 \cdot \dots \cdot T_{q_k}(t_k) \cdot u_k = u'_0 \cdot T_{q'_1}(t_1) \cdot u'_1 \cdot \dots \cdot T_{q'_k}(t_k) \cdot u'_k = T_{q'}(t),$$

which by IH give us the following equality

$$u_i \cdot T_{q_{i+1}}(t_{i+1}) \cdot u_{i+1} \cdot \dots \cdot T_{q_k}(t_k) \cdot u_k = u'_i \cdot T_{q'_{i+1}}(t_{i+1}) \cdot u'_{i+1} \cdot \dots \cdot T_{q'_k}(t_k) \cdot u'_k.$$

Because of this equality we have that u_i is a prefix of u'_i or u'_i is a prefix of u_i . W.l.o.g. we assume that u_i is a prefix of u'_i i.e., $u'_i = u_i \cdot v$. Then, v is also a prefix of $L_{q_{i+1}} \cdot u_{i+1} \cdot \dots \cdot L_{q_k} \cdot u_k$ and by **(E₂)** $v = \varepsilon$. Thus, $u_i = u'_i$.

Now, we fix $t = f(t_1, \dots, t_k) \in \text{dom}(T_q)$ and let $w = u_{i+1} \cdot T_{q_{i+2}}(t_{i+2}) \cdot \dots \cdot T_{q_k}(t_k) \cdot u_k$ and $w' = u'_{i+1} \cdot T_{q'_{i+2}}(t_{i+2}) \cdot \dots \cdot T_{q'_k}(t_k) \cdot u'_k$. Note that w is a suffix of w' or w' is a suffix of w . W.l.o.g. we assume that w is a suffix of w' i.e., $w' = v \cdot w$. We observe that for every $t'_{i+1} \in (f, i+1)^{-1} \text{dom}(T_q) = (f, i+1)^{-1} \text{dom}(T_{q'})$ we have

$$T_{q_{i+1}}(t'_{i+1}) \cdot w = T_{q'_{i+1}}(t'_{i+1}) \cdot w'$$

and then

$$T_{q_{i+1}}(t'_{i+1}) = T_{q'_{i+1}}(t'_{i+1}) \cdot v,$$

which implies that v is a suffix of $L_{q_{i+1}}$ (since $(f, i+1)^{-1} \text{dom}(T_q) = \text{dom}(T_{q_{i+1}})$). By **(E₁)** we have, however, that $v = \varepsilon$. Thus, $T_{q_{i+1}} = T_{q'_{i+1}}$. $\square$

A.7 Proof of Theorem 4

We show NP-completeness of the following decision problem.

Problem: MINIMIZE_{STW}

Input: An STW M and a natural number K .

Question: Is there an STW M' equivalent to M and of size $\leq K$?

We reduce the following variant of SAT to MINIMIZE_{STW}.

Problem: SAT₁

Input: $\varphi \in \text{3CNF}$, i.e. a conjunction of clauses, where each clause is a disjunction of 3 literals (a literal is a variable or its negation).

Question: Does there exists a valuation satisfying φ and such that in every clause of φ exactly one literal is true?

Take any 3CNF formula $\varphi = c_1 \wedge \dots \wedge c_k$ over the set of Boolean variables $p_1, \dots, p_n$. For technical reasons, we assume that $k \geq 1$ and that no two clauses are the same, i.e. $c_i \not\equiv c_j$ for $i \neq j$. We construct a STW M over the input alphabet $\Sigma = \Sigma^{(3)} \cup \Sigma^{(2)} \cup \Sigma^{(0)}$

$$\Sigma^{(3)} = \{f_j\}_{j=1}^k, \quad \Sigma^{(2)} = \{g_i^1, g_i^2, g_i^3\}_{i=1}^n, \quad \Sigma^{(0)} = \{o_i, z_i\}_{i=1}^n \cup \{h\},$$

and the output alphabet $\Delta = \{s\}$. The constructed transducer accepts on the input trees with root f that represent the clauses of the formula φ . Additionally, the transducer accepts trees with root g_i^1, g_i^2, g_i^3 that represent dummy clauses (i.e. $p_i \vee \neg p_i$ independently three times). For all of those trees the transducer outputs s . Finally, to prevent retracting the output string s to the initial rule, the transducer accepts also a tree h for which it output the empty string. Formally, the constructed transducer M_φ has its initial rule q_0 and the following transitions:

1. $\delta(q_0, f_j) = s \cdot q_{\ell_{j,1}} \cdot q_{\ell_{j,2}} \cdot q_{\ell_{j,3}}$ for every $1 \leq j \leq k$ with $c_j = \ell_{j,1} \wedge \ell_{j,2} \wedge \ell_{j,3}$;
2. $\delta(q_0, g_i^j) = s \cdot q_{p_i} \cdot q_{\neg p_i}$ for $1 \leq i \leq n$ and $1 \leq j \leq 3$;
3. $\delta(q_{p_i}, o_i) = \varepsilon$ and $\delta(q_{\neg p_i}, z_i) = \varepsilon$ for $1 \leq i \leq n$;
4. $\delta(q_0, h) = \varepsilon$.

Let T_φ be the transformation defined by M_φ . We observe that the size of M_φ is $5k + 14n + 2$. Now, we claim that

$$\varphi \in \text{SAT}_1 \iff (M_\varphi, 4k + 12n + 2) \in \text{MINIMIZE}_{\text{STW}}.$$

For the *only if* part we take the valuation $V : \{p_1, \dots, p_n\} \rightarrow \{\mathbf{true}, \mathbf{false}\}$ witnessing $\varphi \in \text{SAT}_1$ and construct an STW M_V obtained from M_φ by *pushing down* the symbols s to the subtrees that correspond to the literals satisfied by V . Formally, the transducer M_V differs from M_φ only on the transition function:

1. $\delta_V(q_0, f_j) = q_{\ell_{j,1}} \cdot q_{\ell_{j,2}} \cdot q_{\ell_{j,3}}$ for every $j \in \{1, \dots, k\}$ with $c_j = \ell_{j,1} \wedge \ell_{j,2} \wedge \ell_{j,3}$;
2. $\delta_V(q_0, g_i^j) = q_{p_i} \cdot q_{\neg p_i}$ for $1 \leq i \leq n$ and $1 \leq j \leq 3$;
3. $\delta_V(q_{p_i}, o_i) = s$ and $\delta_V(q_{\neg p_i}, z_i) = \varepsilon$ for $1 \leq i \leq n$ such that $V(p_i) = \mathbf{true}$;
4. $\delta_V(q_{p_i}, o_i) = \varepsilon$ and $\delta_V(q_{\neg p_i}, z_i) = s$ for $1 \leq i \leq n$ such that $V(p_i) = \mathbf{false}$;
5. $\delta_V(q_0, h) = \varepsilon$.

M_φ and M_V are equivalent because V is the witness of $\varphi \in \text{SAT}_1$ and we observe that the size of M_V is exactly $4k + 12n + 2$.

For the *if* part, take the STW M' that is equivalent to M_φ and whose size is $4k + 6n + 2$. Let $w_0 \cdot q_0(x_0) \cdot w_1$ be its initial rule, and observe that both w_0 and w_1 are ε because $T_\varphi(h) = \varepsilon$.

We observe that M needs to have at least $1 + 2n$ states because it is the Myhill-Nerode index of the tree language $\text{dom}(T_\varphi)$. Consequently, M needs to have at least

1. 1 initial rule of size 1,
2. k transition rules of arity 3 with the symbols f_j ,
3. $3n$ transition rules of arity 2 with the symbols g_i^j ,
4. n transition rules of arity 0 with the symbols o_i ,
5. n transition rules of arity 0 with the symbols z_i , and
6. 1 transition rules of arity 0 with the symbol h .

The size of all rules without counting the output strings is (at least) $4k + 11n + 2$. Thus, it remains n *tokens* s that need to be distributed among the rules. We claim that tokens are attributed to the rules with the symbols o_i and z_i only: one token to either o_i or z_i only for every $1 \leq i \leq n$. This follows from an involved combinatorial argument omitted here. The distribution of the tokens is used to construct a valuation V witnessing $\varphi \in \mathbf{SAT}_1$.

NP-completeness follows from the fact that testing the equivalence of STWs is known to be in PTIME [16]. Thus a non deterministic Turing machine needs to guess an STW M' whose size is lower than $\min\{|M|, K\}$ and then test the equivalence of M and M' . $\square$