



UNIVERSITÀ
DEGLI STUDI
DI UDINE

Università degli studi di Udine

One-way definability of two-way word transducers

Original

Availability:

This version is available <http://hdl.handle.net/11390/1174106> since 2021-03-19T11:35:24Z

Publisher:

Published

DOI:10.23638/LMCS-14(4:22)2018

Terms of use:

The institutional repository of the University of Udine (<http://air.uniud.it>) is provided by ARIC services. The aim is to enable open access to all the world.

Publisher copyright

(Article begins on next page)

ONE-WAY DEFINABILITY OF TWO-WAY WORD TRANSDUCERS

FÉLIX BASCHENIS, OLIVIER GAUWIN, ANCA MUSCHOLL, AND GABRIELE PUPPIS

Department of Mathematics and Informatics, University Bremen

LaBRI, University of Bordeaux, CNRS

LaBRI, University of Bordeaux, CNRS

CNRS, LaBRI

ABSTRACT. Functional transductions realized by two-way transducers (or, equally, by streaming transducers or MSO transductions) are the natural and standard notion of “regular” mappings from words to words. It was shown in 2013 that it is decidable if such a transduction can be implemented by some one-way transducer, but the given algorithm has non-elementary complexity. We provide an algorithm of different flavor solving the above question, that has doubly exponential space complexity. In the special case of sweeping transducers the complexity is one exponential less. We also show how to construct an equivalent one-way transducer, whenever it exists, in doubly or triply exponential time, again depending on whether the input transducer is sweeping or two-way. In the sweeping case our construction is shown to be optimal.

1. INTRODUCTION

Since the early times of computer science, transducers have been identified as a fundamental computational model. Numerous fields of computer science are ultimately concerned with transformations, ranging from databases to image processing, and an important challenge is to perform transformations with low costs, whenever possible.

The most basic form of transformations is obtained using devices that process an input with finite memory and produce outputs during the processing. Such devices are called finite-state transducers. Word-to-word finite-state transducers were considered in very early work in formal language theory [1, 16, 32], and it was soon clear that they are much more challenging than finite-state word acceptors (classical finite-state automata). One essential difference between transducers and automata over words is that the capability to process the input in both directions strictly increases the expressive power in the case of transducers, whereas this does not for automata [30, 33]. In other words, two-way word transducers are strictly more expressive than one-way word transducers.

Key words and phrases: Regular word transductions, Two-way transducers, Sweeping transducers, One-way definability.

* The results presented here were published in [4] and [6]. The work was partially supported by the projects ExStream (ANR-13-JS02-0010) and DeLTA (ANR-16-CE40-0007).

We consider in this paper functional transducers, that compute functions from words to words. Two-way word transducers capture very nicely the notion of regularity in this setting. Regular word functions, in other words, functions computed by functional two-way transducers¹, inherit many of the characterizations and algorithmic properties of the robust class of regular languages. Engelfriet and Hoogetboom [17] showed that monadic second-order definable graph transductions, restricted to words, are equivalent to two-way transducers — this justifies the notation “regular” word functions, in the spirit of classical results in automata theory and logic by Büchi, Elgot, Rabin and others. Recently, Alur and Cerný [2] proposed an enhanced version of one-way transducers called streaming transducers, and showed that they are equivalent to the two previous models. A streaming transducer processes the input word from left to right, and stores (partial) output words in some given write-only registers, that are updated through concatenation and constant updates.

Two-way transducers raise challenging questions about resource requirements. One crucial resource is the number of times the transducer needs to re-process the input word. In particular, the case where the input can be processed in a single pass, from left to right, is very attractive as it corresponds to the setting of *streaming*, where the (possibly very large) inputs do not need to be stored in order to be processed. The *one-way definability* of a functional two-way transducer, that is, the question whether the transducer is equivalent to some one-way transducer, was considered quite recently: [19] shows that one-way definability of string transducers is a decidable property. However, the decision procedure of [19] has non-elementary complexity, which raises the question about the intrinsic complexity of this problem.

In this paper we provide an algorithm of elementary complexity solving the one-way definability problem. Our decision algorithm has single or doubly exponential space complexity, depending on whether the input transducer is sweeping (namely, it performs reversals only at the extremities of the input word) or genuinely two-way. We also describe an algorithm that constructs an equivalent one-way transducer, whenever it exists, in doubly or triply exponential time, again depending on whether the input transducer is sweeping or two-way. For the construction of an equivalent one-way transducer we obtain a doubly exponential lower bound, which is tight for sweeping transducers. Note that for the decision problem, the best lower bound known is only polynomial space [19].

Our initial interest in sweeping transducers was the fact that they provide a simpler setting for characterizing one-way definability. Later it turned out that sweeping transducers enjoy interesting and useful connections with streaming transducers: they have the same expressiveness as streaming transducers where concatenation of registers is disallowed. The connection goes even further, since the number of sweeps corresponds exactly to the number of registers [5]. The results of this paper were refined in [5], and used to determine the minimal number of registers required by functional streaming transducers without register concatenation.

Related work. Besides the papers mentioned above, there are several recent results around the expressivity and the resources of two-way transducers, or equivalently, streaming transducers. First-order definable transductions were shown to be equivalent to transductions defined by aperiodic streaming transducers [20] and to aperiodic two-way transducers [10]. An effective characterization of aperiodicity for one-way transducers was obtained in [18].

¹We know from [15, 17] that deterministic and non-deterministic functional two-way transducers have the same expressive power, though the non-deterministic variants are usually more succinct.

Register minimization for right-appending deterministic streaming transducers was shown to be decidable in [14]. An algebraic characterization of (not necessarily functional) two-way transducers over unary alphabets was provided in [12]. It was shown that in this case sweeping transducers have the same expressivity. The expressivity of non-deterministic input-unary or output-unary two-way transducers was investigated in [22].

In [35] a pumping lemma for two-way transducers is proposed, and used to investigate properties of the output languages of two-way transducers. In this paper we also rely on pumping arguments over runs of two-way transducers, but we require loops of a particular form, that allows to identify periodicities in the output.

The present paper unifies results on one-way definability obtained in [4] and [6]. Compared to the conference versions, some combinatorial proofs have been simplified, and the complexity of the procedure presented in [4] has been improved by one exponential.

Overview. Section 2 introduces basic notations for two-way and sweeping transducers, and Section 3 states the main result and provides a roadmap of the proofs. For better readability our paper is divided in two parts: in the first part we consider the easier case of sweeping transducers, and in the second part the general case. The two proofs are similar, but the general case is more involved since we need to deal with special loops (see Section 6). Since the high-level ideas of the proofs are the same and the sweeping case illustrates them in a simpler way, the proof in that setting is a preparation for the general case. Both proofs have the same structure: first we introduce some combinatorial arguments (see Section 4 for the sweeping case and Section 7 for the general case), then we provide the characterization of one-way definability (see Section 5 for the sweeping case and Section 8 for the general case). Finally, Section 9 establishes the complexity of our algorithms.

2. PRELIMINARIES

We start with some basic notations and definitions for two-way automata and transducers. We assume that every input word $w = a_1 \cdots a_n$ has two special delimiting symbols $a_1 = \vdash$ and $a_n = \dashv$ that do not occur elsewhere: $a_i \notin \{\vdash, \dashv\}$ for all $i = 2, \dots, n - 1$.

A *two-way automaton* is a tuple $A = (Q, \Sigma, \vdash, \dashv, \Delta, I, F)$, where

- Q is a finite set of states,
- Σ is a finite alphabet (including \vdash, \dashv),
- $\Delta \subseteq Q \times \Sigma \times Q \times \{\text{left}, \text{right}\}$ is a transition relation,
- $I, F \subseteq Q$ are sets of initial and final states, respectively.

By convention, left transitions on \vdash are not allowed; on the other hand, right transitions on \dashv are allowed, but, as we will see, they will necessarily appear as last transitions of successful runs. A *configuration* of A has the form $u q v$, with $uv \in \vdash \Sigma^* \dashv$ and $q \in Q$. A configuration $u q v$ represents the situation where the current state of A is q and its head reads the first symbol of v (on input uv). If $(q, a, q', \text{right}) \in \Delta$, then there is a transition from any configuration of the form $u q a v$ to the configuration $u a q' v$; we denote such a transition by $u q a v \xrightarrow{a, \text{right}} u a q' v$. Similarly, if $(q, a, q', \text{left}) \in \Delta$, then there is a transition from any configuration of the form $u b q a v$ to the configuration $u q' b a v$, denoted as $u b q a v \xrightarrow{a, \text{left}} u q' b a v$. A *run* on w is a sequence of transitions. It is *successful* if it starts in an initial configuration $q w$, with $q \in I$, and ends in a final configuration $w q'$, with $q' \in F$ — note that this latter configuration does not allow additional transitions. The *language* of A is the set of words that admit a successful run of A .

The definition of *two-way transducers* is similar to that of two-way automata, with the only difference that now there is an additional output alphabet Γ and the transition relation is a finite subset of $Q \times \Sigma \times \Gamma^* \times Q \times \{\text{left}, \text{right}\}$, which associates an output over Γ with each transition of the underlying two-way automaton. For a two-way transducer $T = (Q, \Sigma, \vdash, \dashv, \Gamma, \Delta, I, F)$, we have a transition of the form $ubqav \xrightarrow{a,d|w} u'q'v'$, outputting w , whenever $(q, a, w, q', d) \in \Delta$ and either $u' = uba \wedge v' = v$ or $u' = u \wedge v' = bav$, depending on whether $d = \text{right}$ or $d = \text{left}$. The *output* associated with a run $\rho = u_1 q_1 v_1 \xrightarrow{a_1, d_1 | w_1} \dots \xrightarrow{a_n, d_n | w_n} u_{n+1} q_{n+1} v_{n+1}$ of T is the word $\text{out}(\rho) = w_1 \cdots w_n$. A transducer T defines a relation $\mathcal{L}(T)$ consisting of all pairs (u, w) such that $w = \text{out}(\rho)$, for some successful run ρ on u . The *domain* of T , denoted $\text{dom}(T)$, is the set of input words that have a successful run. For transducers T, T' , we write $T' \subseteq T$ to mean that $\text{dom}(T') \subseteq \text{dom}(T)$ and the transductions computed by T, T' coincide on $\text{dom}(T)$.

A transducer is called *one-way* if it does not contain transition rules of the form $(q, a, w, q', \text{left})$. It is called *sweeping* if it can perform reversals only at the borders of the input word. A transducer that is equivalent to some one-way (resp. sweeping) transducer is called *one-way definable* (resp. *sweeping definable*).

The *size* of a transducer takes into account both the state space and the transition relation, and thus includes the length of the output of each transition.

Crossing sequences. The first notion that we use throughout the paper is that of crossing sequence. We follow the convenient presentation from [24], which appeals to a graphical representation of runs of a two-way transducer, where each configuration is seen as a point (location) in a two-dimensional space. Let $u = a_1 \cdots a_n$ be an input word (recall that $a_1 = \vdash$ and $a_n = \dashv$) and let ρ be a run of a two-way automaton (or transducer) on u . The *positions* of ρ are the numbers from 0 to n , corresponding to “cuts” between two consecutive letters of the input. For example, position 0 is just before the first letter a_1 , position n is just after the last letter a_n , and any other position x , with $1 \leq x < n$, is between the letters a_x and a_{x+1} . We will denote by $u[x_1, x_2]$ the factor of u between the positions x_1 and x_2 (both included).

Each configuration uqv of a two-way run ρ has a specific position associated with it. For technical reasons we need to distinguish leftward and rightward transitions. If the configuration uqv is the target of a rightward transition then the position associated with uqv is $x = |u|$. The same definition also applies when uqv is the initial configuration, for which we have $u = \varepsilon$ and $x = |u| = 0$. Otherwise, if uqv is the target of a leftward transition then the position associated with uqv is $x = |u| + 1$. Note that in both cases, the letter read by the transition leading to uqv is a_x . A *location* of ρ is any pair (x, y) , where x is the position of some configuration of ρ and y is any non-negative integer for which there are at least $y + 1$ configurations in ρ with the same position x . The second component y of a location is called *level*. For example, in Figure 1 we represent a possible run of a two-way automaton together with its locations $(0, 0)$, $(1, 0)$, $(2, 0)$, $(2, 1)$, etc. Each location is naturally associated with a configuration, and thus a state. Formally, we say that q is the *state at location* $\ell = (x, y)$ in ρ , and we denote this by writing $\rho(\ell) = q$, if the $(y + 1)$ -th configuration of ρ with position x has state q . Finally, we define the *crossing sequence* of ρ at position x as the tuple $\rho|x = (q_0, \dots, q_h)$, where the q_y 's are all the states at locations of the form (x, y) , for $y = 0, \dots, h$.

As shown in Figure 1, a two-way run can be represented as an path between locations annotated by the associated states. We observe in particular that if a location (x, y) is the

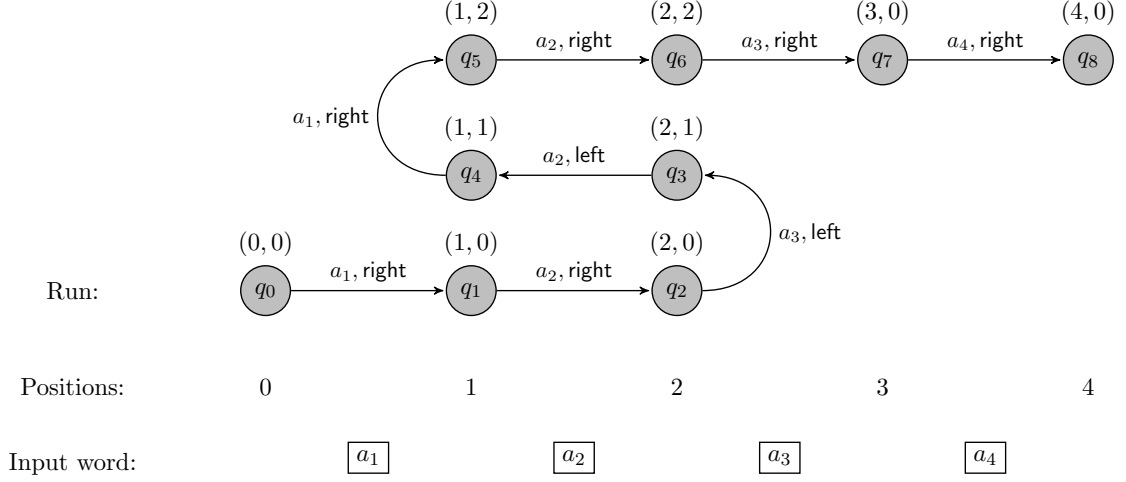


Figure 1. Graphical presentation of a run by means of crossing sequences.

target of a rightward transition, then this transition has read the symbol a_x ; similarly, if (x, y) is the target of a leftward transition, then the transition has read the symbol a_{x+1} . We also observe that, in any successful run ρ , every crossing sequence has odd length and every rightward (resp. leftward) transition reaches a location with even (resp. odd) level. In particular, we can identify four types of transitions between locations, depending on the parities of the levels (the reader may refer again to Figure 1):

$$\begin{array}{ccc}
 (x, 2y) \xrightarrow{a_{x+1}, \text{right}} (x+1, 2y') & & \begin{array}{c} (x, 2y+1) \\ \leftarrow a_{x+1}, \text{left} \\ (x, 2y) \end{array} \\
 \\
 (x, 2y+1) \xleftarrow{a_{x+1}, \text{left}} (x+1, 2y'+1) & & \begin{array}{c} a_x, \text{right} \curvearrowright \\ (x, 2y+2) \\ (x, 2y+1) \end{array}
 \end{array}$$

Hereafter, we will identify runs with the corresponding annotated paths between locations.

It is also convenient to define a total order \trianglelefteq on the locations of a run ρ by letting $\ell_1 \trianglelefteq \ell_2$ if ℓ_2 is reachable from ℓ_1 by following the path described by ρ — the order \trianglelefteq on locations is called *run order*. Given two locations $\ell_1 \trianglelefteq \ell_2$ of a run ρ , we write $\rho[\ell_1, \ell_2]$ for the factor of the run that starts in ℓ_1 and ends in ℓ_2 . Note that the latter is also a run and hence the notation $\text{out}(\rho[\ell_1, \ell_2])$ is permitted. We will often reason with factors of runs up to isomorphism, that is, modulo shifting the coordinates of their locations while preserving the parity of the levels. Of course, when the last location of (a factor of) a run ρ_1 coincides with the first location of (another factor of) a run ρ_2 , then ρ_1 and ρ_2 can be concatenated to form a longer run, denoted by $\rho_1\rho_2$. This operation can be performed even if the two locations, say (x_1, y_1) and (x_2, y_2) , are different, provided that $y_1 = y_2 \pmod 2$: in this case it suffices to shift the positions (resp. levels) of the locations of the first run ρ_1 by x_2 (resp. y_2) and, similarly, the positions (resp. levels) of the locations of the second run ρ_2 by x_1 (resp. y_1).

Intercepted factors. For simplicity, we will denote by ω the maximal position of the input word. We will consider *intervals of positions* of the form $I = [x_1, x_2]$, with $0 \leq x_1 < x_2 \leq \omega$. The *containment* relation \subseteq on intervals is defined expected, as $[x_3, x_4] \subseteq [x_1, x_2]$ if $x_1 \leq x_3 < x_4 \leq x_2$. A *factor* of a run ρ is a contiguous subsequence of ρ . A factor of ρ *intercepted* by an interval $I = [x_1, x_2]$ is a maximal factor that visits only positions $x \in I$, and never uses a left transition from position x_1 or a right transition from position x_2 . Figure 2 on the right shows the factors $\alpha, \beta, \gamma, \delta, \zeta$ intercepted by an interval I . The numbers that annotate the endpoints of the factors represent their levels.




Figure 2. Intercepted factors.

Functionality. We say that a transducer is *functional* (equivalently, one-valued, or single-valued) if for each input u , at most one output w can be produced by any possible successful run on u . Of course, every deterministic transducer is functional, while the opposite implication fails in general. To the best of our knowledge determining the precise complexity of the determinization of a two-way transducer (whenever an equivalent deterministic one-way transducer exists) is still open. From classical bounds on determinization of finite automata, we only know that the size of a determinized transducer may be exponential in the worst-case. One solution to this question, which is probably not the most efficient one, is to check one-way definability: if an equivalent one-way transducer is constructed, one can check in PTIME if it can be determinized [7, 11, 36].

The following result, proven in Section 9, is the reason to consider only functional transducers:

Proposition 9.3. The one-way definability problem for *non-functional* sweeping transducers is undecidable.

Unless otherwise stated, hereafter we tacitly assume that all transducers are functional. Note that functionality is a decidable property, as shown below. The proof of this result is similar to the decidability proof for the equivalence problem of deterministic two-way transducers [23], as it reduces the functionality problem to the reachability problem of a 1-counter automaton of exponential size. A matching PSPACE lower bound follows by a reduction of the emptiness problem for the intersection of finite-state automata [28].

Proposition 2.1. *Functionality of two-way transducers can be decided in polynomial space. This problem is PSPACE-hard already for sweeping transducers.*

A (successful) run of a two-way transducer is called *normalized* if it never visits two locations with the same position, the same state, and both either at even or at odd level. It is easy to see that if a successful run ρ of a *functional* transducer visits two locations $\ell_1 = (x, y)$ and $\ell_2 = (x, y')$ with the same state $\rho(\ell_1) = \rho(\ell_2)$ and with $y = y' \pmod{2}$, then the output produced by ρ between ℓ_1 and ℓ_2 is empty: otherwise, by repeating the non-empty factor $\rho[\ell_1, \ell_2]$, we would contradict functionality. So, by deleting the factor $\rho[\ell_1, \ell_2]$ we obtain a successful run that produces the same output. Iterating this operation leads to an equivalent, normalized run.

Normalized runs are interesting because their crossing sequences have bounded length (at most $\mathbf{H} = 2|Q| - 1$). Throughout the paper we will implicitly assume that successful runs are normalized. The latter property can be easily checked on crossing sequences.

3. ONE-WAY DEFINABILITY: OVERVIEW

In this section we state our main result, which is the existence of an elementary algorithm for checking whether a two-way transducer is equivalent to some one-way transducer. We call such transducers *one-way definable*. Before stating our result, we start with a few examples illustrating the reasons that may prevent a transducer to be one-way definable.

Example 3.1. We consider two-way transducers that accept any input u from a given regular language R and produce as output the word uu . We will argue how, depending on R , these transducers may or may not be one-way definable.

- (1) If $R = (a + b)^*$, then there is no equivalent one-way transducer, as the output language is not regular. If R is finite, however, then the transduction mapping $u \in R$ to uu can be implemented by a one-way transducer that stores the input u (this requires at least as many states as the cardinality of R), and outputs uu at the end of the computation.
- (2) A special case of transduction with finite domain is obtained from the language $R_n = \{a_0 w_0 \cdots a_{2^n-1} w_{2^n-1} : a_i \in \{a, b\}\}$, where n is a fixed natural number, the input alphabet is $\{a, b, 0, 1\}$, and each w_i is the binary encoding of the index $i = 0, \dots, 2^n - 1$ (hence $w_i \in \{0, 1\}^n$). According to Proposition 3.5 below, the transduction mapping $u \in R_n$ to uu can be implemented by a two-way transducer of size $\mathcal{O}(n^2)$, but every equivalent one-way transducer has size (at least) doubly exponential in n .
- (3) Consider now the periodic language $R = (abc)^*$. The function that maps $u \in R$ to uu can be easily implemented by a one-way transducer: it suffices to output alternatively ab, ca, bc for each input letter, while checking that the input is in R .

Example 3.2. We consider now a slightly more complicated transduction that is defined on input words of the form $u_1 \# \cdots \# u_n$, where each factor u_i is over the alphabet $\Sigma = \{a, b, c\}$. The associated output has the form $w_1 \# \cdots \# w_n$, where each w_i is either $u_i u_i$ or just u_i , depending on whether or not $u_i \in (abc)^*$ and $|u_{i+1}|$ has even length, with $u_{n+1} = \varepsilon$ by convention.

The natural way to implement this transduction is by means of a two-way transducer that performs multiple passes on the factors of the input: a first left-to-right pass is performed on $u_i \# u_{i+1}$ to produce the first copy of u_i and to check whether $u_i \in (abc)^*$ and $|u_{i+1}|$ is even; if so, a second pass on u_i is performed to produce another copy of u_i .

Observe however that the above transduction can also be implemented by a one-way transducer, using non-determinism: when entering a factor u_i , the transducer guesses whether or not $u_i \in (abc)^*$ and $|u_{i+1}|$ is even; depending on this it outputs either $(abc abc)^{\frac{|u_i|}{3}}$ or u_i , and checks that the guess is correct while proceeding to read the input.

The main result of our paper is an elementary algorithm that decides whether a functional transducer is one-way definable:

Theorem 3.3. *There is an algorithm that takes as input a functional two-way transducer T and outputs in 3EXPTIME a one-way transducer T' satisfying the following properties:*

- (1) $T' \subseteq T$,
- (2) $\text{dom}(T') = \text{dom}(T)$ if and only if T is one-way definable.
- (3) $\text{dom}(T') = \text{dom}(T)$ can be checked in 2EXPSPACE .

Moreover, if T is a sweeping transducer, then T' can be constructed in 2EXPTIME and $\text{dom}(T') = \text{dom}(T)$ is decidable in EXPSPACE .

Remark 3.4. The transducer T' constructed in the above theorem is in a certain sense maximal: for every $v \in \text{dom}(T) \setminus \text{dom}(T')$ and every one-way transducer T'' with $\text{dom}(T') \subseteq \text{dom}(T'') \subseteq \text{dom}(T)$ there exists some witness input v' obtained from v such that $v' \in \text{dom}(T) \setminus \text{dom}(T'')$. We will make this more precise at the end of Section 8.

We also provide a two-exponential lower bound for the size of the equivalent transducer. As the lower bound is achieved by a sweeping transduction (even a deterministic one), this gives a tight lower bound on the size of any one-way transducer equivalent to some sweeping transducer.

Proposition 3.5. *There is a family $(f_n)_{n \in \mathbb{N}}$ of transductions such that*

- (1) f_n can be implemented by a deterministic sweeping transducer of size $\mathcal{O}(n^2)$,
- (2) f_n can be implemented by a one-way transducer,
- (3) every one-way transducer that implements f_n has size $\Omega(2^{2^n})$.

Proof. The family of transformations is precisely the one described in Example 3.1 (2), where f_n maps inputs of the form $u = a_0 w_0 \cdots a_{2^n-1} w_{2^n-1}$ to outputs of the form $u u$, where $a_i \in \{a, b\}$ and $w_i \in \{0, 1\}^n$ is the binary encoding of i . A deterministic sweeping transducer implementing f_n first checks that the binary encodings w_i , for $i = 0, \dots, 2^n - 1$, are correct. This can be done with n passes: the j -th pass uses $\mathcal{O}(n)$ states to check the correctness of the j -th bits of the binary encodings. Then, the sweeping transducer performs two additional passes to copy the input twice. Overall, the sweeping transducer has size $\mathcal{O}(n^2)$.

As already mentioned, every one-way transducer that implements f_n needs to remember input words u of exponential length in order to output $u u$, which roughly requires doubly exponentially many states. A more formal argument providing a lower bound for the size of a one-way transducer implementing f_n goes as follows.

First of all, one observes that given a one-way transducer T , the language of its outputs, i.e., $L_T^{\text{out}} = \{w : (u, w) \in \mathcal{L}(T) \text{ for some } u\}$ is regular. More precisely, if T has size N , then the language L_T^{out} is recognized by an automaton of size linear in N . Indeed, while parsing w , the automaton can guess an input word u and a run on u , together with a factorization of w in which the i -th factor corresponds to the output of the transition on the i -th letter of u . Basically, this requires storing as control states the transition rules of T and the suffixes of outputs.

Now, suppose that the function f_n is implemented by a one-way transducer T of size N . The language $L_T^{\text{out}} = \{u u : u \in \text{dom}(f_n)\}$ is then recognized by an automaton of size $\mathcal{O}(N)$. Finally, we recall a result from [21], which shows that, given a sequence of pairs of words (u_i, v_i) , for $i = 1, \dots, M$, every non-deterministic automaton that separates the language $\{u_i v_i : 1 \leq i \leq M\}$ from the language $\{u_i u_j : 1 \leq i \neq j \leq M\}$ must have at least M states. By applying this result to our language L_T^{out} , where $u_i = v_i$ for all $i = 1, \dots, M = 2^{2^n}$, we get that N must be at least linear in M , and hence $N \in \Omega(2^{2^n})$. \square

The proof of Theorem 3.3 will be developed in the next sections. The main idea is to decompose a run of the two-way transducer T into factors that can be easily simulated in

a one-way manner. We defer the formal definition of such a decomposition to Section 5, while here we refer to it simply as a “ \mathbf{B} -decomposition”, where \mathbf{B} is a suitable number computed from T . The reader can refer to Figure 7 on page 18, which provides some intuitive account of a \mathbf{B} -decomposition for a sweeping run. Roughly speaking, each factor of a \mathbf{B} -decomposition either already looks like a run of a one-way transducer (e.g. the factors D_1 and D_2 of Figure 7), or it produces a periodic output, where the period is bounded by \mathbf{B} (e.g. the factor between ℓ_1 and ℓ_2). Identifying factors that look like runs of one-way transducers is rather easy. On the other hand, to identify factors with periodic outputs we rely on a notion of “inversion” of a run. Again, we defer the formal definition and the important combinatorial properties of inversions to Section 4. The reader can refer to Figure 4 on page 11 for an example of an inversion of a run of a sweeping transducer. Intuitively, this is a portion of run that is potentially difficult to simulate in a one-way manner, due to existence of long factors of the output that are generated following the opposite order of the input. Finally, the complexity of the decision procedure in Theorem 3.3 is analyzed in Section 9.

Roadmap. In order to provide a roadmap of our proofs, we state below the equivalence between the key properties related to one-way definability, inversions of runs, and existence of decompositions:

Theorem 3.6. *Given a functional two-way transducer T , an integer \mathbf{B} can be computed such that the following are equivalent:*

- P1)** T is one-way definable,
- P2)** for every successful run of T and every inversion in it, the output produced amid the inversion has period at most \mathbf{B} ,
- P3)** every input has a successful run of T that admits a \mathbf{B} -decomposition.

As the notions of inversion and \mathbf{B} -decomposition are simpler to formalize for sweeping transducers, we will first prove the theorem assuming that T is a sweeping transducer; we will focus later on unrestricted two-way transducers. Specifically, in Section 4 we introduce the basic combinatorics on words and the key notion of inversion for a run of a sweeping transducer, and we prove the implication **P1** \Rightarrow **P2**. In Section 5 we define \mathbf{B} -decompositions of runs of sweeping transducers, prove the implication **P2** \Rightarrow **P3**, and sketch a proof of **P3** \Rightarrow **P1** (as a matter of fact, this latter implication can be proved in a way that is independent of whether T is sweeping or not, which explains why we only sketch the proof in the sweeping case). Section 6 lays down the appropriate definitions concerning loops of two-way transducers, and analyzes in detail the effect of pumping special idempotent loops. In Section 7 we further develop the combinatorial arguments that are used to prove the implication **P1** \Rightarrow **P2** in the general case. Finally, in Section 8 we prove the implications **P2** \Rightarrow **P3** \Rightarrow **P1** in the general setting, and show how to decide the condition $\text{dom}(T') = \text{dom}(T)$ of Theorem 3.3.

4. BASIC COMBINATORICS FOR SWEEPING TRANSDUCERS

We fix for the rest of the section a functional *sweeping transducer* T , an input word u , and a (normalized) successful run ρ of T on u .

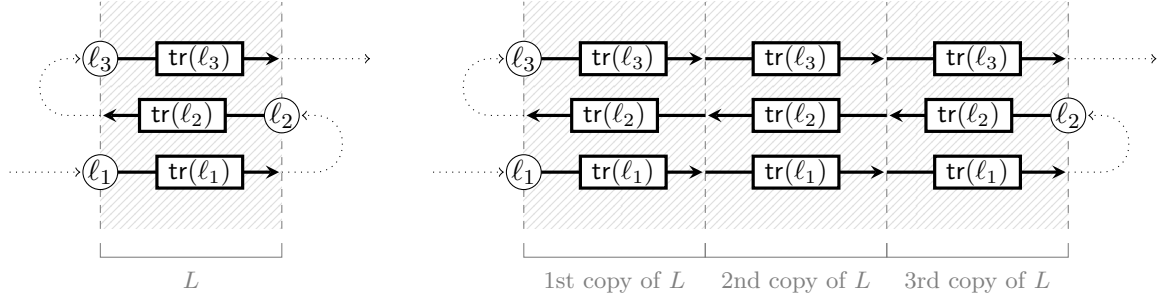


Figure 3. A loop L with 3 anchor points, and the result of pumping.

Pumping loops. Loops turn out to be a basic concept for characterizing one-way definability. Formally, a *loop* of ρ is an interval $L = [x_1, x_2]$ such that $\rho|_{x_1} = \rho|_{x_2}$, namely, with the same crossing sequences at the extremities. The run ρ can be pumped at any loop $L = [x_1, x_2]$, and this gives rise to new runs with iterated factors. Below we study precisely the shape of these pumped runs.

Definition 4.1 (anchor point, trace). Given a loop L and a location ℓ of ρ , we say that ℓ is an *anchor point in L* if ℓ is the first location of some factor of ρ that is intercepted by L ; this factor is then denoted² as $\text{tr}(\ell)$ and called the *trace of ℓ* .

Observe that a loop can have at most $\mathbf{H} = 2|Q| - 1$ anchor points, since we consider only normalized runs.

Given a loop L of ρ and a number $n \in \mathbb{N}$, we can replicate n times the factor $u[x_1, x_2]$ of the input, obtaining a new input of the form

$$\text{pump}_L^{n+1}(u) = u[1, x_1] \cdot (u[x_1 + 1, x_2])^{n+1} \cdot u[x_2 + 1, |u|]. \quad (4.1)$$

Similarly, we can replicate n times the intercepted factors $\text{tr}(\ell)$ of ρ , for all anchor points ℓ of L . In this way we obtain a successful run on $\text{pump}_L^{n+1}(u)$ that is of the form

$$\text{pump}_L^{n+1}(\rho) = \rho_0 \text{tr}(\ell_1)^n \rho_1 \dots \rho_{k-1} \text{tr}(\ell_k)^n \rho_k \quad (4.2)$$

where $\ell_1 \trianglelefteq \dots \trianglelefteq \ell_k$ are all the anchor points in L (listed according to the run order \trianglelefteq), ρ_0 is the prefix of ρ ending at ℓ_1 , ρ_k is the suffix of ρ starting at ℓ_k , and for all $i = 1, \dots, k - 1$, ρ_i is the factor of ρ between ℓ_i and ℓ_{i+1} . Note that $\text{pump}_L^1(\rho)$ coincides with the original run ρ . As a matter of fact, one could define in a similar way the run $\text{pump}_L^0(\rho)$ obtained from removing the loop L from ρ . However, we do not need this, and we will always parametrize the operation pump_L by a positive number $n + 1$.

An example of a pumped run $\text{pump}_{L_1}^3(\rho)$ is given in Figure 3, together with the indication of the anchor points ℓ_i and the intercepted factors $\text{tr}(\ell_i)$.

Output minimality. We are interested into factors of the run ρ that lie on a single level and that contribute to the final output, but in a minimal way, in the sense that is formalized by the following definition:

Definition 4.2. Consider a factor $\alpha = \rho[\ell, \ell']$ of ρ . We say that α is *output-minimal* if $\ell = (x, y)$ and $\ell' = (x', y)$, and all loops $L \subsetneq [x, x']$ produce empty output at level y .

²This is a slight abuse of notation, since the factor $\text{tr}(\ell)$ is not determined by ℓ alone, but requires also the knowledge of the loop L , which is usually clear from the context.




Figure 4. An inversion and the effect of pumping in an equivalent one-way transducer T' .

From now on, we set the constant $\mathbf{B} = \mathbf{C}|Q|^{\mathbf{H}} + 1$, where \mathbf{C} is the capacity of the transducer, that is, the maximal length of an output produced on a single transition (recall that $|Q|^{\mathbf{H}}$ is the maximal number of crossing sequences). As shown below, \mathbf{B} bounds the length of the output produced by an output-minimal factor:

Lemma 4.3. *For all output-minimal factors α , $|\text{out}(\alpha)| \leq \mathbf{B}$.*

Proof. Suppose by contradiction that $|\text{out}(\alpha)| > \mathbf{B}$, with $\alpha = \rho[\ell, \ell']$, $\ell = (x, y)$ and $\ell' = (x', y)$.

Let X be the set of all positions x'' , with $\min(x, x') < x'' < \max(x, x')$, that are sources of transitions of α that produce non-empty output. Clearly, the total number of letters produced by the transitions that depart from locations in $X \times \{y\}$ is strictly larger than $\mathbf{B} - 1$. Moreover, since each transition emits at most \mathbf{C} symbols, we have $|X| > \frac{\mathbf{B}-1}{\mathbf{C}} = |Q|^{\mathbf{H}}$. Now, recall that crossing sequences are sequences of states of length at most \mathbf{H} . Since $|X|$ is larger than the number of crossing sequences, X contains two positions $x_1 < x_2$ such that $\rho|x_1 = \rho|x_2$. In particular, $L = [x_1, x_2]$ is a loop strictly between x, x' with non-empty output on level y . This shows that $\rho[\ell, \ell']$ is not output-minimal. \square

Inversions and periodicity. Next, we define the crucial notion of inversion. Intuitively, an inversion in a run identifies a part of the run that is potentially difficult to simulate in a one-way manner because the order of generating the output is reversed w.r.t. the input. Inversions arise naturally in transducers that reverse arbitrarily long portions of the input, as well as in transducers that produce copies of arbitrarily long portions of the input.

Definition 4.4. An *inversion* of the run ρ is a tuple $(L_1, \ell_1, L_2, \ell_2)$ such that

- (1) L_1, L_2 are loops of ρ ,
- (2) $\ell_1 = (x_1, y_1)$ and $\ell_2 = (x_2, y_2)$ are anchor points of L_1 and L_2 , respectively,
- (3) $\ell_1 \triangleleft \ell_2$ and $x_1 > x_2$
(namely, ℓ_2 follows ℓ_1 in the run, but the position of ℓ_2 precedes the position of ℓ_1),
- (4) for both $i = 1$ and $i = 2$, $\text{out}(\text{tr}(\ell_i)) \neq \varepsilon$ and $\text{tr}(\ell_i)$ is output-minimal.

The left hand-side of Figure 4 gives an example of an inversion, assuming that the outputs $v_1 = \text{tr}(\ell_1)$ and $v_2 = \text{tr}(\ell_2)$ are non-empty and the intercepted factors are output-minimal.

The rest of the section is devoted to prove the implication $\mathbf{P1} \Rightarrow \mathbf{P2}$ of Theorem 3.6. We recall that a word $w = a_1 \cdots a_n$ has *period* p if for every $1 \leq i \leq |w| - p$, we have $a_i = a_{i+p}$. For example, the word $abcabcab$ has period 3.

We remark that, thanks to Lemma 4.3, for every inversion $(L_1, \ell_1, L_2, \ell_2)$, the outputs $\text{out}(\text{tr}(\ell_1))$ and $\text{out}(\text{tr}(\ell_2))$ have length at most \mathbf{B} . By pairing this with the assumption that the transducer T is one-way definable, and by using some classical word combinatorics, we show that the output produced between the anchor points of every inversion has period that divides the lengths of $\text{out}(\text{tr}(\ell_1))$ and $\text{out}(\text{tr}(\ell_2))$. In particular, this period is at most \mathbf{B} . The proposition below shows a slightly stronger periodicity property, which refers to the output produced between the anchor points ℓ_1, ℓ_2 of an inversion, but extended on both sides with the words $\text{out}(\text{tr}(\ell_1))$ and $\text{out}(\text{tr}(\ell_2))$. We will exploit this stronger periodicity property later, when dealing with overlapping portions of the run delimited by different inversions (cf. Lemma 5.5).

Proposition 4.5. *If T is one-way definable, then the following property **P2** holds:*

For all inversions $(L_1, \ell_1, L_2, \ell_2)$ of ρ , the period p of the word

$$\text{out}(\text{tr}(\ell_1)) \text{out}(\rho[\ell_1, \ell_2]) \text{out}(\text{tr}(\ell_2))$$

divides both $|\text{out}(\text{tr}(\ell_1))|$ and $|\text{out}(\text{tr}(\ell_2))|$. Moreover, $p \leq \mathbf{B}$.

The above proposition thus formalizes the implication **P1** \Rightarrow **P2** of Theorem 3.6. Its proof relies on a few combinatorial results. The first one is Fine and Wilf's theorem [29]. In short, this theorem says that, whenever two periodic words w_1, w_2 share a sufficiently long factor, then they have the same period. Here we use a slightly stronger variant of Fine and Wilf's theorem, which additionally shows how to align a common factor of the two words w_1, w_2 so as to form a third word containing a prefix of w_1 and a suffix of w_2 . This variant of Fine-Wilf's theorem will be particularly useful in the proof of Lemma 5.5, while for all other applications the classical statement suffices.

Theorem 4.6 (Fine-Wilf's theorem). *If $w_1 = w'_1 w w''_1$ has period p_1 , $w_2 = w'_2 w w''_2$ has period p_2 , and the common factor w has length at least $p_1 + p_2 - \gcd(p_1, p_2)$, then w_1, w_2 , and $w_3 = w'_1 w w''_2$ have period $\gcd(p_1, p_2)$.*

The second combinatorial result required in our proof concerns periods of words with iterated factors, like those that arise from considering outputs of pumped runs, and it is formalized precisely by the lemma below. To improve readability, we often highlight the important iterations of factors inside a word.

Lemma 4.7. *Assume that $v_0 \mathbf{v}_1^n v_2 \cdots v_{k-1} \mathbf{v}_k^n v_{k+1}$ has period p for some $n > p$. Then $v_0 \mathbf{v}_1^{n_1} v_2 \cdots v_{k-1} \mathbf{v}_k^{n_k} v_{k+1}$ has period p for all $n_1, \dots, n_k \in \mathbb{N}$.*

Proof. Assume that $w = v_0 \mathbf{v}_1^n v_2 \cdots v_{k-1} \mathbf{v}_k^n v_{k+1}$ has period p , and that $n > p$. Consider an arbitrary factor v_i^p of w . Since v_i^p has periods p and $|v_i|$, it has also period $r = \gcd(p, |v_i|)$. By Fine-Wilf (Theorem 4.6), we know that w has period r as well. Moreover, since the length of v_i is multiple of r , changing the number of repetitions of v_i inside w does not affect the period r of w . Since v_i was chosen arbitrarily, this means that, for all $n_1, \dots, n_k \in \mathbb{N}$, $v_0 \mathbf{v}_1^{n_1} v_2 \cdots v_{k-1} \mathbf{v}_k^{n_k} v_{k+1}$ has period r , and hence period p as well. \square

Recall that our goal is to show that the output produced amid every inversion has period bounded by \mathbf{B} . The general idea is to pump the loops of the inversion and compare the outputs of the two-way transducer T with those of an equivalent one-way transducer T' . The comparison leads to an equation between words with iterated factors, where the iterations are parametrized by two unknowns n_1, n_2 that occur in opposite order in the left, respectively right hand-side of the equation. Our third and last combinatorial result considers a word

equation of this precise form, and derives from it a periodicity property. For the sake of brevity, we use the notation $v^{(n_1, n_2)}$ to represent words with factors iterated n_1 or n_2 times, namely, words of the form $v_0 v_1^{n_{i_1}} v_2 \cdots v_{k-1} v_k^{n_{i_k}} v_{k+1}$, where the $v_0, v_1, v_2, \dots, v_{k-1}, v_k, v_{k+1}$ are fixed words (possibly empty) and each index among i_1, \dots, i_k is either 1 or 2.

Lemma 4.8. *Consider a word equation of the form*

$$v_0^{(n_1, n_2)} \mathbf{v}_1^{n_1} v_2^{(n_1, n_2)} \mathbf{v}_3^{n_2} v_4^{(n_1, n_2)} = w_0 \mathbf{w}_1^{n_2} w_2 \mathbf{w}_3^{n_1} w_4$$

where n_1, n_2 are the unknowns and v_1, v_3 are non-empty words. If the above equation holds for all $n_1, n_2 \in \mathbb{N}$, then

$$\mathbf{v}_1 \mathbf{v}_1^{n_1} v_2^{(n_1, n_2)} \mathbf{v}_3^{n_2} \mathbf{v}_3$$

has period $\gcd(|v_1|, |v_3|)$ for all $n_1, n_2 \in \mathbb{N}$.

Proof. The idea of the proof is to let the parameters n_1, n_2 of the equation grow independently, and apply Fine and Wilf's theorem (Theorem 4.6) a certain number of times to establish periodicities in overlapping factors of the considered words.

We begin by fixing n_1 large enough so that the factor $\mathbf{v}_1^{n_1}$ of the left hand-side of the equation becomes longer than $|w_0| + |w_1|$ (this is possible because v_1 is non-empty). Now, if we let n_2 grow arbitrarily large, we see that the length of the periodic word $\mathbf{w}_1^{n_2}$ is almost equal to the length of the left hand-side term $v_0^{(n_1, n_2)} \mathbf{v}_1^{n_1} v_2^{(n_1, n_2)} \mathbf{v}_3^{n_2} v_4^{(n_1, n_2)}$: indeed, the difference in length is given by the constant $|w_0| + |w_2| + n_1 \cdot |w_3| + |w_4|$. In particular, this implies that $\mathbf{w}_1^{n_2}$ covers arbitrarily long prefixes of $\mathbf{v}_1 v_2^{(n_1, n_2)} \mathbf{v}_3^{n_2+1}$, which in its turn contains long repetitions of the word v_3 . Hence, by Theorem 4.6, the word $\mathbf{v}_1 v_2^{(n_1, n_2)} \mathbf{v}_3^{n_2+1}$ has period $|v_3|$.

We remark that the periodicity shown so far holds for a large enough n_1 and for all but finitely many n_2 , where the threshold for n_2 depends on n_1 : once n_1 is fixed, n_2 needs to be larger than $f(n_1)$, for a suitable function f . In fact, by using Lemma 4.7, with n_1 fixed and $n = n_2$ large enough, we deduce that the periodicity holds for large enough $n_1 \in \mathbb{N}$ and for all $n_2 \in \mathbb{N}$.

We could also apply a symmetric reasoning: we choose n_2 large enough and let n_1 grow arbitrarily large. Doing so, we prove that for a large enough n_2 and for all but finitely many n_1 , the word $\mathbf{v}_1^{n_1+1} v_2^{(n_1, n_2)} \mathbf{v}_3$ is periodic with period $|v_1|$. As before, with the help of Lemma 4.7, this can be strengthened to hold for large enough $n_2 \in \mathbb{N}$ and for all $n_1 \in \mathbb{N}$.

Putting together the results proven so far, we get that for all but finitely many n_1, n_2 ,

$$\underbrace{\mathbf{v}_1^{n_1} \cdot \mathbf{v}_1 \cdot v_2^{(n_1, n_2)} \cdot \mathbf{v}_3 \cdot \mathbf{v}_3^{n_2}}_{\text{period } |v_1|} \cdot \mathbf{v}_3$$

Finally, we observe that the prefix $\mathbf{v}_1^{n_1+1} \cdot v_2^{(n_1, n_2)} \cdot \mathbf{v}_3$ and the suffix $\mathbf{v}_1 \cdot v_2^{(n_1, n_2)} \cdot \mathbf{v}_3^{n_2+1}$ share a common factor of length at least $|v_1| + |v_3|$. By Theorem 4.6, we derive that $\mathbf{v}_1^{n_1+1} \cdot v_2^{(n_1, n_2)} \cdot \mathbf{v}_3^{n_2+1}$ has period $\gcd(|v_1|, |v_3|)$ for all but finitely many n_1, n_2 . Finally, by using again Lemma 4.7, we conclude that the periodicity holds for all $n_1, n_2 \in \mathbb{N}$. \square

We are now ready to prove the implication **P1** \Rightarrow **P2**:

Proof of Proposition 4.5. Let T' be a one-way transducer equivalent to T , and consider an inversion $(L_1, \ell_1, L_2, \ell_2)$ of the successful run ρ of T on input u . The reader may refer to Figure 4 to get basic intuition about the proof technique. For simplicity, we assume that the loops L_1 and L_2 are disjoint, as shown in the figure. If this were not the case, we would have at least $\max(L_1) > \min(L_2)$, since the anchor point ℓ_1 is strictly to the right of the anchor point ℓ_2 . We could then consider the pumped run $\text{pump}_{L_1}^k(\rho)$ for a large enough $k > 1$ in such a way that the rightmost copy of L_1 turns out to be disjoint from and strictly to the right of L_2 . We could thus reason as we do below, by replacing everywhere (except in the final part of the proof, cf. *Transferring periodicity to the original run*) the run ρ with the pumped run $\text{pump}_{L_1}^k(\rho)$, and the formal parameter m_1 with $m_1 + k$.

Inducing loops in T' . We begin by pumping the run ρ and the underlying input u , on the loops L_1 and L_2 , in order to induce new loops L'_1 and L'_2 that are also loops in a successful run of T' . Assuming that L_1 is strictly to the right of L_2 , we define for all numbers $m_1, m_2 \in \mathbb{N}$:

$$\begin{aligned} u^{(m_1, m_2)} &= \text{pump}_{L_1}^{m_1+1}(\text{pump}_{L_2}^{m_2+1}(u)) \\ \rho^{(m_1, m_2)} &= \text{pump}_{L_1}^{m_1+1}(\text{pump}_{L_2}^{m_2+1}(\rho)). \end{aligned}$$

In the pumped run $\rho^{(m_1, m_2)}$, we identify the positions that mark the endpoints of the occurrences of L_1, L_2 . More precisely, if $L_1 = [x_1, x_2]$ and $L_2 = [x_3, x_4]$, with $x_1 > x_4$, then the sets of these positions are

$$\begin{aligned} X_2^{(m_1, m_2)} &= \{x_3 + i(x_4 - x_3) : 0 \leq i \leq m_2 + 1\} \\ X_1^{(m_1, m_2)} &= \{x_1 + j(x_2 - x_1) + m_2(x_4 - x_3) : 0 \leq j \leq m_1 + 1\}. \end{aligned}$$

Periodicity of outputs of pumped runs. We use now the fact that T' is a one-way transducer equivalent to T . We first recall (see, for instance, [8, 16]) that every functional one-way transducer can be made unambiguous, namely, can be transformed into an equivalent one-way transducer that admits at most one successful run on each input. This means that, without loss of generality, we can assume that T' too is unambiguous, and hence it admits exactly one successful run, say $\lambda^{(m_1, m_2)}$, on each input $u^{(m_1, m_2)}$.

Since T' has finitely many states, we can find, for a large enough number k_0 two positions $x'_1 < x'_2$, both in $X_1^{(k_0, k_0)}$, such that $L'_1 = [x'_1, x'_2]$ is a loop of $\lambda^{(k_0, k_0)}$. Similarly, we can find two positions $x'_3 < x'_4$, both in $X_2^{(k_0, k_0)}$, such that $L'_2 = [x'_3, x'_4]$ is a loop of $\lambda^{(k_0, k_0)}$. By construction L'_1 (resp. L'_2) consists of $k_1 \leq k_0$ (resp. $k_2 \leq k_0$) copies of L_1 (resp. L_2), and hence L'_1, L'_2 are also loops of $\rho^{(k_0, k_0)}$. In particular, this implies that for all $n_1, n_2 \in \mathbb{N}$:

$$\begin{aligned} \text{pump}_{L'_1}^{n_1+1}(\text{pump}_{L'_2}^{n_2+1}(u^{(k_0, k_0)})) &= u^{(f(n_1), g(n_2))} \\ \text{pump}_{L'_1}^{n_1+1}(\text{pump}_{L'_2}^{n_2+1}(\rho^{(k_0, k_0)})) &= \rho^{(f(n_1), g(n_2))} \\ \text{pump}_{L'_1}^{n_1+1}(\text{pump}_{L'_2}^{n_2+1}(\lambda^{(k_0, k_0)})) &= \lambda^{(f(n_1), g(n_2))}. \end{aligned}$$

where $f(n_1) = k_1 n_1 + k_0$ and $g(n_2) = k_2 n_2 + k_0$.

Now recall that $\rho^{(f(n_1), g(n_2))}$ and $\lambda^{(f(n_1), g(n_2))}$ are runs of T and T' on the same word $u^{(f(n_1), g(n_2))}$, and they produce the same output. Let us denote this output by $w^{(f(n_1), g(n_2))}$. Below, we show two possible factorizations of $w^{(f(n_1), g(n_2))}$ based on the shapes of the pumped runs $\lambda^{(f(n_1), g(n_2))}$ and $\rho^{(f(n_1), g(n_2))}$. For the first factorization, we recall that L'_2

precedes L'_1 , according to the ordering of positions, and that the run $\lambda^{(f(n_1),g(n_2))}$ is one-way (in particular loops have only one anchor point and one trace). We thus obtain:

$$w^{(f(n_1),g(n_2))} = w_0 \mathbf{w}_1^{n_2} w_2 \mathbf{w}_3^{n_1} w_4 \quad (4.3)$$

where

- w_0 is the output produced by the prefix of $\lambda^{(k_0,k_0)}$ ending at the only anchor point of L'_2 ,
 - \mathbf{w}_1 is the trace of L'_2 ,
 - w_2 is the output produced by the factor of $\lambda^{(k_0,k_0)}$ between the anchor points of L'_2 and L'_1 ,
 - \mathbf{w}_3 is the trace of L'_1 ,
 - w_4 is the output produced by the suffix of $\lambda^{(k_0,k_0)}$ starting at the anchor point of L'_1 .
- Hence, $w_0 w_2 w_4$ is the output of $\lambda^{(k_0,k_0)}$.

For the second factorization, we consider L'_1 and L'_2 as loops of $\rho^{(k_0,k_0)}$. We recall that ℓ_1, ℓ_2 are anchor points of the loops L_1, L_2 of ρ , and that there are corresponding copies of these anchor points in the pumped run $\rho^{(f(n_1),g(n_2))}$. We define ℓ'_1 (resp. ℓ'_2) to be the first (resp. last) location in $\rho^{(f(n_1),g(n_2))}$ that corresponds to ℓ_1 (resp. ℓ_2) and that is an anchor point of a copy of L'_1 (resp. L'_2). For example, if $\ell_1 = (x_1, y_1)$, with y_1 even, then $\ell'_1 = (x_1 + f(n_2)(x_4 - x_3), y_1)$. Thanks to Equation 4.2 we know that the output produced by $\rho^{(f(n_1),g(n_2))}$ is of the form

$$w^{(f(n_1),g(n_2))} = v_0^{(n_1,n_2)} \mathbf{v}_1^{n_1} v_2^{(n_1,n_2)} \mathbf{v}_3^{n_2} v_4^{(n_1,n_2)} \quad (4.4)$$

where

- $\mathbf{v}_1 = \text{out}(\text{tr}(\ell'_1))$, where ℓ'_1 is seen as an anchor point in a copy of L'_1 ,
- $\mathbf{v}_3 = \text{out}(\text{tr}(\ell'_2))$, where ℓ'_2 is seen as an anchor point in a copy of L'_2 (note that the words v_1, v_3 depend on k_0 , but not on n_1, n_2),
- $v_0^{(n_1,n_2)}$ is the output produced by the prefix of $\rho^{(f(n_1),g(n_2))}$ that ends at ℓ'_1 (this word may depend on the parameters n_1, n_2 since the loops L'_1, L'_2 may be traversed several times before reaching the first occurrence of $\text{tr}(\ell'_1)$),
- $v_2^{(n_1,n_2)}$ is the output produced by the factor of $\rho^{(f(n_1),g(n_2))}$ that starts at ℓ'_1 and ends at ℓ'_2 ,
- $v_4^{(n_1,n_2)}$ is the output produced by the suffix of $\rho^{(f(n_1),g(n_2))}$ that starts at ℓ'_2 .

Putting together Equations (4.3) and (4.4), we get

$$v_0^{(n_1,n_2)} \mathbf{v}_1^{n_1} v_2^{(n_1,n_2)} \mathbf{v}_3^{n_2} v_4^{(n_1,n_2)} = w_0 \mathbf{w}_1^{n_2} w_2 \mathbf{w}_3^{n_1} w_4. \quad (4.5)$$

Recall that the definition of inversions (Definition 4.4) states that the words v_1, v_3 are non-empty. This allows us to apply Lemma 4.8, which shows that the word $\mathbf{v}_1 \mathbf{v}_1^{n_1} v_2^{(n_1,n_2)} \mathbf{v}_3^{n_2} \mathbf{v}_3$ has period $p = \gcd(|v_1|, |v_3|)$, for all $n_1, n_2 \in \mathbb{N}$.

Note that the latter period p still depends on T' , since the words v_1, v_3 were obtained from loops L'_1, L'_2 on the run $\lambda^{(k_0,k_0)}$ of T' . However, because each loop L'_i consists of k_i copies of the original loop L_i , we also know that $v_1 = (\text{out}(\text{tr}(\ell_1)))^{k_1}$ and $v_3 = (\text{out}(\text{tr}(\ell_2)))^{k_2}$. By Theorem 4.6, this implies that for all $n_1, n_2 \in \mathbb{N}$, the word

$$(\text{out}(\text{tr}(\ell_1))) (\text{out}(\text{tr}(\ell_1)))^{k_1 n_1} v_2^{(n_1,n_2)} (\text{out}(\text{tr}(\ell_2)))^{k_2 n_2} (\text{out}(\text{tr}(\ell_2)))$$

has a period that divides $|\text{out}(\text{tr}(\ell_1))|$ and $|\text{out}(\text{tr}(\ell_2))|$.

Transferring periodicity to the original run. The last part of the proof amounts at showing a similar periodicity property for the output produced by the original run ρ . By construction, the iterated factors inside $v_2^{(n_1, n_2)}$ in the previous word are all of the form $v^{k_1 n_1 + k_0}$ or $v^{k_2 n_2 + k_0}$, for some words v . By taking out the constant factors v^{k_0} from the latter repetitions, we can write $v_2^{(n_1, n_2)}$ as a word with iterated factors of the form $v^{k_1 n_1}$ or $v^{k_2 n_2}$, namely, as $v_2^{(k_1 n_1, k_2 n_2)}$. So the word

$$(\text{out}(\text{tr}(\ell_1))) (\text{out}(\text{tr}(\ell_1)))^{k'_1} v_2^{(k'_1, k'_2)} (\text{out}(\text{tr}(\ell_2)))^{k'_2} (\text{out}(\text{tr}(\ell_2)))$$

is periodic, with period that divides $|\text{out}(\text{tr}(\ell_1))|$ and $|\text{out}(\text{tr}(\ell_2))|$, for all $k'_1 \in \{k_1 n : n \in \mathbb{N}\}$ and all $k'_2 \in \{k_2 n : n \in \mathbb{N}\}$. We now apply Lemma 4.7, once with $n = k'_1$ and once with $n = k'_2$, to conclude that the latter periodicity property holds also for $k'_1 = 1$ and $k'_2 = 1$. This shows that the word

$$\text{out}(\text{tr}(\ell_1)) \text{out}(\rho[\ell_1, \ell_2]) \text{out}(\text{tr}(\ell_2))$$

is periodic, with period that divides $|\text{out}(\text{tr}(\ell_1))|$ and $|\text{out}(\text{tr}(\ell_2))|$. \square

5. ONE-WAY DEFINABILITY IN THE SWEEPING CASE

In the previous section we have shown the implication $\mathbf{P1} \Rightarrow \mathbf{P2}$ for a functional sweeping transducer T . Here we close the cycle by proving the implications $\mathbf{P2} \Rightarrow \mathbf{P3}$ and $\mathbf{P3} \Rightarrow \mathbf{P1}$. In particular, we show how to derive the existence of successful runs admitting a \mathbf{B} -decomposition and construct a one-way transducer T' that simulates T on those runs. This will basically prove Theorem 3.6 in the sweeping case.

Run decomposition. We begin by giving the definition of \mathbf{B} -decomposition for a run ρ of T . Intuitively, a \mathbf{B} -decomposition of ρ identifies factors of ρ that can be easily simulated in a one-way manner. The definition below describes precisely the shape of these factors.

First we need to recall the notion of almost periodicity: a word w is *almost periodic with bound p* if $w = w_0 w_1 w_2$ for some words w_0, w_2 of length at most p and some word w_1 of period at most p .

We need to work with subsequences of the run ρ that are induced by particular sets of locations, not necessarily consecutive. Recall that $\rho[\ell, \ell']$ denotes the factor of ρ delimited by two locations $\ell \leq \ell'$. Similarly, given any set Z of locations, we denote by $\rho|Z$ the subsequence of ρ induced by Z . Note that, even though $\rho|Z$ might not be a valid run of the transducer, we can still refer to the number of transitions in it and to the size of the produced output $\text{out}(\rho|Z)$. Formally, a transition in $\rho|Z$ is a transition from some ℓ to ℓ' , where both ℓ, ℓ' belong to Z . The output $\text{out}(\rho|Z)$ is the concatenation of the outputs of the transitions of $\rho|Z$ (according to the order given by ρ).

Definition 5.1. Consider a factor $\rho[\ell, \ell']$ of a run ρ of T , where $\ell = (x, y)$, $\ell' = (x', y')$ are two locations with $x \leq x'$. We call $\rho[\ell, \ell']$

- a *floor* if $y = y'$ is even, namely, if $\rho[l, \ell']$ lies entirely on the same level and is rightward oriented;
- a *B-diagonal* if there is a sequence $\ell = \ell_0 \leq \ell_1 \leq \dots \leq \ell_{2n+1} = \ell'$, where each $\rho[\ell_{2i+1}, \ell_{2i+2}]$ is a floor, each $\rho[\ell_{2i}, \ell_{2i+1}]$ produces an output of length at most $2HB$, and the position of each ℓ_{2i} is to the left of the position of ℓ_{2i+1} ;
- a *B-block* if the output produced by $\rho[l, \ell']$ is almost periodic with bound $2B$, and the output produced by the subsequence $\rho|Z$, where $Z = [l, \ell'] \setminus ([x, x'] \times [y, y'])$, has length at most $2HB$.




Figure 5. A diagonal




Figure 6. A block.

Before continuing we give some intuition on the above definitions. The simplest concept is that of floor, which is a rightward oriented factor of a run. Diagonals are sequences of consecutive floors interleaved by factors that produce small (bounded) outputs. We see an example of a diagonal in Figure 5, where we marked the important locations and highlighted with thick arrows the two floors that form the diagonal. The factors of the diagonal that are not floors are represented instead by dotted arrows. The third notion is that of a block. An important constraint in the definition of a block is that the produce output must be almost periodic, with small enough bound. In Figure 6, the periodic output is represented by the thick arrows, either solid or dotted, that go from location ℓ to location ℓ' . In addition, the block must satisfy a constraint on the length of the output produced by the subsequence $\rho|Z$, where $Z = [l, \ell'] \setminus ([x, x'] \times [y, y'])$. The latter set Z consists of location that are either to the left of the position of ℓ or to the right of the position of ℓ' . For example, in Figure 6 the set Z coincides with the area outside the hatched rectangle. Accordingly, the portion of the subsequence $\rho|Z$ is represented by the dotted bold arrows.

Diagonals and blocks are used as key objects to derive a notion of decomposition for a run of a sweeping transducer. We formalize this notion below.

Definition 5.2. A *B-decomposition* of a run ρ of T is a factorization $\prod_i \rho[\ell_i, \ell_{i+1}]$ of ρ into *B-diagonals* and *B-blocks*.

Figure 7 gives an example of such a decomposition. Each factor is either a diagonal D_i or a block B_i . Inside each factor we highlight by thick arrows the portions of the run that can be simulated by a one-way transducer, either because they are produced from left to right (thus forming diagonals) or because they are periodic (thus forming blocks). We also recall from Definition 5.1 that most of the output is produced inside the hatched rectangles, since the output produced by a diagonal or a block outside the corresponding blue or red hatched rectangle has length at most $2H^2B$. Finally, we observe that the locations delimiting the factors of the decomposition are arranged following the natural order of positions and levels.

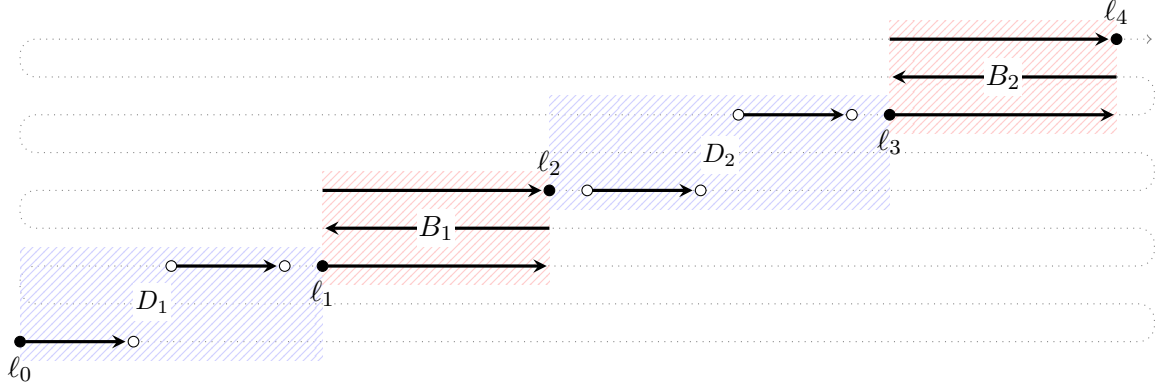


Figure 7. Decomposition of a run into diagonals and blocks.

All together, this means that the output produced by a run that enjoys a decomposition can be simulated in a one-way manner.

From periodicity of inversions to existence of decompositions. Now that we set up the definition of B -decomposition, we turn towards proving the implication $\mathbf{P2} \Rightarrow \mathbf{P3}$ of Theorem 3.6. In fact, we will prove a slightly stronger result than $\mathbf{P2} \Rightarrow \mathbf{P3}$, which is stated further below. Formally, when we say that a run ρ satisfies $\mathbf{P2}$ we mean that for every inversion $(L_1, \ell_1, L_2, \ell_2)$ of ρ , the word $\text{out}(\text{tr}(\ell_1)) \text{out}(\rho[\ell_1, \ell_2]) \text{out}(\text{tr}(\ell_2))$ has period $\gcd(|\text{out}(\text{tr}(\ell_1))|, |\text{out}(\text{tr}(\ell_2))|) \leq B$. We aim at proving that every run that satisfies $\mathbf{P2}$ enjoys a decomposition, independently of whether other runs do or do not satisfy $\mathbf{P2}$:

Proposition 5.3. *If ρ is a run of T that satisfies $\mathbf{P2}$, then ρ admits a B -decomposition.*

Let us fix a run ρ of T and assume that it satisfies $\mathbf{P2}$. To show that ρ admits a B -decomposition, we will identify the blocks of the decomposition as equivalence classes of a suitable relation based on inversions (cf. Definition 5.4). Then, we will use combinatorial arguments (notably, Lemmas 4.3 and 5.5) to prove that the constructed blocks satisfy the desired properties. Finally, we will show how the resulting equivalence classes form all the necessary blocks of the decomposition, in the sense that the factors in between those classes are diagonals.

We begin by introducing the equivalence relation by means of which we can then identify the blocks of a decomposition of ρ .

Definition 5.4. Let \mathbf{S} be the relation that pairs every two locations ℓ, ℓ' of ρ whenever there is an inversion $(L_1, \ell_1, L_2, \ell_2)$ of ρ such that $\ell_1 \preceq \ell, \ell' \preceq \ell_2$, namely, whenever ℓ and ℓ' occur within the same inversion. Let \mathbf{S}^* be the reflexive and transitive closure of \mathbf{S} .

It is easy to see that every equivalence class of \mathbf{S}^* is a convex subset with respect to the run order on locations of ρ . Moreover, every *non-singleton* equivalence class of \mathbf{S}^* is a union of a series of inversions that are two-by-two overlapping. One can refer to Figure 8 for an intuitive account of what we mean by two-by-two overlapping: the thick arrows represent factors of the run that lie entire inside an \mathbf{S}^* -equivalence class, each inversion is identified by a pair of consecutive anchor points with the same color. According to the run order, between every pair of anchor points with the same color, there is at least one anchor point of different color: this means that the inversions corresponding to the two colors are overlapping.




Figure 8. A non-singleton \mathbf{S}^* -equivalence class seen as a series of overlapping inversions.

Formally, we say that an inversion $(L_1, \ell_1, L_2, \ell_2)$ *covers* a location ℓ when $\ell_1 \trianglelefteq \ell \trianglelefteq \ell_2$. We say that two inversions $(L_1, \ell_1, L_2, \ell_2)$ and $(L_3, \ell_3, L_4, \ell_4)$ are *overlapping* if $(L_1, \ell_1, L_2, \ell_2)$ covers ℓ_3 and $(L_3, \ell_3, L_4, \ell_4)$ covers ℓ_2 (or the other way around).

The next lemma uses the fact that ρ satisfies **P2** to deduce that the output produced inside every \mathbf{S}^* -equivalence class has period at most \mathbf{B} . Note that the proof below does not exploit the fact that the transducer is sweeping.

Lemma 5.5. *If ρ satisfies **P2** and $\ell \trianglelefteq \ell'$ are two locations of ρ such that $\ell \mathbf{S}^* \ell'$, then the output $\text{out}(\rho[\ell, \ell'])$ produced between these locations has period at most \mathbf{B} .*

Proof. The claim for $\ell = \ell'$ holds trivially, so we assume that $\ell \triangleleft \ell'$. Since the \mathbf{S}^* -equivalence class that contains ℓ, ℓ' is non-singleton, we know that there is a series of inversions

$$(L_0, \ell_0, L_1, \ell_1) \quad (L_2, \ell_2, L_3, \ell_3) \quad \dots \quad (L_{2k}, \ell_{2k}, L_{2k+1}, \ell_{2k+1})$$

that are two-by-two overlapping and such that $\ell_0 \trianglelefteq \ell \triangleleft \ell' \trianglelefteq \ell_{2k+1}$. Without loss of generality, we can assume that every inversion $(L_{2i}, \ell_{2i}, L_{2i+1}, \ell_{2i+1})$ is *maximal* in the following sense: there is no other inversion $(\tilde{L}, \tilde{\ell}, \tilde{L}', \tilde{\ell}') \neq (L_{2i}, \ell_{2i}, L_{2i+1}, \ell_{2i+1})$ such that $\tilde{\ell} \trianglelefteq \ell_{2i} \trianglelefteq \ell_{2i+1} \trianglelefteq \tilde{\ell}'$.

For the sake of brevity, let $v_i = \text{out}(\text{tr}(\ell_i))$ and $p_i = |v_i|$. Since ρ satisfies **P2** (recall Proposition 4.5), we know that, for all $i = 0, \dots, n$, the word

$$v_{2i} \text{ out}(\rho[\ell_{2i}, \ell_{2i+1}]) v_{2i+1}$$

has period that divides both p_{2i} and p_{2i+1} and is at most \mathbf{B} . In order to show that the period of $\text{out}(\rho[\ell, \ell'])$ is also bounded by \mathbf{B} , it suffices to prove the following claim by induction on i :

Claim 5.6. For all $i = 0, \dots, k$, the word $\text{out}(\rho[\ell_0, \ell_{2i+1}]) v_{2i+1}$ has period at most \mathbf{B} that divides p_{2i+1} .

The base case $i = 0$ follows immediately from our hypothesis, since $(L_0, \ell_0, L_1, \ell_1)$ is an inversion. For the inductive step, we assume that the claim holds for $i < k$, and we prove it

for $i + 1$. First of all, we factorize our word as follows:

$$\text{out}(\rho[\ell_0, \ell_{2i+3}]) v_{2i+3} = \underbrace{\text{out}(\rho[\ell_0, \ell_{2i+2}])}_{\text{period } p_{2i+1}} \overbrace{\text{out}(\rho[\ell_{2i+2}, \ell_{2i+1}]) \text{out}(\rho[\ell_{2i+1}, \ell_{2i+3}])}^{\text{periods } p_{2i+2} \text{ and } p_{2i+3}} v_{2i+3}.$$

By the inductive hypothesis, the output produced between ℓ_0 and ℓ_{2i+1} , extended to the right with v_{2i+1} , has period that divides p_{2i+1} . Moreover, because ρ satisfies **P2** and $(L_{2i+2}, \ell_{2i+2}, L_{2i+3}, \ell_{2i+3})$ is an inversion, the output produced between the locations ℓ_{2i+2} and ℓ_{2i+3} , extended to the left with v_{2i+2} and to the right with v_{2i+3} , has period that divides both p_{2i+2} and p_{2i+3} . Note that this is not yet sufficient for applying Fine-Wilf's theorem, since the common factor $\text{out}(\rho[\ell_{2i+2}, \ell_{2i+1}])$ might be too short (possibly just equal to v_{2i+2}). The key argument here is to prove that the interval $[\ell_{2i+2}, \ell_{2i+1}]$ is covered by an inversion which is different from those that we considered above, namely, i.e. $(L_{2i+2}, \ell_{2i+2}, L_{2i+1}, \ell_{2i+1})$. For example, $[\ell_2, \ell_1]$ in Figure 8 is covered by the inversion $(L_2, \ell_2, L_1, \ell_1)$.

For this, we have to prove that the anchor points ℓ_{2i+2} and ℓ_{2i+1} are correctly ordered w.r.t. \trianglelefteq and the ordering of positions (recall Definition 4.4). First, we observe that $\ell_{2i+2} \trianglelefteq \ell_{2i+1}$, since $(L_{2i}, \ell_{2i}, L_{2i+1}, \ell_{2i+1})$ and $(L_{2i+2}, \ell_{2i+2}, L_{2i+3}, \ell_{2i+3})$ are overlapping inversions. Next, we prove that the position of ℓ_{2i+1} is strictly to the left of the position of ℓ_{2i+2} . By way of contradiction, suppose that this is not the case, namely, $\ell_{2i+1} = (x_{2i+1}, y_{2i+1})$, $\ell_{2i+2} = (x_{2i+2}, y_{2i+2})$, and $x_{2i+1} \geq x_{2i+2}$. Because $(L_{2i}, \ell_{2i}, L_{2i+1}, \ell_{2i+1})$ and $(L_{2i+2}, \ell_{2i+2}, L_{2i+3}, \ell_{2i+3})$ are inversions, we know that ℓ_{2i+3} is strictly to the left of ℓ_{2i+2} and ℓ_{2i+1} is strictly to the left of ℓ_{2i} . This implies that ℓ_{2i+3} is strictly to the left of ℓ_{2i} , and hence $(L_{2i}, \ell_{2i}, L_{2i+3}, \ell_{2i+3})$ is also an inversion. Moreover, recall that $\ell_{2i} \trianglelefteq \ell_{2i+2} \trianglelefteq \ell_{2i+1} \trianglelefteq \ell_{2i+3}$. This contradicts the maximality of $(L_{2i}, \ell_{2i}, L_{2i+1}, \ell_{2i+1})$, which we assumed at the beginning of the proof. Therefore, we must conclude that ℓ_{2i+1} is strictly to the left of ℓ_{2i+2} .

Now that we know that $\ell_{2i+2} \trianglelefteq \ell_{2i+1}$ and that ℓ_{2i+1} is to the left of ℓ_{2i+2} , we derive the existence of the inversion $(L_{2i+2}, \ell_{2i+2}, L_{2i+1}, \ell_{2i+1})$. Again, because ρ satisfies **P2**, we know that the word $v_{2i+2} \text{out}(\rho[\ell_{2i+2}, \ell_{2i+1}]) v_{2i+1}$ has period at most \mathbf{B} that divides p_{2i+2} and p_{2i+1} . Summing up, we have:

- (1) $w_1 = \text{out}(\rho[\ell_0, \ell_{2i+1}]) v_{2i+1}$ has period p_{2i+1} ,
- (2) $w_2 = v_{2i+2} \text{out}(\rho[\ell_{2i+2}, \ell_{2i+1}]) v_{2i+1}$ has period $p = \gcd(p_{2i+2}, p_{2i+1})$,
- (3) $w_3 = v_{2i+2} \text{out}(\rho[\ell_{2i+2}, \ell_{2i+3}]) v_{2i+3}$ has period $p' = \gcd(p_{2i+2}, p_{2i+3})$.

We are now ready to exploit our stronger variant of Fine-Wilf's theorem, that is, Theorem 4.6.

We begin with (1) and (2) above. Let $w = \text{out}(\rho[\ell_{2i+2}, \ell_{2i+1}]) v_{2i+1}$ be the common suffix of w_1 and w_2 . First note that since p divides p_{2i+2} , the word w is also a prefix of w_2 , thus we can write $w_2 = w w'_2$. Second, note that the length of w is at least $|v_{2i+1}| = p_{2i+1} = p_{2i+1} + p - \gcd(p_{2i+1}, p)$. We can apply now Theorem 4.6 to $w_1 = w'_1 w$ and $w_2 = w w'_2$ and obtain:

- (4) $w_4 = w'_1 w w'_2 = \text{out}(\rho[\ell_0, \ell_{2i+2}]) v_{2i+2} \text{out}(\rho[\ell_{2i+2}, \ell_{2i+1}]) v_{2i+1}$ has period p .

We apply next Theorem 4.6 to (2) and (3), namely, to the words w_2 and w_3 with v_{2i+2} as common factor. It is not difficult to check that $|v_{2i+2}| = p_{2i+2} \geq p + p' - p''$ with $p'' = \gcd(p, p')$, using the definitions of p and p' : we can write $p_{2i+2} = p'' r q = p'' r' q'$ with $p = p'' r$ and $p' = p'' r'$. It suffices to check that $p'' r q + p'' r' q' \geq 2(p'' r + p'' r' - p'')$, hence that $r q + r' q' \geq 2r + 2r' - 2$. This is clear if $\min(q, q') > 1$. Otherwise the inequality $p_{2i+2} \geq p + p' - p''$ follows easily because $p = p''$ or $p' = p''$ holds. Hence we obtain that w_2 and w_3 have both period p'' .

Applying once more Theorem 4.6 to w_3 and w_4 with v_{2i+2} as common factor, yields period p'' for the word

$$w_5 = \text{out}(\rho[\ell_0, \ell_{2i+2}]) v_{2i+2} \text{out}(\rho[\ell_{2i+2}, \ell_{2i+3}]) v_{2i+3}$$

Finally, the periodicity is not affected when we remove factors of length multiple than the period. In particular, by removing the factor v_{2i+2} from w_5 , we obtain the desired word $\text{out}(\rho[\ell_0, \ell_{2i+3}]) v_{2i+3}$, whose period p'' divides p_{2i+3} . This proves the claim for the inductive step, and completes the proof of the proposition. \square

The \mathbf{S}^* -classes considered so far cannot be directly used to define the blocks in the desired decomposition of ρ , since the x -coordinates of their endpoints might not be in the appropriate order. The next definition takes care of this, by enlarging the \mathbf{S}^* -classes according to x -coordinates of the anchor points in the equivalence class.

Definition 5.7. Consider a non-singleton \mathbf{S}^* -equivalence class $K = [\ell, \ell']$. Let $\text{an}(K)$ be the restriction of K to the anchor points occurring in some inversion, and $X_{\text{an}(K)} = \{x : \exists y (x, y) \in \text{an}(K)\}$ be the projection of $\text{an}(K)$ on positions. We define $\text{block}(K) = [\tilde{\ell}, \tilde{\ell}']$, where

- $\tilde{\ell}$ is the latest location $(\tilde{x}, \tilde{y}) \trianglelefteq \ell$ such that $\tilde{x} = \min(X_{\text{an}(K)})$,
- $\tilde{\ell}'$ is the earliest location $(\tilde{x}, \tilde{y}) \trianglerighteq \ell'$ such that $\tilde{x} = \max(X_{\text{an}(K)})$

(note that the locations $\tilde{\ell}, \tilde{\ell}'$ exist since ℓ, ℓ' are anchor points in some inversion).




Figure 9. Block construction.

Lemma 5.8. *If K is a non-singleton \mathbf{S}^* -equivalence class, then $\rho|_{\text{block}(K)}$ is a \mathbf{B} -block.*

Proof. Consider a non-singleton \mathbf{S}^* -class $K = [\ell, \ell']$ and let $\text{an}(K)$, $X_{\text{an}(K)}$, and $\text{block}(K) = [\tilde{\ell}, \tilde{\ell}']$ be as in Definition 5.7. The reader can refer to Figure 9 to quickly recall the notation. We need to verify that $\rho|_{[\tilde{\ell}, \tilde{\ell}]}$ is a \mathbf{B} -block (cf. Definition 5.1), namely, that:

- $\tilde{\ell} = (\tilde{x}, \tilde{y})$, $\tilde{\ell}' = (\tilde{x}', \tilde{y}')$, with $\tilde{x} \leq \tilde{x}'$,
- the output produced by $\rho|_{[\tilde{\ell}, \tilde{\ell}]}$ is almost periodic with bound $2\mathbf{B}$,
- the output produced by the subsequence $\rho|_Z$, where $Z = [\tilde{\ell}, \tilde{\ell}] \setminus ([\tilde{x}, \tilde{x}') \times [\tilde{y}, \tilde{y}']$, has length at most $2\mathbf{HB}$.

The first condition $\tilde{x} \leq \tilde{x}'$ follows immediately from the definition of \tilde{x} and \tilde{x}' as $\min(X_{\text{an}(K)})$ and $\max(X_{\text{an}(K)})$, respectively.

Next, we prove that the output produced by the factor $\rho|_{[\tilde{\ell}, \tilde{\ell}]}$ is almost periodic with bound $2\mathbf{B}$. By Definition 5.7, we have $\tilde{\ell} \trianglelefteq \ell \triangleleft \ell' \triangleleft \tilde{\ell}'$, and by Lemma 5.5 we know that $\text{out}(\rho[\ell, \ell'])$ is periodic with period at most \mathbf{B} ($\leq 2\mathbf{B}$). So it suffices to show that the lengths of the words $\text{out}(\rho[\tilde{\ell}, \ell])$ and $\text{out}(\rho[\ell', \tilde{\ell}'])$ are at most $2\mathbf{B}$. We shall focus on the former word, as the arguments for the latter are similar.

First, we note that the factor $\rho|_{[\tilde{\ell}, \ell]}$ lies entirely to the right of position \tilde{x} , and in particular, it starts at an even level \tilde{y} . This follows from the definition of $\tilde{\ell}$, and whether ℓ itself is at an odd/even level. In particular, the location ℓ is either at the same level as $\tilde{\ell}$, or just one level above.

Now, suppose, by way of contradiction, that $|\text{out}(\rho[\tilde{\ell}, \ell])| > 2\mathbf{B}$. We head towards a contradiction by finding a location $\ell'' \triangleleft \ell$ that is \mathbf{S}^* -equivalent to the first location ℓ of the \mathbf{S}^* -equivalence class K . Since the location ℓ is either at the same level as $\tilde{\ell}$, or just above it, the factor $\rho[\tilde{\ell}, \ell]$ is of the form $\alpha\beta$, where α is rightward factor lying on the same level as $\tilde{\ell}$ and β is either empty or a leftward factor on the next level. Moreover, since $|\text{out}(\rho[\tilde{\ell}, \ell])| > 2\mathbf{B}$, we know that either $|\text{out}(\alpha)| > \mathbf{B}$ or $|\text{out}(\beta)| > \mathbf{B}$. Thus, Lemma 4.3 says that one of the two factors α, β is not output-minimal. In particular, there is a loop L_1 , strictly to the right of \tilde{x} , that intercepts a subfactor γ of $\rho[\tilde{\ell}, \ell]$, with $\text{out}(\gamma)$ non-empty and output-minimal.

Let ℓ'' be the first location of the factor γ . Clearly, ℓ'' is an anchor point of L and $\text{out}(\text{tr}(\ell'')) \neq \varepsilon$. Further recall that $\tilde{x} = \min(X_{\text{an}(K)})$ is the leftmost position of locations in the class $K = [\ell, \ell']$ that are also anchor points of inversions. In particular, there is a loop L_2 with some anchor point $\ell''_2 = (\tilde{x}, y''_2) \in \text{an}(K)$, and such that $\text{tr}(\ell'')$ is non-empty and output-minimal. Since $\ell'' \triangleleft \ell \leq \ell''_2$ and the position of ℓ'' is to the right of the position of ℓ''_2 , we know that $(L_1, \ell'', L_2, \ell''_2)$ is also an inversion, and hence $\ell'' \mathbf{S}^* \ell''_2 \mathbf{S}^* \ell$. But since $\ell'' \triangleleft \ell$, we get a contradiction with the assumption that ℓ is the first location of a \mathbf{S}^* -class. In this way we have shown that $|\text{out}(\rho[\ell_1, \ell])| \leq 2\mathbf{B}$.

It remains to show that the output produced by the subsequence $\rho|Z$, where $Z = [\tilde{\ell}, \tilde{\ell}'] \setminus ([\tilde{x}, \tilde{x}'] \times [\tilde{y}, \tilde{y}'])$, has length at most $2\mathbf{H}\mathbf{B}$. For this it suffices to prove that $|\text{out}(\alpha)| \leq \mathbf{B}$ for every factor α of $\rho[\tilde{\ell}, \tilde{\ell}']$ that lies at a single level and either to the left of \tilde{x} or to the right of \tilde{x}' . By symmetry, we consider only one of the two types of factors. Suppose, by way of contradiction, that there is a factor α at level y'' , to the left of \tilde{x} , and such that $|\text{out}(\alpha)| > \mathbf{B}$. By Lemma 4.3 we know that α is not output-minimal, so there is some loop L_2 strictly to the left of \tilde{x} that intercepts an output-minimal subfactor β of α with non-empty output. Let ℓ''' be the first location of β . We know that $\tilde{\ell} \triangleleft \ell''' \leq \tilde{\ell}'$. Since the level \tilde{y} is even, this means that the level of ℓ''' is strictly greater than \tilde{y} . Since we also know that ℓ is an anchor point of some inversion, we can take a suitable loop L_1 with anchor point ℓ and obtain that $(L_1, \ell, L_2, \ell''')$ is an inversion, so $\ell''' \mathbf{S}^* \ell$. But this contradicts the fact that \tilde{x} is the leftmost position of $\text{an}(K)$. We thus conclude that $|\text{out}(\alpha)| \leq \mathbf{B}$, and this completes the proof that $\rho|\text{block}(K)$ is a \mathbf{B} -block. \square

The next lemma shows that blocks do not overlap along the input axis:

Lemma 5.9. *Suppose that K_1 and K_2 are two different non-singleton \mathbf{S}^* -classes such that $\ell \triangleleft \ell'$ for all $\ell \in K_1$ and $\ell' \in K_2$. Let $\text{block}(K_1) = [\ell_1, \ell_2]$ and $\text{block}(K_2) = [\ell_3, \ell_4]$, with $\ell_2 = (x_2, y_2)$ and $\ell_3 = (x_3, y_3)$. Then $x_2 < x_3$.*

Proof. Suppose by contradiction that K_1 and K_2 are as in the statement, but $x_2 \geq x_3$. By Definition 5.7, $x_2 = \max(X_{\text{an}(K_1)})$ and $x_3 = \min(X_{\text{an}(K_2)})$. This implies the existence of some inversions (L, ℓ, L', ℓ') and $(L'', \ell'', L''', \ell''')$ such that $\ell = (x_2, y)$ and $\ell''' = (x_3, y''')$. Moreover, since $\ell \leq \ell'''$ and $x_2 \geq x_3$, we know that (L, ℓ, L''', ℓ''') is also an inversion, thus implying that $K_1 = K_2$. \square

For the sake of brevity, we call \mathbf{S}^* -block any factor of the form $\rho|\text{block}(K)$ that is obtained by applying Definition 5.7 to a non-singleton \mathbf{S}^* -class K . The results obtained so far imply that every location covered by an inversion is also covered by an \mathbf{S}^* -block (Lemma 5.8), and that the order of occurrence of \mathbf{S}^* -blocks is the same as the order of positions (Lemma 5.9). So the \mathbf{S}^* -blocks can be used as factors for the \mathbf{B} -decomposition of ρ we

are looking for. Below, we show that the remaining factors of ρ , which do not overlap the \mathbf{S}^* -blocks, are \mathbf{B} -diagonals. This will complete the construction of a \mathbf{B} -decomposition of ρ .

Formally, we say that a factor $\rho[\ell, \ell']$ overlaps another factor $\rho[\ell'', \ell''']$ if $[\ell, \ell'] \cap [\ell'', \ell'''] \neq \emptyset$, $\ell' \neq \ell''$, and $\ell \neq \ell'''$.

Lemma 5.10. *Let $\rho[\ell, \ell']$ be a factor of ρ , with $\ell = (x, y)$, $\ell' = (x', y')$, and $x \leq x'$, that does not overlap any \mathbf{S}^* -block. Then $\rho[\ell, \ell']$ is a \mathbf{B} -diagonal.*

Proof. Consider a factor $\rho[\ell, \ell']$, with $\ell = (x, y)$, $\ell' = (x', y')$, and $x \leq x'$, that does not overlap any \mathbf{S}^* -block. We will focus on locations ℓ'' with $\ell \leq \ell'' \leq \ell'$ that are anchor points of some loop with $\text{out}(\text{tr}(\ell'')) \neq \varepsilon$. We denote by A the set of all such locations.

First, we show that the locations in A are monotonic w.r.t. the position order. Formally, we prove that for all $\ell_1, \ell_2 \in A$, if $\ell_1 = (x_1, y_1) \leq \ell_2 = (x_2, y_2)$, then $x_1 \leq x_2$. Suppose that this were not the case, namely, that A contained two anchor points $\ell_1 = (x_1, y_1)$ and $\ell_2 = (x_2, y_2)$ with $\ell_1 \triangleleft \ell_2$ and $x_1 > x_2$. Let L_1, L_2 be the loops of ℓ_1, ℓ_2 , respectively, and recall that $\text{out}(\text{tr}(\ell_1)), \text{out}(\text{tr}(\ell_2)) \neq \varepsilon$. This means that $(L_1, \ell_1, L_2, \ell_2)$ is an inversion, and hence $\ell_1 \mathbf{S}^* \ell_2$. But this contradicts the hypothesis that $\rho[\ell, \ell']$ does not overlap any \mathbf{S}^* -block.

Next, we identify the floors of our diagonal. Let y_0, y_1, \dots, y_{n-1} be all the even levels that have locations in A . For each $i = 0, \dots, n-1$, let ℓ_{2i+1} (resp. ℓ_{2i+2}) be the first (resp. last) anchor point of A at level y_i . Further let $\ell_0 = \ell$ and $\ell_{2n+1} = \ell'$. Clearly, each factor $\rho[\ell_{2i+1}, \ell_{2i+2}]$ is a floor. Moreover, thanks to the previous arguments, each location ℓ_{2i} is to the left of the location ℓ_{2i+1} .

It remains to prove that each factor $\rho[\ell_{2i}, \ell_{2i+1}]$ produces an output of length at most $2\mathbf{HB}$. By construction, A contains no anchor point at an even level and strictly between ℓ_{2i} and ℓ_{2i+1} . By Lemma 4.3 this means that the outputs produced by subfactors of $\rho[\ell_{2i}, \ell_{2i+1}]$ that lie entirely at an even level have length at most \mathbf{B} . Let us now consider the subfactors α of $\rho[\ell_{2i}, \ell_{2i+1}]$ that lie entirely at an odd level, and let us prove that they produce outputs of length at most $2\mathbf{B}$. Suppose that this is not the case, namely, that $|\text{out}(\alpha)| > 2\mathbf{B}$. In this case we show that an inversion would exist at this level. Formally, we can find two locations $\ell'' \triangleleft \ell'''$ in α such that the prefix of α that ends at location ℓ'' and the suffix of α that starts at location ℓ''' produce outputs of length greater than \mathbf{B} . By Lemma 4.3, those two factors would not be output-minimal, and hence α would contain disjoint loops L_1, L_2 with anchor points ℓ''_1, ℓ''_2 forming an inversion $(L_1, \ell''_1, L_2, \ell''_2)$. But this would imply that ℓ''_1, ℓ''_2 belong to the same non-singleton \mathbf{S}^* -equivalence class, which contradicts the hypothesis that $\rho[\ell, \ell']$ does not overlap any \mathbf{S}^* -block. We must conclude that the subfactors of $\rho[\ell_{2i}, \ell_{2i+1}]$ produce outputs of length at most $2\mathbf{B}$.

Overall, this shows that the output produced by each factor $\rho[\ell_{2i}, \ell_{2i+1}]$ has length at most $2\mathbf{HB}$. \square

We have just shown how to construct a \mathbf{B} -decomposition of the run ρ that satisfies **P2**. This proves Proposition 5.3, as well as the implication **P2** \Rightarrow **P3** of Theorem 3.6.

From existence of decompositions to an equivalent one-way transducer. We now focus on the last implication **P3** \Rightarrow **P1** of Theorem 3.6. More precisely, we show how to construct a one-way transducer T' that simulates the outputs produced by the successful runs of T that admit \mathbf{B} -decompositions. In particular, T' turns out to be equivalent to T when T is one-way definable. Here we will only give a proof sketch of this construction

(as there is no real difference between the sweeping and two-way cases) assuming that T is a sweeping transducer; a fully detailed construction of T' from an arbitrary two-way transducer T will be given in Section 8 (Proposition 8.6), together with a procedure for deciding one-way definability of T (Proposition 9.2).

Proposition 5.11. *Given a functional sweeping transducer T a one-way transducer T' can be constructed in 2EXPTIME such that the following hold:*

- (1) $T' \subseteq T$,
- (2) $\text{dom}(T')$ contains all words that induce successful runs of T admitting \mathbf{B} -decompositions.

In particular, T' is equivalent to T iff T is one-way definable.

Proof sketch. Given an input word u , the one-way transducer T' needs to guess a successful run ρ of T on u that admits a \mathbf{B} -decomposition. This can be done by guessing the crossing sequences of ρ at each position, together with a sequence of locations ℓ_i that identify the factors of a \mathbf{B} -decomposition of ρ . To check the correctness of the decomposition, T' also needs to guess a bounded amount of information (words of bounded length) to reconstruct the outputs produced by the \mathbf{B} -diagonals and the \mathbf{B} -blocks. For example, while scanning a factor of the input underlying a diagonal, T' can easily reproduce the outputs of the floors and the guessed outputs of factors between them. In a similar way, while scanning a factor of the input underlying a block, T' can simulate the almost periodic output by guessing its repeating pattern and the bounded prefix and suffix of it, and by emitting the correct amount of letters, as it is done in Example 3.2. In particular, one can verify that the capacity of T' is linear in \mathbf{HB} . Moreover, because the guessed objects are of size linear in \mathbf{HB} and \mathbf{HB} is a simple exponential in the size of T , the size of the one-way transducer T' has doubly exponential size in that of T . \square

6. THE STRUCTURE OF TWO-WAY LOOPS

While loop pumping in a sweeping transducer is rather simple, we need a much better understanding when it comes to pump loops of unrestricted two-way transducers. This section is precisely devoted to untangling the structure of two-way loops. We will focus on specific loops, called idempotent, that generate repetitions with a “nice shape”, very similar to loops of sweeping transducers.

We fix throughout this section a functional two-way transducer T , an input word u , and a (normalized) successful run ρ of T on u . As usual, $\mathbf{H} = 2|Q| - 1$ is the maximal length of a crossing sequence of ρ , and \mathbf{C} is the maximal number of letters output by a single transition.

Flows and effects. We start by analyzing the shape of factors of ρ intercepted by an interval $I = [x_1, x_2]$. We identify four types of factors α intercepted by I depending on the first location (x, y) and the last location (x', y') :

- α is an LL-factor if $x = x' = x_1$,
- α is an RR-factor if $x = x' = x_2$,
- α is an LR-factor if $x = x_1$ and $x' = x_2$,
- α is an RL-factor if $x = x_2$ and $x' = x_1$.

In Figure 2 we see that α is an LL-factor, β, δ are LR-factors, ζ is an RR-factor, and γ is an RL-factor.

Definition 6.1. Let $I = [x_1, x_2]$ be an interval of ρ and h_i the length of the crossing sequence $\rho|_{x_i}$, for both $i = 1$ and $i = 2$.

The *flow* F_I of I is the directed graph with set of nodes $\{0, \dots, \max(h_1, h_2) - 1\}$ and set of edges consisting of all (y, y') such that there is a factor of ρ intercepted by I that starts at location (x_i, y) and ends at location (x_j, y') , for $i, j \in \{1, 2\}$.

The *effect* E_I of I is the triple (F_I, c_1, c_2) , where $c_i = \rho|_{x_i}$ is the crossing sequence at x_i .

For example, the interval I of Figure 2 has the flow graph $0 \mapsto 1 \mapsto 3 \mapsto 4 \mapsto 2 \mapsto 0$.

It is easy to see that every node of a flow F_I has at most one incoming and at most one outgoing edge. More precisely, if $y < h_1$ is even, then it has one outgoing edge (corresponding to an LR- or LL-factor intercepted by I), and if it is odd it has one incoming edge (corresponding to an RL- or LL-factor intercepted by I). Similarly, if $y < h_2$ is even, then it has one incoming edge (corresponding to an LR- or RR-factor), and if it is odd it has one outgoing edge (corresponding to an RL- or RR-factor).

In the following we consider effects that are not necessarily associated with intervals of specific runs. The definition of such effects should be clear: they are triples consisting of a graph (called flow) and two crossing sequences of lengths $h_1, h_2 \leq \mathbf{H}$, with sets of nodes of the form $\{0, \dots, \max(h_1, h_2) - 1\}$, that satisfy the in/out-degree properties stated above. It is convenient to distinguish the edges in a flow based on the parity of the source and target nodes. Formally, we partition any flow F into the following subgraphs:

- F_{LR} consists of all edges of F between pairs of even nodes,
- F_{RL} consists of all edges of F between pairs of odd nodes,
- F_{LL} consists of all edges of F from an even node to an odd node,
- F_{RR} consists of all edges of F from an odd node to an even node.

We denote by \mathcal{F} (resp. \mathcal{E}) the set of all flows (resp. effects) augmented with a dummy element \perp . We equip both sets \mathcal{F} and \mathcal{E} with a semigroup structure, where the corresponding products \circ and \odot are defined below (similar definitions appear in [9]). Later we will use the semigroup structure to identify the *idempotent loops*, that play a crucial role in our characterization of one-way definability.

Definition 6.2. For two graphs G, G' , we denote by $G \cdot G'$ the graph with edges of the form (y, y'') such that (y, y') is an edge of G and (y', y'') is an edge of G' , for some node y' that belongs to both G and G' . Similarly, we denote by G^* the graph with edges (y, y') such that there exists a (possibly empty) path in G from y to y' .

The product of two flows F, F' is the unique flow $F \circ F'$ (if it exists) such that:

- $(F \circ F')_{\text{LR}} = F_{\text{LR}} \cdot (F'_{\text{LL}} \cdot F'_{\text{RR}})^* \cdot F'_{\text{LR}}$,
- $(F \circ F')_{\text{RL}} = F'_{\text{RL}} \cdot (F_{\text{RR}} \cdot F_{\text{LL}})^* \cdot F_{\text{RL}}$,
- $(F \circ F')_{\text{LL}} = F_{\text{LL}} \cup F_{\text{LR}} \cdot (F'_{\text{LL}} \cdot F'_{\text{RR}})^* \cdot F'_{\text{LL}} \cdot F_{\text{RL}}$,
- $(F \circ F')_{\text{RR}} = F'_{\text{RR}} \cup F'_{\text{RL}} \cdot (F_{\text{RR}} \cdot F_{\text{LL}})^* \cdot F_{\text{RR}} \cdot F'_{\text{LR}}$.

If no flow $F \circ F'$ exists with the above properties, then we let $F \circ F' = \perp$.

The product of two effects $E = (F, c_1, c_2)$ and $E' = (F', c'_1, c'_2)$ is either the effect $E \odot E' = (F \circ F', c_1, c'_2)$ or the dummy element \perp , depending on whether $F \circ F' \neq \perp$ and $c_2 = c'_1$.

For example, let F be the flow of interval I in Figure 2. Then $(F \circ F)_{\text{LL}} = \{0 \mapsto 1, 2 \mapsto 3\}$, $(F \circ F)_{\text{RR}} = \{1 \mapsto 2, 3 \mapsto 4\}$, and $(F \circ F)_{\text{LR}} = \{4 \mapsto 0\}$ — one can quickly verify this with the help of Figure 10.

It is also easy to see that (\mathcal{F}, \circ) and (\mathcal{E}, \odot) are finite semigroups, and that for every run ρ and every pair of consecutive intervals $I = [x_1, x_2]$ and $J = [x_2, x_3]$ of ρ , $F_{I \cup J} = F_I \circ F_J$




Figure 10. Pumping a loop in a two-way run.

and $E_{I \cup J} = E_I \odot E_J$. In particular, the function E that associates each interval I of ρ with the corresponding effect E_I can be seen as a semigroup homomorphism.

Loops and components. Recall that a loop is an interval $L = [x_1, x_2]$ with the same crossing sequences at x_1 and x_2 . We will follow techniques similar to those presented in Section 4 to show that the outputs generated in non left-to-right manner are essentially periodic. However, differently from the sweeping case, we will consider only special types of loops:

Definition 6.3. A loop L is *idempotent* if $E_L = E_L \odot E_L$ and $E_L \neq \perp$.

For example, the interval I of Figure 2 is a loop, if one assumes that the crossing sequences at the borders of I are the same. By comparing with Figure 10, it is easy to see that I is not idempotent. On the other hand, the loop consisting of 2 copies of I is idempotent.

As usual, given a loop $L = [x_1, x_2]$ and a number $n \in \mathbb{N}$, we can introduce n new copies of L and connect the intercepted factors in the obvious way. This results in a new run $\text{pump}_L^{n+1}(\rho)$ on the word $\text{pump}_L^{n+1}(u)$. Figure 10 shows how to do this for $n = 1$ and $n = 2$. Below, we analyze in detail the shape of the pumped run $\text{pump}_L^{n+1}(\rho)$ (and the produced output as well) when L is an *idempotent* loop. We will focus on idempotent loops because pumping non-idempotent loops may induce permutations of factors that are difficult to handle. For example, if we consider again the non-idempotent loop I to the left of Figure 10, the factor of the run between β and γ (to the right of I , highlighted in red) precedes the factor between γ and δ (to the left of I , again in red), but this ordering is reversed when a new copy of I is added.

When pumping a loop L , subsets of factors intercepted by L are glued together to form factors intercepted by the replication of L . The notion of component introduced below identifies groups of factors that are glued together.

Definition 6.4. A *component* of a loop L is any strongly connected component of its flow F_L (note that this is also a cycle, since every node in it has in/out-degree 1).

Given a component C , we denote by $\min(C)$ (resp. $\max(C)$) the minimum (resp. maximum) node in C . We say that C is *left-to-right* (resp. *right-to-left*) if $\min(C)$ is even (resp., odd).




Figure 11. Pumping an idempotent loop with three components.

An (L, C) -factor is a factor of the run that is intercepted by L and that corresponds to an edge of C .

We will usually list the (L, C) -factors based on their order of occurrence in the run. For example, the loop I of Figure 10 contains a single component $C = 0 \mapsto 1 \mapsto 3 \mapsto 4 \mapsto 2 \mapsto 0$ which is left-to-right. Another example is given in Figure 11, where the loop L has three components C_1, C_2, C_3 (colored in blue, red, and green, respectively): $\alpha_1, \alpha_2, \alpha_3$ are the (L, C_1) -factors, $\beta_1, \beta_2, \beta_3$ are the (L, C_2) -factors, and γ_1 is the unique (L, C_3) -factor.

Below, we show that the levels of each component of a loop (not necessarily idempotent) form an interval.

Lemma 6.5. *Let C be a component of a loop $L = [x_1, x_2]$. The nodes of C are precisely the levels in the interval $[\min(C), \max(C)]$. Moreover, if C is left-to-right (resp. right-to-left), then $\max(C)$ is the smallest level $\geq \min(C)$ such that between $(x_1, \min(C))$ and $(x_2, \max(C))$ (resp. $(x_2, \min(C))$ and $(x_1, \max(C))$) there are equally many LL-factors and RR-factors intercepted by L .*

Proof idea. The proof of this lemma is rather technical and deferred to Appendix A, since the lemma is not at the core of the proof of the main result. Let us first note that with the definition of $\max(C)$ stated in the lemma it is rather easy to see that the interval $[\min(C), \max(C)]$ is a union of cycles (i.e., components). This can be shown by arguing that every node in $[\min(C), \max(C)]$ has in-degree and out-degree one. What is much less obvious is that $[\min(C), \max(C)]$ is connected, thus consists of a single cycle.

The crux is thus to show that the nodes visited by every cycle of the flow (or, equally, every component) form an interval. For this, we use an induction based on portions of the cycle, namely, on *paths* of the flow. The difficulty underlying the formalization of the inductive invariant comes from the fact that, differently from cycles, paths of a flow may visit sets of levels that do not form intervals. An example is given in Figure 12, which represents some edges of a flow forming a path from \vec{y}_i to $\vec{y}_i + 1$ and covering a non-convex set of nodes: note that there could be a large gap between the nodes \vec{y}_i and \overleftarrow{y}_i due to the unbalanced numbers of LL-factors and RR-factors below \overleftarrow{y}_i .

Essentially, the first part of the proof of the lemma amounts at identifying the sources \vec{y}_i (resp. \overleftarrow{y}_i) of the LR-factors (resp. RL-factors), and at showing that the latter factors are precisely of the form $\vec{y}_i \rightarrow \overleftarrow{y}_{i-1} + 1$ (resp. $\overleftarrow{y}_i \rightarrow \vec{y}_i + 1$). Once these nodes are identified, we




Figure 12. Some factors intercepted by L and the corresponding edges in the flow.

show by induction on i that every two consecutive nodes \vec{y}_i and $\vec{y}_i + 1$ must be connected by a path whose intermediate nodes form the interval $[\vec{y}_{i-1} + 1, \vec{y}_i]$. Finally, we argue that every cycle (or component) C visits all and only the nodes in the interval $[\min(C), \max(C)]$. \square

The next lemma describes the precise shape and order of the intercepted factors when the loop L is idempotent.

Lemma 6.6. *If C is a left-to-right (resp. right-to-left) component of an idempotent loop L , then the (L, C) -factors are in the following order: k LL-factors (resp. RR-factors), followed by one LR-factor (resp. RL-factor), followed by k RR-factors (resp. LL-factors), for some $k \geq 0$.*

Proof. Suppose that C is a left-to-right component of L . We show by way of contradiction that C has only one LR-factor and no RL-factor. By Lemma 6.5 this will yield the claimed shape. Figure 13 can be used as a reference example for the arguments that follow.

We begin by listing the (L, C) -factors. As usual, we order them based on their occurrences in the run ρ . Let γ be the first (L, C) -factor that is not an LL-factor, and let β_1, \dots, β_k be the (L, C) -factors that precede γ (these are all LL-factors). Because γ starts at an even level, it must be an LR-factor. Suppose that there is another (L, C) -factor, say ζ , that comes after γ and it is neither an RR-factor nor an LL-factor. Because ζ starts at an odd level, it must be an RL-factor. Further let $\delta_1, \dots, \delta_{k'}$ be the intercepted RR-factors that occur between γ and ζ . We claim that $k' < k$, namely, that the number of RR-factors between γ and ζ is strictly less than the number of LL-factors before γ . Indeed, if this were not the case, then, by Lemma 6.5, the level where ζ starts would not belong to the component C .

Now, consider the pumped run $\rho' = \text{pump}_L^2(\rho)$, obtained by adding a new copy of L . Let L' be the loop of ρ' obtained from the union of L and its copy. Since L is idempotent, the components of L are isomorphic to the components of L' . In particular, we can denote by C' the component of L' that is isomorphic to C . Let us consider the (L', C') -factors of ρ' . The first k factors are isomorphic to the k LL-factors β_1, \dots, β_k from ρ . However, the $(k + 1)$ -th element has a different shape: it is isomorphic to $\gamma \beta_1 \delta_1 \beta_2 \dots \delta_{k'} \beta_{k'+1} \zeta$, and in particular it is an LL-factor. This implies that the $(k + 1)$ -th edge of C' is of the form $(y, y + 1)$, while the $(k + 1)$ -th edge of C is of the form $(y, y - 2k)$. This contradiction comes from having assumed the existence of the RL-factor ζ , and is illustrated in Figure 13. \square




Figure 13. Pumping a loop L with a wrong shape and showing it is not idempotent.

Remark 6.7. Note that every loop in the sweeping case is idempotent. Moreover, the (L, C) -factors are precisely the factors intercepted by the loop L .

Pumping idempotent loops. To describe in a formal way the run obtained by pumping an idempotent loop, we need to generalize the notion of anchor point in the two-way case (the reader may compare this with the analogous definitions in Section 4 for the sweeping case). Intuitively, the anchor point of a component C of an idempotent loop L is the source location of the unique LR- or RL-factor intercepted by L that corresponds to an edge of C (recall Lemma 6.6):

Definition 6.8. Let C be a component of an idempotent loop $L = [x_1, x_2]$. The *anchor point* of C inside L , denoted³ $\text{an}(C)$, is either the location $(x_1, \max(C))$ or the location $(x_2, \max(C))$, depending on whether C is left-to-right or right-to-left.

We will usually depict anchor points by black circles (like, for instance, in Figure 11).

It is also convenient to redefine the notation $\text{tr}(\ell)$ for representing an appropriate sequence of transitions associated with each anchor point ℓ of an idempotent loop:

Definition 6.9. Let C be a component of some idempotent loop L , let $\ell = \text{an}(C)$ be the anchor point of C inside L , and let $i_0 \mapsto i_1 \mapsto i_2 \mapsto \dots \mapsto i_k \mapsto i_{k+1}$ be a cycle of C , where $i_0 = i_{k+1} = \max(C)$. For every $j = 0, \dots, k$, further let β_j be the factor intercepted by L that corresponds to the edge $i_j \mapsto i_{j+1}$ of C .

The *trace* of ℓ inside L is the run $\text{tr}(\ell) = \beta_0 \beta_1 \dots \beta_k$.

Note that $\text{tr}(\ell)$ is not necessarily a factor of the original run ρ . However, $\text{tr}(\ell)$ is indeed a run, since L is a loop and the factors β_i are concatenated according to the flow. As we will see below, $\text{tr}(\ell)$ will appear as (iterated) factor of the pumped version of ρ , where the loop L is iterated.

As an example, by referring again to the components C_1, C_2, C_3 of Figure 11, we have the following traces: $\text{tr}(\text{an}(C_1)) = \alpha_2 \alpha_1 \alpha_3$, $\text{tr}(\text{an}(C_2)) = \beta_2 \beta_1 \beta_3$, and $\text{tr}(\text{an}(C_3)) = \gamma_1$.

The next proposition shows the effect of pumping idempotent loops. The reader can note the similarity with the sweeping case.

³In denoting the anchor point — and similarly the trace — of a component C inside a loop L , we omit the annotation specifying L , since this is often understood from the context.

Proposition 6.10. *Let L be an idempotent loop of ρ with components C_1, \dots, C_k , listed according to the order of their anchor points: $\ell_1 = \text{an}(C_1) \triangleleft \dots \triangleleft \ell_k = \text{an}(C_k)$. For all $n \in \mathbb{N}$, we have*

$$\text{pump}_L^{n+1}(\rho) = \rho_0 \text{tr}(\ell_1)^n \rho_1 \cdots \rho_{k-1} \text{tr}(\ell_k)^n \rho_k$$

where

- ρ_0 is the prefix of ρ that ends at the first anchor point ℓ_1 ,
- ρ_k is the suffix of ρ that starts at the last anchor point ℓ_k ,
- ρ_i is the factor $\rho[\ell_i, \ell_{i+1}]$, for all $1 \leq i < k$.

Proof. Along the proof we sometimes refer to Figure 11 to ease the intuition of some definitions and arguments. For example, in the left hand-side of Figure 11, the run ρ_0 goes until the first location marked by a black circle; the runs ρ_1 and ρ_2 , resp., are between the first and the second black dot, and the second and third black dot; finally, ρ_3 is the suffix starting at the last black dot. The pumped run $\text{pump}_L^{n+1}(\rho)$ for $n = 2$ is depicted to the right of the figure.

Let $L = [x_1, x_2]$ be an idempotent loop and, for all $i = 0, \dots, n$, let $L'_i = [x'_i, x'_{i+1}]$ be the i -th copy of the loop L in the pumped run $\rho' = \text{pump}_L^{n+1}(\rho)$, where $x'_i = x_1 + i \cdot (x_2 - x_1)$ (the “0-th copy of L ” is the loop L itself). Further let $L' = L'_0 \cup \dots \cup L'_n = [x'_0, x'_{n+1}]$, that is, L' is the loop of ρ' that spans across the $n + 1$ occurrences of L . As L is idempotent, the loops L'_0, \dots, L'_n and L' have all the same effect as L . In particular, the components of L'_0, \dots, L'_n , and L' are isomorphic to and in same order as those of L . We denote these components by C_1, \dots, C_k .

We let $\ell_j = \text{an}(C_j)$ be the anchor point of each component C_j inside the loop L of ρ (these locations are marked by black dots in the left hand-side of Figure 11). Similarly, we let $\ell'_{i,j}$ (resp. ℓ'_j) be the anchor point of C_j inside the loop L'_i (resp. L'). From Definition 6.8, we have that either $\ell'_j = \ell'_{1,j}$ or $\ell'_j = \ell'_{n,j}$, depending on whether C_j is left-to-right or right-to-left (or, equally, on whether j is odd or even).

Now, let us consider the factorization of the pumped run ρ' induced by the locations $\ell'_{i,j}$, for all $i = 0, \dots, n$ and for $j = 1, \dots, k$ (these locations are marked by black dots in the right hand-side of the figure). By construction, the prefix of ρ' that ends at location $\ell'_{0,1}$ coincides with the prefix of ρ that ends at ℓ_1 , i.e. ρ_0 in the statement of the proposition. Similarly, the suffix of ρ' that starts at location $\ell'_{n,k}$ is isomorphic to the suffix of ρ that starts at ℓ_k , i.e. ρ_k in the statement. Moreover, for all odd (resp. even) indices j , the factor $\rho'[\ell'_{n,j}, \ell'_{n,j+1}]$ (resp. $\rho'[\ell'_{0,j}, \ell'_{0,j+1}]$) is isomorphic to $\rho[\ell_j, \ell_{j+1}]$, i.e. the ρ_j of the statement.

The remaining factors of ρ' are those delimited by the pairs of locations $\ell'_{i,j}$ and $\ell'_{i+1,j}$, for all $i = 0, \dots, n - 1$ and all $j = 1, \dots, k$. Consider one such factor $\rho'[\ell'_{i,j}, \ell'_{i+1,j}]$, and assume that the index j is odd (the case of an even j is similar). This factor can be seen as a concatenation of factors intercepted by L that correspond to edges of C_j inside L'_i . More precisely, $\rho'[\ell'_{i,j}, \ell'_{i+1,j}]$ is obtained by concatenating the unique LR-factor of C_j — recall that by Lemma 6.6 there is exactly one such factor — with an interleaving of the LL-factors and the RR-factors of C_j . As the components are the same for all L'_i 's and for L , this corresponds precisely to the trace $\text{tr}(\ell_j)$ (cf. Definition 6.9). Now that we know that $\rho'[\ell'_{i,j}, \ell'_{i+1,j}]$ is isomorphic to $\text{tr}(\ell_j)$, we can conclude that $\rho'[\ell'_{0,j}, \ell'_{n,j}] = \rho'[\ell'_{0,j}, \ell'_{1,j}] \cdots \rho'[\ell'_{n-1,j}, \ell'_{n,j}]$ is isomorphic to $\text{tr}(\ell_j)^n$. \square

7. COMBINATORICS IN THE TWO-WAY CASE

In this section we develop the main combinatorial techniques required in the general case. In particular, we will show how to derive the existence of idempotent loops with bounded outputs using Ramsey-based arguments, and we will use this to derive periodicity properties for the outputs produced between inversions.

As usual, ρ is a fixed successful run of T on some input word u .

Ramsey-type arguments. We start with a technique used for bounding the lengths of the outputs of certain factors, or subsequences of a two-way run. This technique is a Ramsey-type argument, more precisely it relies on Simon’s “factorization forest” theorem [13, 34], which is recalled below. The classical version of Ramsey theorem would yield a similar result, but without the tight bounds that we get here.

Let X be a set of positions of ρ . A *factorization forest* for X is an unranked tree, where the nodes are intervals I with endpoints in X and labeled with the corresponding effect E_I , the ancestor relation is given by the containment order on intervals, the leaves are the minimal intervals $[x_1, x_2]$, with x_2 successor of x_1 in X , and for every internal node I with children J_1, \dots, J_k , we have:

- $I = J_1 \cup \dots \cup J_k$,
- $E_I = E_{J_1} \odot \dots \odot E_{J_k}$,
- if $k > 2$, then $E_I = E_{J_1} = \dots = E_{J_k}$ is an idempotent of the semigroup (\mathcal{E}, \odot) .

Recall that in a normalized run there are at most $|Q|^H$ distinct crossing sequences. Moreover, a flow contains at most H edges, and each edge has one of the 4 possible types LL, LR, RL, RR, so they are at most 4^H different flows. Hence, the effect semigroup (\mathcal{E}, \odot) has size at most $E = 4^H \cdot |Q|^{2H} \cdot |Q|^{2H} = (2|Q|)^{2H}$. Further recall that C is the maximum number of letters output by a single transition of T . Like we did in the sweeping case, we define the constant $B = C \cdot H \cdot (2^{3E} + 4) + 4C$ that will be used to bound the lengths of some outputs of T . Note that now B is doubly exponential with respect to the size of T , due to the size of the effect semigroup.

Theorem 7.1 (Factorization forest theorem [13, 34]). *For every set X of positions of ρ , there is a factorization forest for X of height at most $3E$.*

The above theorem can be used to show that if ρ produces an output longer than B , then it contains an idempotent loop and a trace with non-empty output. Below, we present a result in the same spirit, but refined in a way that it can be used to find anchor points inside specific intervals. To formally state the result, we consider subsequences of ρ induced by sets of locations that are not necessarily contiguous. Recall the notation $\rho|Z$ introduced on page 16: $\rho|Z$ is the subsequence of ρ induced by the location set Z . For example, Figure 14 depicts a set $Z = [\ell_1, \ell_2] \cap (I \times \mathbb{N})$ by a hatched area, together with the induced subrun $\rho|Z$, represented by thick arrows.




Figure 14. A subrun $\rho|Z$.

Theorem 7.2. *Let $I = [x_1, x_2]$ be an interval of positions, $K = [\ell_1, \ell_2]$ an interval of locations, and $Z = K \cap (I \times \mathbb{N})$. If $|\text{out}(\rho|Z)| > B$, then there exists an idempotent loop L and an anchor point ℓ of L such that*




Figure 15. Two consecutive idempotent loops with the same effect.

- (1) $x_1 < \min(L) < \max(L) < x_2$ (in particular, $L \subsetneq I$),
- (2) $\ell_1 \triangleleft \ell \triangleleft \ell_2$ (in particular, $\ell \in K$),
- (3) $\text{out}(\text{tr}(\ell)) \neq \varepsilon$.

Proof. Let I, K, Z be as in the statement, and suppose that $|\text{out}(\rho|Z)| > \mathbf{B}$. We define $Z' = Z \setminus (\{\ell_1, \ell_2\} \cup \{x_1, x_2\} \times \mathbb{N})$ and we observe that there are at most $2\mathbf{H} + 2$ locations that are missing from Z' . This means that $\rho|Z'$ contains all but $4\mathbf{H} + 4$ transitions of $\rho|Z$, and because each transition outputs at most \mathbf{C} letters, we have $|\text{out}(\rho|Z')| > \mathbf{B} - 4\mathbf{C} \cdot \mathbf{H} - 4\mathbf{C} = \mathbf{C} \cdot \mathbf{H} \cdot 2^{3\mathbf{E}}$.

For every level y , let X_y be the set of positions x such that (x, y) is the source location of some transition of $\rho|Z'$ that produces non-empty output. For example, if we refer to Figure 15, the vertical dashed lines represent the positions of X_y for a particular level y ; accordingly, the circles in the figure represent the locations of the form (x, y) , for all $x \in X_y$. Since each transition outputs at most \mathbf{C} letters, we have $\sum_y |X_y| > \mathbf{H} \cdot 2^{3\mathbf{E}}$. Moreover, since there are at most \mathbf{H} levels, there is a level y (which we fix hereafter) such that $|X_y| > 2^{3\mathbf{E}}$.

We now prove the following:

Claim 7.3. There are two consecutive loops $L_1 = [x, x']$ and $L_2 = [x', x'']$ with endpoints $x, x', x'' \in X_y$ and such that $E_{L_1} = E_{L_2} = E_{L_1 \cup L_2}$.

Proof. By Theorem 7.1, there is a factorization forest for X_y of height at most $3\mathbf{E}$. Since ρ is a valid run, the dummy element \perp of the effect semigroup does not appear in this factorization forest. Moreover, since $|X_y| > 2^{3\mathbf{E}}$, we know that the factorization forest contains an internal node $L' = [x'_1, x'_{k+1}]$ with $k > 2$ children, say $L_1 = [x'_1, x'_2], \dots, L_k = [x'_k, x'_{k+1}]$. By definition of factorization forest, the effects $E_{L'}, E_{L_1}, \dots, E_{L_k}$ are all equal and idempotent. In particular, the effect $E_{L'} = E_{L_1} = \dots = E_{L_k}$ is a triple of the form $(F_{L'}, c_1, c_2)$, where $c_i = \rho|x'_i$ is the crossing sequence at x'_i . Finally, since $E_{L'}$ is idempotent, we have that $c_1 = c_2$ and this is equal to the crossing sequences of ρ at the positions x'_1, \dots, x'_{k+1} . This shows that L_1, L_2 are idempotent loops. \square

Turning back to the proof of the theorem, we know from the above claim that there are two consecutive idempotent loops $L_1 = [x, x']$ and $L_2 = [x', x'']$ with the same effect and with endpoints $x, x', x'' \in X_y \subseteq I \setminus \{x_1, x_2\}$ (see again Figure 15).

Let $\tilde{\ell}_1 = (x, y)$ and $\tilde{\ell}_2 = (x'', y)$, and observe that $\tilde{\ell}_1, \tilde{\ell}_2 \in Z'$. In particular, $\tilde{\ell}_1$ and $\tilde{\ell}_2$ are strictly between ℓ_1 and ℓ_2 . Suppose by symmetry that $\tilde{\ell}_1 \preceq \tilde{\ell}_2$. Further let C be the component of $L_1 \cup L_2$ (or, equally, of L_1 or L_2) that contains the node y . Below, we focus on

the factors of $\rho[\tilde{\ell}_1, \tilde{\ell}_2]$ that are intercepted by $L_1 \cup L_2$: these are represented in Figure 15 by the thick arrows. By Lemma 6.6 all these factors correspond to edges of the same component C , namely, they are $(L_1 \cup L_2, C)$ -factors.

Let us fix an arbitrary factor α of $\rho[\tilde{\ell}_1, \tilde{\ell}_2]$ that is intercepted by $L_1 \cup L_2$, and assume that $\alpha = \beta_1 \cdots \beta_k$, where β_1, \dots, β_k are the factors intercepted by either L_1 or L_2 .

Claim 7.4. If β, β' are two factors intercepted by $L_1 = [x, x']$ and $L_2 = [x', x'']$, with $E_{L_1} = E_{L_2} = E_{L_1 \cup L_2}$, and β, β' are adjacent in the run ρ (namely, they share an endpoint at position x'), then β, β' correspond to edges in the same component of L_1 (or, equally, L_2).

Proof. Let C be the component of L_1 and $y_1 \mapsto y_2$ the edge of C that corresponds to the factor β intercepted by L_1 . Similarly, let C' be the component of L_2 and $y_3 \mapsto y_4$ the edge of C' that corresponds to the factor β' intercepted by L_2 . Since β and β' share an endpoint at position x' , we know that $y_2 = y_3$. This shows that $C \cap C' \neq \emptyset$, and hence $C = C'$. \square

The above claim shows that any two adjacent factors β_i, β_{i+1} correspond to edges in the same component of L_1 and L_2 , respectively. Thus, by transitivity, all factors β_1, \dots, β_k correspond to edges in the same component, say C' . We claim that $C' = C$. Indeed, if β_1 is intercepted by L_1 , then $C' = C$ because α and β_1 start from the same location and hence they correspond to edges of the flow that depart from the same node. The other case is where β_1 is intercepted by L_2 , for which a symmetric argument can be applied.

So far we have shown that every factor of $\rho[\tilde{\ell}_1, \tilde{\ell}_2]$ intercepted by $L_1 \cup L_2$ can be factorized into some (L_1, C) -factors and some (L_2, C) -factors. We conclude the proof with the following observations:

- By construction, both loops L_1, L_2 are contained in the interval of positions $I = [x_1, x_2]$, and have endpoints different from x_1, x_2 .
- Both anchor points of C inside L_1, L_2 belong to the interval of locations $K \setminus \{\ell_1, \ell_2\}$. This holds because $\rho[\tilde{\ell}_1, \tilde{\ell}_2]$ contains a factor α that is intercepted by $L_1 \cup L_2$ and spans across all the positions from x to x'' , namely, an LR-factor. This factor starts at the anchor point of C inside L_1 and visits the anchor point of C inside L_2 . Moreover, by construction, α is also a factor of the subsequence $\rho|Z'$. This shows that the anchor points of C inside L_1 and L_2 belong to Z' , and in particular to $K \setminus \{\ell_1, \ell_2\}$.
- The first factor of $\rho[\tilde{\ell}_1, \tilde{\ell}_2]$ that is intercepted by $L_1 \cup L_2$ starts at $\tilde{\ell}_1 = (x, y)$, which by construction is the source location of some transition producing non-empty output. By the previous arguments, this factor is a concatenation of (L_1, C) -factors and (L_2, C) -factors. This implies that the trace of the anchor point of C inside L_1 , or the trace of C inside L_2 produces non-empty output. \square

Inversions and periodicity. The first important notion that is used to characterize one-way definability is that of inversion. It turns out that the definition of inversion in the sweeping case (see page 11) can be reused almost verbatim in the two-way setting. The only difference is that here we require the loops to be idempotent and we do not enforce output-minimality (we will discuss this latter choice further below, with a formal definition of output-minimality at hand).

Definition 7.5. An *inversion* of the run ρ is a tuple $(L_1, \ell_1, L_2, \ell_2)$ such that

- (1) L_1, L_2 are idempotent loops,
- (2) $\ell_1 = (x_1, y_1)$ and $\ell_2 = (x_2, y_2)$ are anchor points inside L_1 and L_2 , respectively,




Figure 16. An example of an inversion $(L_1, \ell_1, L_2, \ell_2)$ of a two-way run.

- (3) $\ell_1 \triangleleft \ell_2$ and $x_1 > x_2$,
- (4) for both $i = 1$ and $i = 2$, $\text{out}(\text{tr}(\ell_i)) \neq \varepsilon$.

Figure 16 gives an example of an inversion involving the idempotent loop L_1 with anchor point ℓ_1 , and the idempotent loop L_2 with anchor point ℓ_2 . The intercepted factors that form the corresponding traces are represented by thick arrows; the one highlighted in red are those that produce non-empty output.

The implication **P1** \Rightarrow **P2** of Theorem 3.6 in the two-way case is formalized below exactly as in Proposition 4.5, and the proof is very similar to the sweeping case. More precisely, it can be checked that the proof of the first claim in Proposition 4.5 was shown independently of the sweeping assumption — one just needs to replace the use of Equation 4.2 with Proposition 6.10. The sweeping assumption was used only for deriving the notion of *output-minimal* factor, which was crucial to conclude that the period p is bounded by the specific constant \mathbf{B} . In this respect, the proof of Proposition 7.6 requires a different argument for showing that $p \leq \mathbf{B}$:

Proposition 7.6. *If T is one-way definable, then the following property **P2** holds:*

For all inversions $(L_1, \ell_1, L_2, \ell_2)$ of ρ , the period p of the word

$$\text{out}(\text{tr}(\ell_1)) \text{ out}(\rho[\ell_1, \ell_2]) \text{ out}(\text{tr}(\ell_2))$$

divides both $|\text{out}(\text{tr}(\ell_1))|$ and $|\text{out}(\text{tr}(\ell_2))|$. Moreover, $p \leq \mathbf{B}$.

We only need to show here that $p \leq \mathbf{B}$. Recall that in the sweeping case we relied on the assumption that the factors $\text{tr}(\ell_1)$ and $\text{tr}(\ell_2)$ of an inversion are output-minimal, and on Lemma 4.3. In the general case we need to replace output-minimality by the following notion:

Definition 7.7. Consider pairs (L, C) consisting of an idempotent loop L and a component C of L .

- (1) On such pairs, define the relation \sqsubset by $(L', C') \sqsubset (L, C)$ if $L' \subsetneq L$ and at least one (L', C') -factor is contained in some (L, C) -factor.
- (2) A pair (L, C) is *output-minimal* if $(L', C') \sqsubset (L, C)$ implies $\text{out}(\text{tr}(\text{an}(C'))) = \varepsilon$.

Note that the relation \sqsubset is not a partial order in general (it is however antisymmetric). Moreover, it is easy to see that the notion of output-minimal pair (L, C) generalizes that of output-minimal factor introduced in the sweeping case: indeed, if ℓ is the anchor point of




Figure 17. An inversion $(L_1, \ell_1, L_2, \ell_2)$ with non output-minimal pairs (L_i, C_i) .

a loop L of a sweeping transducer and $\text{tr}(\ell)$ satisfies Definition 4.2, then the pair (L, C) is output-minimal, where C is the unique component whose edge corresponds to $\text{tr}(\ell)$.

The following lemma bounds the length of the output trace $\text{out}(\text{tr}(\text{an}(C)))$ for an output-minimal pair (L, C) :

Lemma 7.8. *For every output-minimal pair (L, C) , $|\text{out}(\text{tr}(\text{an}(C)))| \leq B$.*

Proof. Consider a pair (L, C) consisting of an idempotent loop $L = [x_1, x_2]$ and a component C of L . Suppose by contradiction that $|\text{out}(\text{tr}(\text{an}(C)))| > B$. We will show that (L, C) is not output-minimal.

Recall that $\text{tr}(\text{an}(C))$ is a concatenation of (L, C) -factors, say, $\text{tr}(\text{an}(C)) = \beta_1 \cdots \beta_k$. Let ℓ_1 (resp. ℓ_2) be the first (resp. last) location that is visited by these factors. Further let $K = [\ell_1, \ell_2]$ and $Z = K \cap (L \times \mathbb{N})$. By construction, the subsequence $\rho|Z$ can be seen as a concatenation of the factors β_1, \dots, β_k , possibly in a different order than that of $\text{tr}(\text{an}(C))$. This implies that $|\text{out}(\rho|Z)| > B$.

By Theorem 7.2, we know that there exist an idempotent loop $L' \subsetneq L$ and a component C' of L' such that $\text{an}(C') \in K$ and $\text{out}(\text{tr}(\text{an}(C'))) \neq \varepsilon$. Note that the (L', C') -factor that starts at the anchor point $\text{an}(C')$ (an LR- or RL-factor) is entirely contained in some (L, C) -factor. This implies that $(L', C') \subset (L, C)$, and thus (L, C) is not output-minimal. \square

We remark that the above lemma cannot be used directly to bound the period of the output produced amid an inversion. The reason is that we cannot restrict ourselves to inversions $(L_1, \ell_1, L_2, \ell_2)$ that induce output-minimal pairs (L_i, C_i) for $i = 1, 2$, where C_i is the unique component of the anchor point ℓ_i . An example is given in Figure 16, assuming that the factors depicted in red are the only ones that produce non-empty output, and the lengths of these outputs exceed B . On the one hand $(L_1, \ell_1, L_2, \ell_2)$ is an inversion, but (L_1, C_1) is not output-minimal. On the other hand, it is possible that ρ contains no other inversion than $(L_1, \ell_1, L_2, \ell_2)$: any loop strictly contained in the red factor in L_1 will have the anchor point *after* ℓ_2 .

We are now ready to show the second claim of Proposition 7.6.

Proof of Proposition 7.6. The proof of the second claim requires a refinement of the arguments that involve pumping the run ρ simultaneously on three different loops. As usual, we

assume that the loops L_1, L_2 of the inversion are disjoint (otherwise, we preliminarily pump one of the two loops a few times).

Recall that the word

$$\text{out}(\text{tr}(\ell_1)) \text{out}(\rho[\ell_1, \ell_2]) \text{out}(\text{tr}(\ell_2))$$

has period $p = \gcd(|\text{out}(\text{tr}(\ell_1))|, |\text{out}(\text{tr}(\ell_2))|)$, but that we cannot bound p by assuming that $(L_1, \ell_1, L_2, \ell_2)$ is output-minimal. However, in the pumped run $\rho^{(2,1)}$ we do find inversions with output-minimal pairs. For example, as depicted in the right part of Figure 17, we can consider the left and right copy of L_1 in $\rho^{(2,1)}$, denoted by \overleftarrow{L}_1 and \overrightarrow{L}_1 , respectively. Accordingly, we denote by $\overleftarrow{\ell}_1$ and $\overrightarrow{\ell}_1$ the left and right copy of ℓ_1 in $\rho^{(2,1)}$. Now, let (L_0, C_0) be any *output-minimal* pair such that L_0 is an idempotent loop, $\text{out}(\text{tr}(\text{an}(C_0))) \neq \varepsilon$, and either $(L_0, C_0) = (\overleftarrow{L}_1, C_1)$ or $(L_0, C_0) \sqsubset (\overleftarrow{L}_1, C_1)$. Such a loop L_0 is represented in Figure 17 by the red vertical stripe. Further let $\ell_0 = \text{an}(C_0)$.

We claim that either $(L_0, \ell_0, L_2, \ell_2)$ or $(\overrightarrow{L}_1, \overrightarrow{\ell}_1, L_0, \ell_0)$ is an inversion of the run $\rho^{(2,1)}$, depending on whether ℓ_0 occurs before or after ℓ_2 . First, note that all the loops $L_0, L_2, \overrightarrow{L}_1$ are idempotent and non-overlapping; more precisely, we have $\max(L_2) \leq \min(L_0)$ and $\max(L_0) \leq \min(\overrightarrow{L}_1)$. Moreover, the outputs of the traces $\text{tr}(\ell_0)$, $\text{tr}(\overrightarrow{\ell}_1)$, and $\text{tr}(\ell_2)$ are all non-empty. So it remains to distinguish two cases based on the ordering of the anchor points $\ell_0, \overrightarrow{\ell}_1, \ell_2$. If $\ell_0 \triangleleft \ell_2$, then $(L_0, \ell_0, L_2, \ell_2)$ is an inversion. Otherwise, because $(\overrightarrow{L}_1, \overrightarrow{\ell}_1, L_2, \ell_2)$ is an inversion, we know that $\overrightarrow{\ell}_1 \triangleleft \ell_2 \trianglelefteq \ell_0$, and hence $(\overrightarrow{L}_1, \overrightarrow{\ell}_1, L_0, C_0)$ is an inversion.

Now, we know that $\rho^{(2,1)}$ contains the inversion $(\overrightarrow{L}_1, \overrightarrow{\ell}_1, L_2, \ell_2)$, but also an inversion involving the output-minimal pair (L_0, C_0) , with L_0 strictly between \overrightarrow{L}_1 and L_2 . For all m_0, m_1, m_2 , we define $\rho^{(m_0, m_1, m_2)}$ as the run obtained from $\rho^{(2,1)}$ by pumping m_0, m_1, m_2 times the loops $L_0, \overrightarrow{L}_1, L_2$, respectively. By reasoning as we did in the proof of Proposition 4.5 (cf. *Periodicity of outputs of pumped runs*), one can show that there are arbitrarily large output factors of $\rho^{(m_0, m_1, m_2)}$ that are produced within the inversion on ℓ_0 (i.e. either $(L_0, \ell_0, L_2, \ell_2)$ or $(\overrightarrow{L}_1, \overrightarrow{\ell}_1, L_0, \ell_0)$) and that are periodic with period p' that divides $|\text{out}(\text{tr}(\ell_0))|$. In particular, by Lemma 7.8, we know that $p' \leq \mathbf{B}$. Moreover, large portions of these factors are also produced within the inversion $(\overrightarrow{L}_1, \overrightarrow{\ell}_1, L_2, \ell_2)$, and hence by Theorem 4.6 they have period $\gcd(p, p')$.

To conclude the proof we need to transfer the periodicity property from the pumped runs $\rho^{(m_0, m_1, m_2)}$ to the original run ρ . This is done exactly like in Proposition 4.5 by relying on Lemma 4.7: we observe that the periodicity property holds for large enough parameters m_0, m_1, m_2 , hence for all values of the parameters, and in particular for $m_0 = m_1 = m_2 = 1$. This shows that the word

$$\text{out}(\text{tr}(\ell_1)) \text{out}(\rho[\ell_1, \ell_2]) \text{out}(\text{tr}(\ell_2))$$

has period $\gcd(p, p') \leq \mathbf{B}$. □

So far we have shown that the output produced amid every inversion of a run of a one-way definable two-way transducer is periodic, with period bounded by \mathbf{B} and dividing the lengths of the trace outputs of the inversion. This basically proves the implication **P1** \Rightarrow **P2** of Theorem 3.6. In the next section we will follow a line of arguments similar to that of Section 5 to prove the remaining implications **P2** \Rightarrow **P3** \Rightarrow **P1**.

8. THE CHARACTERIZATION IN THE TWO-WAY CASE

In this section we generalize the characterization of one-way definability of sweeping transducers to the general two-way case. As usual, we fix through the rest of the section a successful run ρ of T on some input word u .

From periodicity of inversions to existence of decompositions. We continue by proving the second implication **P2** \Rightarrow **P3** of Theorem 3.6 in the two-way case. This requires showing the existence of a suitable decomposition of a run ρ that *satisfies* property **P2**. Recall that **P2** says that for every inversion $(L_1, \ell_1, L_2, \ell_2)$, the period of the word $\text{out}(\text{tr}(\ell_1)) \text{out}(\rho[\ell_1, \ell_2]) \text{out}(\text{tr}(\ell_2))$ divides $\text{gcd}(|\text{out}(\text{tr}(\ell_1))|, |\text{out}(\text{tr}(\ell_2))|) \leq B$. The definitions underlying the decomposition of ρ are similar to those given in the sweeping case:

Definition 8.1. Let $\rho[\ell, \ell']$ be a factor of a run ρ of T , where $\ell = (x, y)$, $\ell' = (x', y')$, and $x \leq x'$. We call $\rho[\ell, \ell']$

- a **B-diagonal** if for all $z \in [x, x']$, there is a location ℓ_z at position z such that $\ell \leq \ell_z \leq \ell'$ and the words $\text{out}(\rho|Z_{\ell_z}^{\leftarrow})$ and $\text{out}(\rho|Z_{\ell_z}^{\rightarrow})$ have length at most B , where $Z_{\ell_z}^{\leftarrow} = [\ell_z, \ell'] \cap ([0, z] \times \mathbb{N})$ and $Z_{\ell_z}^{\rightarrow} = [\ell, \ell_z] \cap ([z, \omega] \times \mathbb{N})$;
- a **B-block** if the word $\text{out}(\rho[\ell, \ell'])$ is almost periodic with bound B , and $\text{out}(\rho|Z^{\leftarrow})$ and $\text{out}(\rho|Z^{\rightarrow})$ have length at most B , where $Z^{\leftarrow} = [\ell, \ell'] \cap ([0, x] \times \mathbb{N})$ and $Z^{\rightarrow} = [\ell, \ell'] \cap ([x', \omega] \times \mathbb{N})$.




Figure 18. A diagonal.




Figure 19. A block.

The definition of **B**-decomposition is copied verbatim from the sweeping case, but uses the new notions of **B**-diagonal and **B**-block:

Definition 8.2. A **B**-decomposition of a run ρ of T is a factorization $\prod_i \rho[\ell_i, \ell_{i+1}]$ of ρ into **B**-diagonals and **B**-blocks.

To provide further intuition, we consider the transduction of Example 3.2 and the two-way transducer T that implements it in the most natural way. Figure 20 shows an example of a run of T on an input of the form $u_1 \# u_2 \# u_3 \# u_4$, where $u_2, u_4 \in (abc)^*$, $u_1, u_3 \notin (abc)^*$, and u_3 has even length. The factors of the run that produce long outputs are highlighted by the bold arrows. The first and third factors of the decomposition, i.e. $\rho[\ell_1, \ell_2]$ and $\rho[\ell_3, \ell_4]$, are diagonals (represented by the blue hatched areas); the second and fourth factors $\rho[\ell_2, \ell_3]$ and $\rho[\ell_4, \ell_5]$ are blocks (represented by the red hatched areas).

To identify the blocks of a possible decomposition of ρ , we reuse the equivalence relation \mathbf{S}^* introduced in Definition 5.4. Recall that this is the reflexive and transitive closure of the relation \mathbf{S} that groups any two locations ℓ, ℓ' that occur between ℓ_1, ℓ_2 , for some inversion $(L_1, \ell_1, L_2, \ell_2)$.




Figure 20. A decomposition of a run of a two-way transducer.

The proof that the output produced inside each \mathbf{S}^* -equivalence class is periodic, with period at most \mathbf{B} (Lemma 5.5) carries over in the two-way case without modifications. Similarly, every \mathbf{S}^* -equivalence class can be extended to the left and to the right by using Definition 5.7, which we report here verbatim for the sake of readability, together with an exemplifying figure.

Definition 8.3. Consider a non-singleton \mathbf{S}^* -equivalence class $K = [\ell, \ell']$. Let $\text{an}(K)$ be the restriction of K to the anchor points occurring in some inversion, and $X_{\text{an}(K)} = \{x : \exists y (x, y) \in \text{an}(K)\}$ be the projection of $\text{an}(K)$ on positions. We define $\text{block}(K) = [\tilde{\ell}, \tilde{\ell}']$, where

- $\tilde{\ell}$ is the latest location $(\tilde{x}, \tilde{y}) \preceq \ell$ such that $\tilde{x} = \min(X_{\text{an}(K)})$,
- $\tilde{\ell}'$ is the earliest location $(\tilde{x}, \tilde{y}) \succeq \ell'$ such that $\tilde{x} = \max(X_{\text{an}(K)})$

(note that the locations $\tilde{\ell}, \tilde{\ell}'$ exist since ℓ, ℓ' are anchor points in some inversion).




Figure 21. Block construction.

As usual, we call \mathbf{S}^* -block any factor of ρ of the form $\rho|_{\text{block}(K)}$ that is obtained by applying the above definition to a non-singleton \mathbf{S}^* -class K . Lemma 5.8, which shows that \mathbf{S}^* -blocks can indeed be used as \mathbf{B} -blocks in a decomposition of ρ , generalizes easily to the two-way case:

Lemma 8.4. *If K is a non-singleton \mathbf{S}^* -equivalence class, then $\rho|_{\text{block}(K)}$ is a \mathbf{B} -block.*

Proof. The proof is similar to that of Lemma 5.8. The main difference is that here we will bound the lengths of some outputs using a Ramsey-type argument (Theorem 7.2), instead of output-minimality of factors (Lemma 4.3). To follow the various constructions and arguments the reader can refer to Figure 21.

Let $K = [\ell, \ell']$, $\text{an}(K)$, $X_{\text{an}(K)}$, and $\text{block}(K) = [\tilde{\ell}, \tilde{\ell}']$ be as in Definition 8.3, where $\tilde{\ell} = (\tilde{x}, \tilde{y})$, $\tilde{\ell}' = (\tilde{x}', \tilde{y}')$, $\tilde{x} = \min(X_{\text{an}(K)})$, and $\tilde{x}' = \max(X_{\text{an}(K)})$. We need to verify that $\rho|_{\text{block}(K)}$ is a \mathbf{B} -block, namely, that:

- $\tilde{x} \leq \tilde{x}'$,
- $\text{out}(\rho|_{[\tilde{\ell}, \tilde{\ell}']})$ is almost periodic with bound \mathbf{B} ,

- $\text{out}(\rho|Z^-)$ and $\text{out}(\rho|Z^+)$ have length at most \mathbf{B} , where $Z^- = [\tilde{\ell}, \tilde{\ell}'] \cap ([0, x] \times \mathbb{N})$ and $Z^+ = [\tilde{\ell}, \tilde{\ell}'] \cap ([x', \omega] \times \mathbb{N})$.

The first condition $\tilde{x} \leq \tilde{x}'$ follows immediately from $\tilde{x} = \min(X_{\text{an}(K)})$ and $\tilde{x}' = \max(X_{\text{an}(K)})$.

Next, we prove that the output produced by the factor $\rho[\tilde{\ell}, \tilde{\ell}']$ is almost periodic with bound \mathbf{B} . By Definition 8.3, we have $\tilde{\ell} \trianglelefteq \ell \triangleleft \ell' \trianglelefteq \tilde{\ell}'$, and by Lemma 5.5 we know that $\text{out}(\rho[\ell, \ell'])$ is periodic with period at most \mathbf{B} . So it suffices to bound the length of the words $\text{out}(\rho[\tilde{\ell}, \ell])$ and $\text{out}(\rho[\ell', \tilde{\ell}'])$. We shall focus on the former word, as the arguments for the latter are similar.

First, we show that the factor $\rho[\tilde{\ell}, \ell]$ lies entirely to the right of position \tilde{x} (in particular, it starts at an even level \tilde{y}). Indeed, if this were not the case, there would exist another location $\ell'' = (\tilde{x}, \tilde{y} + 1)$, on the same position as $\tilde{\ell}$, but at a higher level, such that $\tilde{\ell} \triangleleft \ell'' \trianglelefteq \ell$. But this would contradict Definition 8.3 ($\tilde{\ell}$ is the *latest* location $(x, y) \trianglelefteq \ell$ such that $x = \min(X_{\text{an}(K)})$).

Suppose now that the length of $|\text{out}(\rho[\tilde{\ell}, \ell])| > \mathbf{B}$. We head towards a contradiction by finding a location $\ell'' \triangleleft \ell$ that is \mathbf{S}^* -equivalent to the first location ℓ of the \mathbf{S}^* -equivalence class K . Since the factor $\rho[\tilde{\ell}, \ell]$ lies entirely to the right of position \tilde{x} , it is intercepted by the interval $I = [\tilde{x}, \omega]$. So $|\text{out}(\rho[\tilde{\ell}, \ell])| > \mathbf{B}$ is equivalent to saying $|\text{out}(\rho|Z)| > \mathbf{B}$, where $Z = [\tilde{\ell}, \ell] \cap ([\tilde{x}, \omega] \times \mathbb{N})$. Then, Theorem 7.2 implies the existence of an idempotent loop L and an anchor point ℓ'' of L such that

- $\min(L) > \tilde{x}$,
- $\tilde{\ell} \triangleleft \ell'' \triangleleft \ell$,
- $\text{out}(\text{tr}(\ell'')) \neq \varepsilon$.

Further recall that $\tilde{x} = \min(X_{\text{an}(K)})$ is the leftmost position of locations in the class $K = [\ell, \ell']$ that are also anchor points of inversions. In particular, there is an inversion $(L_1, \ell_1'', L_2, \ell_2'')$, with $\ell_2'' = (\tilde{x}, y_2'') \in K$. Since $\ell'' \triangleleft \ell \trianglelefteq \ell_2''$ and the position of ℓ'' is to the right of the position of ℓ_2'' , we know that $(L, \ell'', L_2, \ell_2'')$ is also an inversion, and hence $\ell'' \mathbf{S}^* \ell_2'' \mathbf{S}^* \ell$. But since $\ell'' \neq \ell$, we get a contradiction with the assumption that ℓ is the first location of the \mathbf{S}^* -class K . In this way we have shown that $|\text{out}(\rho[\tilde{\ell}, \ell])| \leq \mathbf{B}$.

It remains to bound the lengths of the outputs produced by the subsequences $\rho|Z^-$ and $\rho|Z^+$, where $Z^- = [\tilde{\ell}, \tilde{\ell}'] \cap ([0, \tilde{x}] \times \mathbb{N})$ and $Z^+ = [\tilde{\ell}, \tilde{\ell}'] \cap ([\tilde{x}', \omega] \times \mathbb{N})$. As usual, we consider only one of the two symmetric cases. Suppose, by way of contradiction, that $|\text{out}(\rho|Z^-)| > \mathbf{B}$. By Theorem 7.2, there exist an idempotent loop L and an anchor point ℓ'' of L such that

- $\max(L) < \tilde{x}$,
- $\tilde{\ell} \triangleleft \ell'' \triangleleft \tilde{\ell}'$,
- $\text{out}(\text{tr}(\ell'')) \neq \varepsilon$.

By following the same line of reasoning as before, we recall that ℓ is the first location of the non-singleton class K . From this we derive the existence an inversion $(L_1, \ell_1'', L_2, \ell_2'')$ where $\ell_1'' = \ell$. We claim that $\ell \trianglelefteq \ell''$. Indeed, if this were not the case, then, because ℓ'' is strictly to the left of \tilde{x} and ℓ is to the right of \tilde{x} , there would exist a location ℓ''' between ℓ'' and ℓ that lies at position \tilde{x} . But $\tilde{\ell} \triangleleft \ell'' \trianglelefteq \ell''' \triangleleft \ell$ would contradict the fact that $\tilde{\ell}$ is the *latest* location before ℓ that lies at the position \tilde{x} . Now that we know that $\ell \trianglelefteq \ell''$ and that ℓ'' is to the left of \tilde{x} , we observe that $(L_1, \ell_1'', L, \ell'')$ is also an inversion, and hence $\ell'' \in \text{an}(K)$. Since ℓ'' is strictly to the left of \tilde{x} , we get a contradiction with the definition of \tilde{x} as leftmost position of the locations in $\text{an}(K)$. So we conclude that $|\text{out}(\rho|Z^-)| \leq \mathbf{B}$. \square

The proof of Lemma 5.9, which shows that \mathbf{S}^* -blocks do not overlap along the input axis, carries over in the two-way case, again without modifications. Finally, we generalize Lemma 5.10 to the new definition of diagonal, which completes the construction of a \mathbf{B} -decomposition for the run ρ :

Lemma 8.5. *Let $\rho[\ell, \ell']$ be a factor of ρ , with $\ell = (x, y)$, $\ell' = (x', y')$, and $x \leq x'$, that does not overlap any \mathbf{S}^* -block. Then $\rho[\ell, \ell']$ is a \mathbf{B} -diagonal.*

Proof. Suppose by way of contradiction that there is some $z \in [x, x']$ such that, for all locations ℓ'' at position z and between ℓ and ℓ' , one of the two conditions holds:

- (1) $|\text{out}(\rho|Z_{\ell''}^{\leftarrow})| > \mathbf{B}$, where $Z_{\ell''}^{\leftarrow} = [\ell'', \ell'] \cap ([0, z] \times \mathbb{N})$,
- (2) $|\text{out}(\rho|Z_{\ell''}^{\leftarrow})| > \mathbf{B}$, where $Z_{\ell''}^{\leftarrow} = [\ell, \ell''] \cap ([z, \omega] \times \mathbb{N})$.

First, we claim that *each* of the two conditions above are satisfied at some locations $\ell'' \in [\ell, \ell']$ at position z . Consider the highest even level y'' such that $\ell'' = (z, y'') \in [\ell, \ell']$ (use Figure 18 as a reference). Since $z \leq x'$, the outgoing transition at ℓ'' is rightward oriented, and the set $Z_{\ell''}^{\leftarrow}$ is empty. This means that condition (1) is trivially violated at ℓ'' , and hence condition (2) holds at ℓ'' by the initial assumption. Symmetrically, condition (1) holds at the location $\ell'' = (z, y'')$, where y'' is the lowest even level with $\ell'' \in [\ell, \ell']$.

Let us now compare the levels where the above conditions hold. Clearly, the lower the level of location ℓ'' , the easier it is to satisfy condition (1), and symmetrically for condition (2). So, let $\ell^+ = (z, y^+)$ (resp. $\ell^- = (z, y^-)$) be the highest (resp. lowest) location in $[\ell, \ell']$ at position z that satisfies condition (1) (resp. condition (2)).

We claim that $y^+ \geq y^-$. For this, we first observe that $y^+ \geq y^- - 1$, since otherwise there would exist a location $\ell'' = (z, y'')$, with $y^+ < y'' < y^-$, that violates both conditions (1) and (2). Moreover, y^+ must be odd, otherwise the transition departing from $\ell^+ = (z, y^+)$ would be rightward oriented and the location $\ell'' = (z, y^+ + 1)$ would still satisfy condition (1), contradicting the definition of highest location ℓ^+ . For similar reasons, y^- must also be odd, otherwise there would be a location $\ell'' = (z, y^- - 1)$ below ℓ^- that satisfies condition (2). But since $y^+ \geq y^- - 1$ and both y^+ and y^- are odd, we need to have $y^+ \geq y^-$.

In fact, from the previous arguments we know that the location $\ell'' = (z, y^+)$ (or equally the location (x, y^-)) satisfies *both* conditions (1) and (2). We can thus apply Theorem 7.2 to the sets $Z_{\ell''}^{\leftarrow}$ and $Z_{\ell''}^{\leftarrow}$, deriving the existence of two idempotent loops L_1, L_2 and two anchor points ℓ_1, ℓ_2 of L_1, L_2 , respectively, such that

- $\max(L_2) < z < \min(L_1)$,
- $\ell \triangleleft \ell_1 \triangleleft \ell'' \triangleleft \ell_2 \triangleleft \ell'$,
- $\text{out}(\text{tr}(\ell_1)), \text{out}(\text{tr}(\ell_2)) \neq \varepsilon$.

In particular, since ℓ_1 is to the right of ℓ_2 w.r.t. the order of positions, we know that $(L_1, \ell_1, L_2, \ell_2)$ is an inversion, and hence $\ell_1 \mathbf{S}^* \ell_2$. But this contradicts the assumption that $\rho[\ell, \ell']$ does not overlap with any \mathbf{S}^* -block. \square

From existence of decompositions to an equivalent one-way transducer. It remains to prove the last implication $\mathbf{P3} \Rightarrow \mathbf{P1}$ of Theorem 3.6, which amounts to construct a one-way transducer T' equivalent to T .

Hereafter, we denote by D the language of words $u \in \text{dom}(T)$ such that *all* successful runs of T on u admit a \mathbf{B} -decomposition. So far, we know that if T is one-way definable ($\mathbf{P1}$), then $D = \text{dom}(T)$ ($\mathbf{P3}$). As a matter of fact, this reduces the one-way definability

problem for T to the containment problem $\text{dom}(T) \subseteq D$. We will see later (in Section 9) how the latter problem can be decided in double exponential space by further reducing it to checking the emptiness of the intersection of the languages $\text{dom}(T)$ and D^c , where D^c is the complement of D .

Below, we show how to construct a one-way transducer T' of triple exponential size such that $T' \subseteq T$ and $\text{dom}(T')$ is the set of all input words that have some successful run admitting a \mathbf{B} -decomposition (hence $\text{dom}(T') \supseteq D$). In particular, we will have that

$$T|_D \subseteq T' \subseteq T.$$

Note that this will prove **P3** to **P1**, as well as the second item of Theorem 3.3, since $D = \text{dom}(T)$ if and only if T is one-way definable. A sketch of the proof of this construction when T is a sweeping transducer was given at the end of Section 5.

Proposition 8.6. *Given a functional two-way transducer T , a one-way transducer T' satisfying*

$$T' \subseteq T \quad \text{and} \quad \text{dom}(T') \supseteq D$$

can be constructed in 3EXPTIME. Moreover, if T is sweeping, then T' can be constructed in 2EXPTIME.

Proof. Given an input word u , the transducer T' will guess (and check) a successful run ρ of T on u , together with a \mathbf{B} -decomposition $\prod_i \rho[\ell_i, \ell_{i+1}]$. The latter decomposition will be used by T' to simulate the output of ρ in left-to-right manner, thus proving that $T' \subseteq T$. Moreover, $u \in D$ implies the existence of a successful run that can be decomposed, thus proving that $\text{dom}(T') \supseteq D$. We now provide the details of the construction of T' .

Guessing the run ρ is standard (see, for instance, [24, 33]): it amounts to guess the crossing sequences $\rho|x$ for each position x of the input. Recall that this is a bounded amount of information for each position x , since the run is normalized. As concerns the decomposition of ρ , it can be encoded by the endpoints ℓ_i of its factors, that is, by annotating the position of each ℓ_i as the level of ℓ_i . In a similar way T' guesses the information of whether each factor $\rho[\ell_i, \ell_{i+1}]$ is a \mathbf{B} -diagonal or a \mathbf{B} -block.

Thanks to the definition of decomposition (see Definition 8.2 and Figure 20), every two distinct factors span across non-overlapping intervals of positions. This means that each position x is covered by exactly one factor of the decomposition. We call this factor the *active factor at position x* . The mode of computation of the transducer will depend on the type of active factor: if the active factor is a diagonal (resp. a block), then we say that T' is in *diagonal mode* (resp. *block mode*). Below we describe the behaviour for these two modes of computation.

Diagonal mode. We recall the key condition satisfied by the diagonal $\rho[\ell, \ell']$ that is active at position x (cf. Definition 8.1 and Figure 18): there is a location $\ell_x = (x, y_x)$ between ℓ and ℓ' such that the words $\text{out}(\rho|Z_{\ell_x}^{\leftarrow})$ and $\text{out}(\rho|Z_{\ell_x}^{\rightarrow})$ have length at most \mathbf{B} , where $Z_{\ell_x}^{\leftarrow} = [\ell_x, \ell'] \cap ([0, x] \times \mathbb{N})$ and $Z_{\ell_x}^{\rightarrow} = [\ell, \ell_x] \cap ([x, \omega] \times \mathbb{N})$.

Besides the run ρ and the decomposition, the transducer T' will also guess the locations $\ell_x = (x, y_x)$, that is, will annotate each x with the corresponding y_x . Without loss of generality, we can assume that the function that associates each position x with the guessed location $\ell_x = (x, y_x)$ is monotone, namely, $x \leq x'$ implies $\ell_x \leq \ell_{x'}$. While the transducer T' is in diagonal mode, the goal is to preserve the following invariant:

After reaching a position x covered by the active diagonal, T' must have produced the output of ρ up to location ℓ_x .

To preserve the above invariant when moving from x to the next position $x+1$, the transducer should output the word $\text{out}(\rho[\ell_x, \ell_{x+1}])$. This word consists of the following parts:

- (1) The words produced by the single transitions of $\rho[\ell_x, \ell_{x+1}]$ with endpoints in $\{x, x+1\} \times \mathbb{N}$. Note that there are at most \mathbf{H} such words, each of them has length at most \mathbf{C} , and they can all be determined using the crossing sequences at x and $x+1$ and the information about the levels of ℓ_x and ℓ_{x+1} . We can thus assume that this information is readily available to the transducer.
- (2) The words produced by the factors of $\rho[\ell_x, \ell_{x+1}]$ that are intercepted by the interval $[0, x]$. Thanks to the definition of diagonal, we know that the total length of these words is at most \mathbf{B} . These words cannot be determined from the information on $\rho|x$, $\rho|x+1$, ℓ_x , and ℓ_{x+1} alone, so they need to be constructed while scanning the input. For this, some additional information needs to be stored.

More precisely, at each position x of the input, the transducer stores all the outputs produced by the factors of ρ that are intercepted by $[0, x]$ and that occur *after* a location of the form $\ell_{x'}$, for any $x' \geq x$ that is covered by a diagonal. This clearly includes the previous words when $x' = x$, but also other words that might be used later for processing other diagonals. Moreover, by exploiting the properties of diagonals, one can prove that those words have length at most \mathbf{B} , so they can be stored with triply exponentially many states. Using classical techniques, the stored information can be maintained while scanning the input u using the guessed crossing sequences of ρ .

- (3) The words produced by the factors of $\rho[\ell_x, \ell_{x+1}]$ that are intercepted by the interval $[x+1, \omega]$. These words must be guessed, since they depend on a portion of the input that has not been processed yet. Accordingly, the guesses need to be stored into memory, in such a way that they can be checked later. For this, the transducer stores, for each position x , the guessed words that correspond to the outputs produced by the factors of ρ intercepted by $[x, \omega]$ and occurring *before* a location of the form $\ell_{x'}$, for any $x' \leq x$ that is covered by a diagonal.

Block mode. Suppose that the active factor $\rho[\ell, \ell']$ is a \mathbf{B} -block. Let $I = [x, x']$ be the set of positions covered by this factor. Moreover, for each position $z \in I$, let $Z_z^- = [\ell, \ell'] \cap ([0, z] \times \mathbb{N})$ and $Z_z^+ = [\ell, \ell'] \cap ([z, \omega] \times \mathbb{N})$. We recall the key property of a block (cf. Definition 8.1 and Figure 19): the word $\text{out}(\rho[\ell, \ell'])$ is almost periodic with bound \mathbf{B} , and the words $\text{out}(\rho|Z_x^-)$ and $\text{out}(\rho|Z_{x'}^+)$ have length at most \mathbf{B} .

For the sake of brevity, suppose that $\text{out}(\rho[\ell, \ell']) = w_1 w_2 w_3$, where w_2 is periodic with period \mathbf{B} and w_1, w_3 have length at most \mathbf{B} . Similarly, let $w_0 = \text{out}(\rho|Z_x^-)$ and $w_4 = \text{out}(\rho|Z_{x'}^+)$. The invariant preserved by T' in block mode is the following:

After reaching a position z covered by the active block $\rho[\ell, \ell']$, T' must have produced the output of the prefix of ρ up to location ℓ , followed by a prefix of $\text{out}(\rho[\ell, \ell']) = w_1 w_2 w_3$ of the same length as $\text{out}(\rho|Z_z^-)$.

The initialization of the invariant is done when reaching the left endpoint x . At this moment, it suffices that T' outputs a prefix of $w_1 w_2 w_3$ of the same length as $w_0 = \text{out}(\rho|Z_x^-)$, thus bounded by \mathbf{B} . Symmetrically, when reaching the right endpoint x' , T' will have produced almost the entire word $\text{out}(\rho[\ell, \ell']) w_1 w_2 w_3$, but without the suffix $w_4 = \text{out}(\rho|Z_{x'}^+)$ of length at most \mathbf{B} . Thus, before moving to the next factor of the decomposition, the transducer will produce the remaining suffix, so as to complete the output of ρ up to location ℓ_{i_x+1} .

It remains to describe how the above invariant can be maintained when moving from a position z to the next position $z + 1$ inside $I = [x, x']$. For this, it is convenient to succinctly represent the word w_2 by its repeating pattern, say v , of length at most \mathbf{B} . To determine the symbols that have to be output at each step, the transducer will maintain a pointer on either $w_1 v$ or w_3 . The pointer is increased in a deterministic way, and precisely by the amount $|\text{out}(\rho|Z_{z+1}^-)| - |\text{out}(\rho|Z_z^-)|$. The only exception is when the pointer lies in $w_1 v$, but its increase would go over $w_1 v$: in this case the transducer has the choice to either bring the pointer back to the beginning of v (representing a periodic output inside w_2), or move it to w_3 . Of course, this is a non-deterministic choice, but it can be validated when reaching the right endpoint of I . Concerning the number of symbols that need to be emitted at each step, this can be determined from the crossing sequences at z and $z + 1$, and from the knowledge of the lowest and highest levels of locations that are at position z and between ℓ and ℓ' . We denote the latter levels by y_z^- and y_z^+ , respectively.

Overall, this shows how to maintain the invariant of the block mode, assuming that the levels y_z^-, y_z^+ are known, as well as the words w_0, w_1, v, w_3, w_4 of bounded length. Like the mapping $z \mapsto \ell_z = (z, y_z)$ used in diagonal mode, the mapping $z \mapsto (y_z^-, y_z^+)$ can be guessed and checked using the crossing sequences. Similarly, the words w_1, v, w_3 can be guessed just before entering the active block, and can be checked along the process. As concerns the words w_0, w_4 , these can be guessed and checked in a way similar to the words that we used in diagonal mode. More precisely, for each position z of the input, the transducer stores the following additional information:

- (1) the outputs produced by the factors of ρ that are intercepted by $[0, z]$ and that occur after the beginning ℓ'' of some block, with $\ell'' = (x'', y'')$ and $x'' \geq z$;
- (2) the outputs produced by the factors of ρ that are intercepted by $[z, \omega]$ and that occur before the ending ℓ''' of a block, where $\ell''' = (x''', y''')$ and $x''' \leq z$.

By the definition of blocks, the above words have length at most \mathbf{B} and can be maintained while processing the input and the crossing sequences. Finally, we observe that the words, together with the information given by the lowest and highest levels y_z^-, y_z^+ , for both $z = x$ and $z = x'$, are sufficient for determining the content of w_0 and w_4 .

We have just shown how to construct a one-way transducer $T' \subseteq T$ such that $\text{dom}(T') \supseteq D$. From the above construction it is easy to see that the number of states and transitions of T' , as well as the number of letters emitted by each transition, are at most exponential in \mathbf{B} . Since \mathbf{B} is doubly exponential in the size of T , this shows that T' can be constructed from T in 3EXPTIME. Note that the triple exponential complexity comes from the lengths of the words that need to be guessed and stored in the control states, and these lengths are bounded by \mathbf{B} . However, if T is a sweeping transducer, then, according to the results proved in Section 5, the bound \mathbf{B} is simply exponential. In particular, in the sweeping case we can construct the one-way transducer T' in 2EXPTIME. \square

Generality of the construction. We conclude the section with a discussion on the properties of the one-way transducer T' constructed from T . Roughly speaking, we would like to show that, even when T is not one-way definable, T' is somehow the *best one-way under-approximation* of T . However, strictly speaking, the latter terminology is meaningless: if T' is a one-way transducer strictly contained in T , then one can always find a better one-way transducer T'' that satisfies $T' \subsetneq T'' \subsetneq T$, for instance by extending T' with a single

input-output pair. Below, we formalize in an appropriate way the notion of “best one-way under-approximation”.

We are interested in comparing the domains of transducers, but only up to a certain amount. In particular, we are interested in languages that are preserved under pumping loops of runs of T . Formally, given a language L , we say that L is T -pumpable if $L \subseteq \text{dom}(T)$ and for all words $u \in L$, all successful runs ρ of T on u , all loops L of ρ , and all positive numbers n , the word $\text{pump}_L^n(u)$ also belongs to L . Clearly, the domain $\text{dom}(T)$ of a transducer T is a regular T -pumpable language.

Another noticeable example of T -pumpable regular language is the domain of the one-way transducer T' , as defined in Proposition 8.6. Indeed, $\text{dom}(T')$ consists of words $u \in \text{dom}(T)$ that induce successful runs with \mathbf{B} -decompositions, and the property of having a \mathbf{B} -decomposition is preserved under pumping.

The following result shows that T' is the best under-approximation of T within the class of one-way transducers with T -pumpable domains:

Corollary 8.7. *Given a functional two-way transducer T , one can construct a one-way transducer T' such that*

- $T' \subseteq T$ and $\text{dom}(T')$ is T -pumpable,
- for all one-way transducers T'' , if $T'' \subseteq T$ and $\text{dom}(T'')$ is T -pumpable, then $T'' \subseteq T'$.

Proof. The transducer T' is precisely the one defined in Proposition 8.6. As already explained, its domain $\text{dom}(T')$ is a T -pumpable language. In particular, T' satisfies the conditions in the first item.

For the conditions in the second item, consider a one-way transducer $T'' \subseteq T$ with a T -pumpable domain $L = \text{dom}(T'')$. Let \tilde{T} be the transducer obtained from T by restricting its domain to L . Clearly, \tilde{T} is one-way definable, and one could apply Proposition 7.6 to \tilde{T} , using T'' as a witness of one-way definability. In particular, when it comes to comparing the outputs of the pumped runs of \tilde{T} and T'' , one could exploit the fact that the domain L of T'' , and hence the domain of \tilde{T} as well, is T -pumpable. This permits to derive periodicities of inversions with the same bound \mathbf{B} as before, but only restricted to the successful runs of T on the input words that belong to L . As a consequence, one can define \mathbf{B} -decompositions of successful runs of T on words in L , thus showing that $L \subseteq \text{dom}(T')$. This proves that $T'' \subseteq T'$. \square

9. COMPLEXITY OF THE ONE-WAY DEFINABILITY PROBLEM

In this section we analyze the complexity of the problem of deciding whether a transducer T is one-way definable. We begin with the case of a functional two-way transducer. In this case, thanks to the results presented in Section 8 page 41, we know that T is one-way definable if and only if $\text{dom}(T) \subseteq D$, where D is the language of words $u \in \text{dom}(T)$ such that all successful runs of T on u admit a \mathbf{B} -decomposition. In particular, the one-way definability problem reduces to an emptiness problem for the intersection of two languages:

$$T \text{ one-way definable} \quad \text{if and only if} \quad \text{dom}(T) \cap D^c = \emptyset.$$

The following lemma exploits the characterization of Theorem 3.6 to show that the language D^c can be recognized by a non-deterministic finite automaton A of triply exponential size w.r.t. T . In fact, this lemma shows that the automaton recognizing D^c can be constructed

using doubly exponential workspace. As before, we gain an exponent when restricting to sweeping transducers.

Lemma 9.1. *Given a functional two-way transducer T , an NFA A recognizing D^c can be constructed in 2EXPSpace . Moreover, when T is sweeping, the NFA A can be constructed in EXPSpace .*

Proof. Consider an input word u . By Theorem 3.6 we know that $u \in D^c$ iff there exist a successful run ρ of T on u and an inversion $\mathcal{I} = (L_1, \ell_1, L_2, \ell_2)$ of ρ such that no positive number $p \leq \mathbf{B}$ is a period of the word

$$w_{\rho, \mathcal{I}} = \text{out}(\text{tr}(\ell_1)) \text{out}(\rho[\ell_1, \ell_2]) \text{out}(\text{tr}(\ell_2)).$$

The latter condition on $w_{\rho, \mathcal{I}}$ can be rephrased as follows: there is a function $f : \{1, \dots, \mathbf{B}\} \rightarrow \{1, \dots, |w_{\rho, \mathcal{I}}|\}$ such that $w_{\rho, \mathcal{I}}(f(p)) \neq w_{\rho, \mathcal{I}}(f(p) + p)$ for all positive numbers $p \leq \mathbf{B}$. In particular, each of the images of the latter function f , that is, $f(1), \dots, f(\mathbf{B})$, can be encoded by a suitable marking of the crossing sequences of ρ . This shows that the run ρ , the inversion \mathcal{I} , and the function f described above can all be guessed within space $\mathcal{O}(\mathbf{B})$: ρ is guessed on-the-fly, the inversion is guessed by marking the anchor points, and for f we only store two symbols and a counter $\leq \mathbf{B}$, for each $1 \leq i \leq \mathbf{B}$. That is, any state of A requires doubly exponential space, resp. simply exponential space, depending on whether T is arbitrary two-way or sweeping. \square

As a consequence of the previous lemma, the emptiness problem for the language $\text{dom}(T) \cap D^c$, and thus the one-way definability problem for T , can be decided in 2EXPSpace or EXPSpace , depending on whether T is two-way or sweeping:

Proposition 9.2. *The problem of deciding whether a functional two-way transducer T is one-way definable is in 2EXPSpace . When T is sweeping, the problem is in EXPSpace .*

The last result of the section shows that functional two-way transducers are close to be the largest class for which a characterization of one-way definability is feasible: as soon as we consider arbitrary transducers (including non-functional ones), the problem becomes undecidable.

Proposition 9.3. *The one-way definability problem for non-functional sweeping transducers is undecidable.*

Proof. The proof uses some ideas and variants of constructions provided in [25], concerning the proof of undecidability of the equivalence problem for one-way non-functional transducers.

We show a reduction from the Post Correspondence Problem (PCP). A *PCP instance* is described by two finite alphabets Σ and Δ and two morphisms $f, g : \Sigma^* \rightarrow \Delta^*$. A *solution* of such an instance is any non-empty word $w \in \Sigma^+$ such that $f(w) = g(w)$. We recall that the problem of testing whether a PCP instance has a solution is undecidable.

Below, we fix a tuple $\tau = (\Sigma, \Delta, f, g)$ describing a PCP instance and we show how to reduce the problem of testing the *non-existence of solutions* of τ to the problem of deciding *one-way definability* of a relation computed by a sweeping transducer. Roughly, the idea is to construct a relation B_τ between words over a suitable alphabet Γ that encodes all the *non-solutions* to the PCP instance τ (this is simpler than encoding solutions because the presence of errors can be easily checked). The goal is to have a relation B_τ that (i) can be computed by a sweeping transducer and (ii) coincides with a trivial one-way definable relation when τ has no solution.

We begin by describing the encodings for the solutions of the PCP instance. We assume that the two alphabets of the PCP instance, Σ and Δ , are disjoint and we use a fresh symbol $\# \notin \Sigma \cup \Delta$. We define the new alphabet $\Gamma = \Sigma \cup \Delta \cup \{\#\}$ that will serve both as input alphabet and as output alphabet for the transduction. We call *encoding* any pair of words over Γ of the form $(w \cdot u, w \cdot v)$, where $w \in \Sigma^+$, $u \in \Delta^*$, and $v \in \{\#\}^*$. We will write the encodings as vectors to improve readability, e.g., as

$$\begin{pmatrix} w \cdot u \\ w \cdot v \end{pmatrix}.$$

We denote by E_τ the set of all encodings and we observe that E_τ is computable by a one-way transducer (note that this transducer needs ε -transitions). We then restrict our attention to the pairs in E_τ that are encodings of valid solutions of the PCP instance. Formally, we call *good encodings* the pairs in E_τ of the form

$$\begin{pmatrix} w \cdot u \\ w \cdot \#^{|u|} \end{pmatrix} \quad \text{where } u = f(w) = g(w).$$

All the other pairs in E_τ are called *bad encodings*. Of course, the relation that contains the good encodings is not computable by a transducer. On the other hand, we can show that the complement of this relation w.r.t. E_τ is computable by a sweeping transducer. Let B_τ be the set of all bad encodings. Consider $(w \cdot u, w \cdot \#^m) \in E_\tau$, with $w \in \Sigma^+$, $u \in \Delta^*$, and $m \in \mathbb{N}$, and we observe that this pair belongs to B_τ if and only if one of the following conditions is satisfied:

- (1) $m < |u|$,
- (2) $m > |u|$,
- (3) $u \neq f(w)$,
- (4) $u \neq g(w)$.

We explain how to construct a sweeping transducer \mathcal{S}_τ that computes B_τ . Essentially, \mathcal{S}_τ guesses which of the above conditions holds and processes the input accordingly. More precisely, if \mathcal{S}_τ guesses that the first condition holds, then it performs a single left-to-right pass, first copying the prefix w to the output and then producing a block of occurrences of the symbol $\#$ that is shorter than the suffix u . This task can be easily performed while reading u : it suffices to emit at most one occurrence of $\#$ for each position in u , and at the same time guarantee that, for at least one such position, no occurrence of $\#$ is emitted. The second condition can be dealt with by a similar strategy: first copy the prefix w , then output a block of $\#$ that is longer than the suffix u . To deal with the third condition, the transducer \mathcal{S}_τ has to perform two left-to-right passes, interleaved by a backward pass that brings the head back to the initial position. During the first left-to-right pass, \mathcal{S}_τ copies the prefix w to the output. During the second left-to-right pass, it reads again the prefix w , but this time he guesses a factorization of it of the form $w_1 a w_2$. On reading w_1 , \mathcal{S}_τ will output $\#^{|f(w_1)|}$. After reading w_1 , \mathcal{S}_τ will store the symbol a and move to the position where the suffix u begins. From there, it will guess a factorization of u of the form $u_1 u_2$, check that u_2 does not begin with $f(a)$, and emit one occurrence of $\#$ for each position in u_2 . The number of occurrences of $\#$ produced in the output is thus $m = |f(w_1)| + |u_2|$, and the fact that u_2 does not begin with $f(a)$ ensures that the factorizations of w and u do not match, i.e.

$$m \neq |f(w)|$$

Note that the described behaviour does not immediately guarantee that $u \neq f(w)$. Indeed, it may still happen that $u = f(w)$, but as a consequence $m \neq |u|$. This case is already

covered by the first and second condition, so the computation is still correct in the sense that it produces only bad encodings. On the other hand, if m happens to be the same as $|u|$, then $|u| = m \neq |f(w)|$ and thus $u \neq f(w)$. A similar behaviour can be used to deal with the fourth condition.

We have just shown that there is a sweeping non-functional transducer \mathcal{S}_τ that computes the relation B_τ containing all the bad encodings. Note that, if the PCP instance τ admits no solution, then all encodings are bad, i.e., $B_\tau = E_\tau$, and hence B_τ is one-way definable. It remains to show that when τ has a solution, B_τ is not one-way definable. Suppose that τ has solution $w \in \Sigma^+$ and let $(w \cdot u, w \cdot \#^{|u|})$ be the corresponding good encoding, where $u = f(w) = g(w)$. Note that every exact repetition of w is also a solution, and hence the pairs $(w^n \cdot u^n, w^n \cdot \#^{n \cdot |u|})$ are also good encodings, for all $n \geq 1$.

Suppose, by way of contradiction, that there is a one-way transducer T that computes the relation B_τ . For every $n, m \in \mathbb{N}$, we define the encoding

$$\alpha_{n,m} = \begin{pmatrix} w^n \cdot u^m \\ w^n \cdot \#^{m \cdot |u|} \end{pmatrix}$$

and we observe that $\alpha_{n,m} \in B_\tau$ if and only if $n \neq m$ (recall that $w \neq \varepsilon$ is the solution of the PCP instance τ and $u = f(w) = g(w)$). Below, we consider bad encodings like the above ones, where the parameter n is supposed to be large enough. Formally, we define the set I of all pairs of indices $(n, m) \in \mathbb{N}^2$ such that (i) $n \neq m$ (this guarantees that $\alpha_{n,m} \in B_\tau$) and (ii) n is larger than the number $|Q|$ of states of T .

We consider some pair $(n, m) \in I$ and we choose a successful run $\rho_{n,m}$ of T that witnesses the membership of $\alpha_{n,m}$ in B_τ , namely, that reads the input $w^n \cdot u^m$ and produces the output $w^n \cdot \#^{m \cdot |u|}$. We can split the run $\rho_{n,m}$ into a prefix $\overleftarrow{\rho}_{n,m}$ and a suffix $\overrightarrow{\rho}_{n,m}$ in such a way that $\overleftarrow{\rho}_{n,m}$ consumes the prefix w^n and $\overrightarrow{\rho}_{n,m}$ consumes the remaining suffix u^m . Since n is larger than the number of state of T , we can find a factor $\hat{\rho}_{n,m}$ of $\overleftarrow{\rho}_{n,m}$ that starts and ends with the same state and consumes a non-empty exact repetition of w , say w^{n_1} , for some $1 \leq n_1 \leq |Q|$. We claim that the output produced by the factor $\hat{\rho}_{n,m}$ must coincide with the consumed part w^{n_1} of the input. Indeed, if this were not the case, then deleting the factor $\hat{\rho}_{n,m}$ from $\rho_{n,m}$ would result in a new successful run that reads $w^{n-n_1} \cdot u^m$ and produces $w^{n-n_1} \cdot \#^{m \cdot |u|}$ as output, for some $n_2 \neq n_1$. This however would contradict the fact that, by definition of encoding, the possible outputs produced by T on input $w^{n-n_1} \cdot u^m$ must agree on the prefix w^{n-n_1} . We also remark that, even if we do not annotate this explicitly, the number n_1 depends on the choice of the pair $(n, m) \in I$. This number, however, range over the fixed finite set $J = [1, |Q|]$.

We can now pump the factor $\hat{\rho}_{n,m}$ of the run $\rho_{n,m}$ any arbitrary number of times. In this way, we obtain new successful runs of T that consume inputs of the form $w^{n+k \cdot n_1} \cdot u^m$ and produce outputs of the form $w^{n+k \cdot n_1} \cdot \#^m$, for all $k \in \mathbb{N}$. In particular, we know that B_τ contains all pairs of the form $\alpha_{n+k \cdot n_1, m}$. Summing up, we can claim the following:

Claim 9.4. There is a function $h : I \rightarrow J$ such that, for all pairs $(n, m) \in I$,

$$\{(n + k \cdot h(n, m), m) \mid k \in \mathbb{N}\} \subseteq I.$$

We can now head towards a contradiction. Let \tilde{n} be the maximum common multiple of the numbers $h(n, m)$, for all $(n, m) \in I$. Let $m = n + \tilde{n}$ and observe that $n \neq m$, whence $(n, m) \in I$. Since \tilde{n} is a multiple of $h(n, m)$, we derive from the above claim that the pair $(n + \tilde{n}, m) = (m, m)$ also belongs to I . However, this contradicts the definition of I , since

we observed earlier that $\alpha_{n,m}$ is a bad encoding if and only if $n \neq m$. We conclude that B_τ is not one-way definable when τ has a solution. \square

10. CONCLUSIONS

It was shown in [19] that it is decidable whether a given two-way transducer can be implemented by some one-way transducer. However, the provided algorithm has non-elementary complexity.

The main contribution of our paper is a new algorithm that solves the above question with elementary complexity, precisely in 2EXPSpace. The algorithm is based on a characterization of those transductions, given as two-way transducers, that can be realized by one-way transducers. The flavor of our characterization is different from that of [19]. The approach of [19] is based on a variant of Rabin and Scott's construction [30] of one-way automata, and on local modifications of the two-way run. Our approach instead relies on the global notion of *inversion* and more involved combinatorial arguments. The size of the one-way transducer that we obtain is triply exponential in the general case, and doubly exponential in the sweeping case, and the latter is optimal. The approach described here was adapted to characterize functional two-way transducers that are equivalent to some sweeping transducer with either known or unknown number of passes (see [5], [3] for details).

Our procedure considers non-deterministic transducers, both for the initial two-way transducer, and for the equivalent one-way transducer, if it exists. Deterministic two-way transducers are as expressive as non-deterministic functional ones. This means that starting from deterministic two-way transducers would address the same problem in terms of transduction classes, but could in principle yield algorithms with better complexity. We also recall that Proposition 3.5 gives a tight lower bound for the size of a one-way transducer equivalent to a given deterministic sweeping transducer. This latter result basically shows that there is no advantage in considering one-way definability for deterministic variants of sweeping transducers, at least with respect to the size of the possible equivalent one-way transducers.

A variant of the one-way definability problem asks whether a given two-way transducer is equivalent to some *deterministic* one-way transducer. A decision procedure for this latter problem is obtained by combining our characterization with the classical algorithm that determines whether a one-way transducer can be determinized [7, 11, 36]. In terms of complexity, it is conceivable that a better algorithm may exist for deciding definability by deterministic one-way transducers, since in this case one can rely on structural properties that characterize deterministic transducers.

ACKNOWLEDGMENT

The authors wish to acknowledge fruitful discussions with Emmanuel Filiot, Ismaël Jecker and Sylvain Salvati. We also thank the referees for their very careful reading and the suggestions for improvement.

REFERENCES

- [1] A. Aho, J. Hopcroft, and J. Ullman. A general theory of translation. *Math. Syst. Theory*, 3(3):193–221, 1969.
- [2] R. Alur and P. Cerný. Expressiveness of streaming string transducers. In *30th IARCS Annual Conference on Foundation of Software Technology and Theoretical Computer Science (FSTTCS)*, volume 8 of *LIPICs*, pages 1–12. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2010.
- [3] F. Baschenis. *Minimizing resources for regular word transductions*. PhD thesis, University of Bordeaux, France, 2017.
- [4] F. Baschenis, O. Gauwin, A. Muscholl, and G. Puppis. One-way definability of sweeping transducers. In *35th IARCS Annual Conference on Foundation of Software Technology and Theoretical Computer Science (FSTTCS)*, volume 45 of *LIPICs*, pages 178–191. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2015.
- [5] F. Baschenis, O. Gauwin, A. Muscholl, and G. Puppis. Minimizing resources of sweeping and streaming string transducers. In *43rd International Colloquium on Automata, Languages, and Programming (ICALP)*, volume 55 of *LIPICs*, pages 114:1–114:14. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2016. Full version available at <https://hal.archives-ouvertes.fr/hal-01274992>.
- [6] F. Baschenis, O. Gauwin, A. Muscholl, and G. Puppis. Untwisting two-way transducers in elementary time. In *32nd Annual ACM/IEEE Symposium on Logic in Computer Science (LICS)*. ACM, 2017. Full version available at <https://arxiv.org/abs/1701.02502>.
- [7] M.-P. Béal and O. Carton. Determinization of transducers over finite and infinite words. *Theoretical Computer Science*, 289(1):225–251, 2002.
- [8] J. Berstel. *Transductions and Context-free Languages*. Springer-Verlag, 2013.
- [9] J. Birget. Two-way automaton computations. *RAIRO - Theoretical Informatics and Applications - Informatique Théorique et Applications*, 24(1):47–66, 1990.
- [10] O. Carton and L. Dartois. Aperiodic two-way transducers and FO-transductions. In *24th EACSL Annual Conference on Computer Science Logic (CSL)*, volume 41 of *LIPICs*, pages 160–174. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2015.
- [11] C. Choffrut. Une caractérisation des fonctions séquentielles et des fonctions sous-séquentielles en tant que relations rationnelles. *Theoretical Computer Science*, 5(3):325–337, 1977.
- [12] C. Choffrut and B. Guillon. An algebraic characterization of unary two-way transducers. In *Mathematical Foundations of Computer Science 2014 - 39th International Symposium (MFCS)*, volume 8634 of *LNCS*, pages 196–207. Springer, 2014.
- [13] T. Colcombet. Factorisation forests for infinite words. In *Fundamentals of Computation Theory (FCT)*, volume 4639 of *LNCS*, pages 226–237. Springer, 2007.
- [14] L. Daviaud, P. Reynier, and J. Talbot. A generalised twinning property for minimisation of cost register automata. In *31st Annual ACM/IEEE Symposium on Logic in Computer Science (LICS)*, pages 857–866. ACM, 2016.
- [15] R. De Souza. Uniformisation of two-way transducers. In *International Conference on Language and Automata Theory and Applications*, pages 547–558. Springer, 2013.
- [16] S. Eilenberg. *Automata, languages, and machines*. Academic press, 1974.
- [17] J. Engelfriet and H. J. Hoogeboom. MSO definable string transductions and two-way finite-state transducers. *ACM Trans. Comput. Logic*, 2(2):216–254, 2001.

- [18] E. Filiot, O. Gauwin, and N. Lhote. First-order definability of rational transductions: An algebraic approach. In *31st Annual ACM/IEEE Symposium on Logic in Computer Science (LICS)*, pages 387–396. ACM, 2016.
- [19] E. Filiot, O. Gauwin, P. Reynier, and F. Servais. From two-way to one-way finite state transducers. In *28th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS)*, pages 468–477. IEEE Computer Society, 2013.
- [20] E. Filiot, S. N. Krishna, and A. Trivedi. First-order definable string transformations. In *34th International Conference on Foundation of Software Technology and Theoretical Computer Science (FSTTCS)*, volume 29 of *LIPICs*, pages 147–159. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2014.
- [21] I. Glaister and J. Shallit. A lower bound technique for the size of nondeterministic finite automata. *Information Processing Letters*, 59:75–77, 1996.
- [22] B. Guillon. Sweeping weakens two-way transducers even with a unary output alphabet. In *Seventh Workshop on Non-Classical Models of Automata and Applications (NCMA)*, volume 318 of *books@ocg.at*, pages 91–108. Österreichische Computer Gesellschaft, 2015.
- [23] E. M. Gurari. The equivalence problem for deterministic two-way sequential transducers is decidable. In *21st Annual Symposium on Foundations of Computer Science (FOCS)*, pages 83–85, 1980.
- [24] J. E. Hopcroft and J. D. Ullman. *Introduction to Automata Theory, Languages, and Computation*. Addison-Wesley, 1979.
- [25] O. H. Ibarra. The unsolvability of the equivalence problem for epsilon-free NGSMS with unary input (output) alphabet and applications. *SIAM J. Comput.*, 7(4):524–532, 1978.
- [26] I. Jecker. Personal communication, 2018.
- [27] J. Kortelainen. On the system of word equations $x_0 u_1^i x_1 u_2^i x_2 \dots u_m^i x_m = y_0 v_1^i y_1 v_2^i y_2 \dots v_m^i y_m$ ($i = 0, 1, 2, \dots$) in a free monoid. *Journal of Automata, Languages and Combinatorics*, 3(1):43–57, 1998.
- [28] D. Kozen. Lower bounds for natural proof systems. In *Proceedings of the 18th Annual Symposium on Foundations of Computer Science, SFCS'77*, pages 254–266. IEEE Computer Society, 1977.
- [29] M. Lothaire. *Combinatorics on words*. Number 17 in Encyclopedia of Mathematics and its Applications. Addison-Wesley Publishing Co., Reading, Mass., 1983.
- [30] M. Rabin and D. Scott. Finite automata and their decision problems. *IBM J. Res. Dev.*, 3(2):114–125, 1959.
- [31] A. Saarela. Systems of word equations, polynomials and linear algebra: a new approach. *European Journal of Combinatorics*, 47(5):1–14, 2015.
- [32] M.-P. Schützenberger. A remark on finite transducers. *Information and Control*, 4(2-3):185–196, 1961.
- [33] J. Shepherdson. The reduction of two-way automata to one-way automata. *IBM J. Res. Dev.*, 3(2):198–200, 1959.
- [34] I. Simon. Factorization forests of finite height. *Theoretical Computer Science*, 72(1):65–94, 1990.
- [35] T. Smith. A pumping lemma for two-way finite transducers. In *39th International Symposium on Mathematical Foundations of Computer Science (MFCS)*, LNCS, pages 523–534. Springer, 2014.
- [36] A. Weber and R. Klemm. Economy of description for single-valued transducers. *Information and Computation*, 118(2):327–340, 1995.

APPENDIX A.

Here we give a fully detailed proof of Lemma 6.5, for which we recall the statement below:

Lemma 6.5. *Let C be a component of a loop $L = [x_1, x_2]$. The nodes of C are precisely the levels in the interval $[\min(C), \max(C)]$. Moreover, if C is left-to-right (resp. right-to-left), then $\max(C)$ is the smallest level $\geq \min(C)$ such that between $(x_1, \min(C))$ and $(x_2, \max(C))$ (resp. $(x_2, \min(C))$ and $(x_1, \max(C))$) there are equally many LL-factors and RR-factors intercepted by L .*

Proof. To ease the understanding the reader may refer to Figure 12, that shows some factors intercepted by L and the corresponding edges in the flow.

We begin the proof by partitioning the set of levels of the flow into suitable intervals as follows. We observe that every loop $L = [x_1, x_2]$ intercepts equally many LL-factors and RR-factors. This is so because the crossing sequences at x_1, x_2 have the same length h . We also observe that the sources of the factors intercepted by L are either of the form (x_1, y) , with y even, or (x_2, y) , with y odd. For any location $\ell \in \{x_1, x_2\} \times \mathbb{N}$ that is the source of an intercepted factor, we define d_ℓ to be the difference between the number of LL-factors and the number of RR-factors intercepted by L that end at a location *strictly before* ℓ . Intuitively, $d_\ell = 0$ when the prefix of the run up to location ℓ has visited equally many times the position x_1 and the position x_2 . For the sake of brevity, we let $d_y = d_{(x_1, y)}$ for an even level y , and $d_y = d_{(x_2, y)}$ for an odd level y . Note that $d_0 = 0$. We also let $d_{h+1} = 0$, by convention.

We now consider the numbers z 's, with $0 \leq z \leq h + 1$, such that $d_z = 0$, that is: $0 = z_0 < z_1 < \dots < z_k = h + 1$. Using a simple induction, we prove that for all $i \leq k$, the parity of z_i is the same as the parity of its index i . The base case $i = 0$ is trivial, since $z_0 = 0$. For the inductive case, suppose that z_i is even (the case of z_i odd is similar). We prove that z_{i+1} is odd by a case distinction based on the type of factor intercepted by L that starts at level z_i . If this factor is an LR-factor, then it ends at the same level z_i , and hence $d_{z_{i+1}} = d_{z_i} = 0$, which implies that $z_{i+1} = z_i + 1$ is odd. Otherwise, if the factor is an LL-factor, then for all levels z strictly between z_i and z_{i+1} , we have $d_z > 0$, and since $d_{z_{i+1}} = 0$, the last factor before z_{i+1} must decrease d_z , that is, must be an RR-factor. This implies that (x_2, z_{i+1}) is the source of an intercepted factor, and thus z_{i+1} is odd.

The levels $0 = z_0 < z_1 < \dots < z_k = h + 1$ induce a partition of the set of nodes of the flow into intervals of the form $Z_i = [z_i, z_{i+1} - 1]$. To prove the lemma, it suffices to show that the subgraph of the flow induced by each interval Z_i is connected. Indeed, because the union of the previous intervals covers all the nodes of the flow, and because each node has one incoming and one outgoing edge, this will imply that the intervals coincide with the components of the flow.

Now, let us fix an interval of the partition, which we denote by Z to avoid clumsy notation. Hereafter, we will focus on the edges of subgraph of the flow induced by Z (we call it *subgraph of Z* for short). We prove a few basic properties of these edges. For the sake of brevity, we call LL-edges the edges of the subgraph of Z that correspond to the LL-factors intercepted by L , and similarly for the RR-edges, LR-edges, and RL-edges.

We make a series of assumption to simplify our reasoning. First, we assume that the edges are ordered based on the occurrences of the corresponding factors in the run. For instance, we may say the first, second, etc. LR-edge (of the subgraph of Z) — from now on, we tacitly assume that the edges are inside the subgraph of Z . Second, we assume that the

first edge of the subgraph of Z starts at an even node, namely, it is an LL-edge or an LR-edge (if this were not the case, one could apply symmetric arguments to prove the lemma). From this it follows that the subgraph contains n LR-edges interleaved by $n - 1$ RL-edges, for some $n > 0$. Third, we assume that $\min(Z) = 0$, in order to avoid clumsy notations (otherwise, we need to add $\min(Z)$ to all the levels considered hereafter).

Now, we observe that, by definition of Z , there are equally many LL-edges and RR-edges: indeed, the difference between the number of LL-edges and the number of RR-edges at the beginning and at the end of Z is the same, namely, $d_z = 0$ for both $z = \min(Z)$ and $z = \max(Z)$. It is also easy to see that the LL-edges and the RR-edges are all of the form $y \rightarrow y + 1$, for some level y . We call these edges *incremental edges*.

For the other edges, we denote by \vec{y}_i (resp. \overleftarrow{y}_i) the source level of the i -th LR-edge (resp. the i -th RL-edge). Clearly, each \vec{y}_i is even, and each \overleftarrow{y}_i is odd, and $i \leq j$ implies $\vec{y}_i < \vec{y}_j$ and $\overleftarrow{y}_i < \overleftarrow{y}_j$. Consider the location (x_1, \vec{y}_i) , which is the source of the i -th LR-edge (e.g. the edge in blue in the figure). The latest location at position x_2 that precedes (x_1, \vec{y}_i) must be of the form $(x_2, \overleftarrow{y}_{i-1})$, provided that $i > 1$. This implies that, for all $1 < i \leq n$, the i -th LR-edge is of the form $\vec{y}_i \rightarrow \overleftarrow{y}_{i-1} + 1$. For $i = 1$, we recall that $\min(Z) = 0$ and observe that the first location at position x_2 that occurs after the location $(x_1, 0)$ is $(x_2, 0)$, and thus the first LR-edge has a similar form: $\vec{y}_1 \rightarrow \overleftarrow{y}_0 + 1$, where $\overleftarrow{y}_0 = -1$ by convention.

Using symmetric arguments, we see that the i -th RL-edge (e.g. the one in red in the figure) is of the form $\overleftarrow{y}_i \rightarrow \vec{y}_i + 1$. In particular, the last LR-edge starts at the level $\vec{y}_n = \max(Z)$.

Summing up, we have just seen that the edges of the subgraph of Z are of the following forms:

- $y \rightarrow y + 1$ (incremental edges),
- $\vec{y}_i \rightarrow \overleftarrow{y}_{i-1} + 1$ (i -th LR-edge, for $i = 1, \dots, n$),
- $\overleftarrow{y}_i \rightarrow \vec{y}_i + 1$ (i -th RL-edge, for $i = 1, \dots, n - 1$).

In addition, we have $\vec{y}_i + 1 = \overleftarrow{y}_i + 2d_{\vec{y}_i}$. Since $d_z > 0$ for all $\min(Z) < z < \max(Z)$, this implies that $\vec{y}_i > \overleftarrow{y}_i$.

The goal is to prove that the subgraph of Z is strongly connected, namely, it contains a cycle that visits all its nodes. As a matter of fact, because components are also strongly connected subgraphs, and because every node in the flow has in-/out-degree 1, this will imply that the considered subgraph coincides with a component C , thus implying that the nodes in C form an interval. Towards this goal, we will prove a series of claims that aim at identifying suitable sets of nodes that are covered by paths in the subgraph of Z . Formally, we say that a path *covers* a set Y if it visits all the nodes in Y , and possibly other nodes. As usual, when we talk of edges or paths, we tacitly understand that they occur inside the subgraph of Z . On the other hand, we do not need to assume $Y \subseteq Z$, since this would follow from the fact that Y is covered by a path inside Z . For example, the right hand-side of Figure 12 shows a path from \vec{y}_i to $\vec{y}_i + 1$ that covers the set $Y = \{\vec{y}_i, \vec{y}_i + 1\} \cup [\overleftarrow{y}_{i-1} + 1, \overleftarrow{y}_i]$.

The covered sets will be intervals of the form

$$Y_i = [\overleftarrow{y}_{i-1} + 1, \overleftarrow{y}_i].$$

Note that the sets Y_i are well-defined for all $i = 1, \dots, n - 1$, but not for $i = n$ since \overleftarrow{y}_n is not defined either (the subgraph of Z contains only $n - 1$ RL-edges).

Claim A.1. For all $i = 1, \dots, n - 1$, there is a path from \vec{y}_i to $\vec{y}_i + 1$ that covers Y_i (for short, we call it an *incremental path*).

Proof. We prove the claim by induction on i . The base case $i = 1$ is rather easy. Indeed, we recall the convention that $\overleftarrow{y}_0 + 1 = \min(Z) = 0$. In particular, the node $\overleftarrow{y}_0 + 1$ is the target of the first LR-edge of the subgraph of Z . Before this edge, according to the order induced by the run, we can only have LL-edges of the form $y \rightarrow y + 1$, with $y = 0, 2, \dots, \overrightarrow{y}_1 - 2$. Similarly, after the LR-edge we have RR-edges of the form $y \rightarrow y + 1$, with $y = 1, 3, \dots, \overrightarrow{y}_1 - 1$. Those incremental edges can be connected to form the path $\overleftarrow{y}_{i-1} + 1 \rightarrow^* \overleftarrow{y}_1$ that covers the interval $[\overleftarrow{y}_0 + 1, \overleftarrow{y}_1]$. By prepending to this path the LR-edge $\overrightarrow{y}_1 \rightarrow \overleftarrow{y}_0 + 1$, and by appending the RL-edge $\overleftarrow{y}_1 \rightarrow \overrightarrow{y}_1 + 1$, we get a path from \overrightarrow{y}_1 to $\overrightarrow{y}_1 + 1$ that covers the interval $[\overleftarrow{y}_0 + 1, \overleftarrow{y}_1]$. The latter interval is precisely the set Y_1 .

For the inductive step, we fix $1 < i < n$ and we construct the desired path from \overrightarrow{y}_i to $\overrightarrow{y}_i + 1$. The initial edge of this path is defined to be the LR-edge $\overrightarrow{y}_i \rightarrow \overleftarrow{y}_{i-1} + 1$. Similarly, the final edge of the path will be the RL-edge $\overleftarrow{y}_i \rightarrow \overrightarrow{y}_i + 1$, which exists since $i < n$. It remains to connect $\overleftarrow{y}_{i-1} + 1$ to \overleftarrow{y}_i . For this, we consider the edges that depart from nodes strictly between \overleftarrow{y}_{i-1} and \overleftarrow{y}_i .

Let y be an arbitrary node in $[\overleftarrow{y}_{i-1} + 1, \overleftarrow{y}_i - 1]$. Clearly, y cannot be of the form \overleftarrow{y}_j , for some j , because it is strictly between \overleftarrow{y}_{i-1} and \overleftarrow{y}_i . So y cannot be the source of an RL-edge. Moreover, recall that the LL-edges and the RR-edges are the of the form $y \rightarrow y + 1$. As these incremental edges do not pose particular problems for the construction of the path, we focus mainly on the LR-edges that depart from nodes inside $[\overleftarrow{y}_{i-1} + 1, \overleftarrow{y}_i - 1]$.

Let $\overrightarrow{y}_j \rightarrow \overleftarrow{y}_{j-1} + 1$ be such an LR-edge, for some j such that $\overrightarrow{y}_j \in [\overleftarrow{y}_{i-1} + 1, \overleftarrow{y}_i - 1]$. If we had $j \geq i$, then we would have $\overrightarrow{y}_j \geq \overrightarrow{y}_i > \overleftarrow{y}_i$, but this would contradict the assumption that $\overrightarrow{y}_j \in [\overleftarrow{y}_{i-1} + 1, \overleftarrow{y}_i - 1]$. So we know that $j < i$. This enables the use of the inductive hypothesis, which implies the existence of an incremental path from \overrightarrow{y}_j to $\overrightarrow{y}_j + 1$ that covers the interval Y_j .

Finally, by connecting the above paths using the incremental edges, and by adding the initial and final edges $\overrightarrow{y}_i \rightarrow \overleftarrow{y}_{i-1} + 1$ and $\overleftarrow{y}_i \rightarrow \overrightarrow{y}_i + 1$, we obtain a path from \overrightarrow{y}_i to $\overrightarrow{y}_i + 1$. It is easy to see that this path covers the interval Y_i . \square

Next, we define

$$Y = [\overleftarrow{y}_{n-1} + 1, \overrightarrow{y}_n] \cup \bigcup_{1 \leq i < n} Y_i.$$

We prove a claim similar to the previous one, but now aiming to cover Y with a cycle. Towards the end of the proof we will argue that the set Y coincides with the full interval Z , thus showing that there is a component C whose set of notes is precisely Z .

Claim A.2. There is a cycle that covers Y .

Proof. It is convenient to construct our cycle starting from the last LR-edge, that is, $\overrightarrow{y}_n \rightarrow \overleftarrow{y}_{n-1} + 1$, since this will cover the upper node $\overrightarrow{y}_n = \max(Z)$. From there we continue to add edges and incremental paths, following an approach similar to the proof of the previous claim, until we reach the node \overrightarrow{y}_n again. More precisely, we consider the edges that depart from nodes strictly between \overleftarrow{y}_{n-1} and \overrightarrow{y}_n . As there are only $n - 1$ RL-edges, we know that every node in the interval $[\overleftarrow{y}_{n-1} + 1, \overrightarrow{y}_n - 1]$ must be source of an LL-edge, an RR-edge, or an LR-edge. As usual, incremental edges do not pose particular problems for the construction of the cycle, so we focus on the LR-edges. Let $\overrightarrow{y}_i \rightarrow \overleftarrow{y}_{i-1} + 1$ be such an LR-edge, with $\overrightarrow{y}_i \in [\overleftarrow{y}_{n-1} + 1, \overrightarrow{y}_n - 1]$. Since $i < n$, we know from the previous claim that there is a path from \overrightarrow{y}_i to $\overrightarrow{y}_i + 1$ that covers Y_i . We can thus build a cycle π by connecting the above paths using the incremental edges and the LR-edge $\overrightarrow{y}_n \rightarrow \overleftarrow{y}_{n-1} + 1$.

By construction, the cycle π covers the interval $[\overleftarrow{y}_{n-1} + 1, \overrightarrow{y}_n]$, and for every $i < n$, if π visits \overrightarrow{y}_i , then π covers Y_i . So to complete the proof — namely, to show that π covers the entire set Y — it suffices to prove that π visits each node \overrightarrow{y}_i , with $i < n$.

Suppose, by way of contradiction, that \overrightarrow{y}_i is the node with the highest index $i < n$ that is not visited by π . Recall that $\overrightarrow{y}_i > \overleftarrow{y}_i$. This shows that

$$\overrightarrow{y}_i \in [\overleftarrow{y}_i + 1, \overrightarrow{y}_n] = \bigcup_{i \leq j < n-1} [\overleftarrow{y}_j + 1, \overleftarrow{y}_{j+1}] \cup [\overleftarrow{y}_{n-1} + 1, \overrightarrow{y}_n].$$

As we already proved that π covers the interval $[\overleftarrow{y}_{n-1} + 1, \overrightarrow{y}_n]$, we know that $\overrightarrow{y}_i \in [\overleftarrow{y}_j + 1, \overleftarrow{y}_{j+1}]$ for some j with $i \leq j < n-1$. Now recall that \overrightarrow{y}_i is the highest node that is not visited by π . This means that \overrightarrow{y}_{j+1} is visited by π . Moreover, since $j+1 < n$, we know that π uses the incremental path from \overrightarrow{y}_{j+1} to $\overrightarrow{y}_{j+1} + 1$, which covers $Y_{j+1} = [\overleftarrow{y}_j + 1, \overleftarrow{y}_{j+1}]$. But this contradicts the fact that \overrightarrow{y}_i is not visited by π , since $\overrightarrow{y}_i \in [\overleftarrow{y}_j + 1, \overleftarrow{y}_{j+1}]$. \square

We know that the set Y is covered by a cycle of the subgraph of Z , and that Z is an interval whose endpoints are consecutive levels $z < z'$, with $d_z = d_{z'} = 0$. For the homestretch, we prove that $Y = Z$. This will imply that the nodes of the cycle are precisely the nodes of the interval Z . Moreover, because the cycle must coincide with a component C of the flow (recall that all the nodes have in-/out-degree 1), this will show that the nodes of C are precisely those of Z .

To prove $Y = Z$ it suffices to recall its definition as the union of the interval $[\overleftarrow{y}_{n-1} + 1, \overrightarrow{y}_n]$ with the sets Y_i , for all $i = 1, \dots, n-1$. Clearly, we have that $Y \subseteq Z$. For the converse inclusion, we also recall that $\overleftarrow{y}_0 + 1 = 0 = \min(Z)$ and $\overrightarrow{y}_n = \max(Z)$. Consider an arbitrary level $z \in Z$. Clearly, we have either $z \leq \overleftarrow{y}_i$, for some $1 \leq i < n$, or $z > \overleftarrow{y}_n$. In the former case, by choosing the smallest index i such that $z \leq \overleftarrow{y}_i$, we get $z \in [\overleftarrow{y}_{i-1} + 1, \overleftarrow{y}_i]$, whence $z \in Y_i \subseteq Y$. In the latter case, we immediately have $z \in Y$, by construction. \square