A Circuit Complexity Approach to Transductions

Michaël Cadilhac[1], Andreas Krebs[1], Michael Ludwig[1], and Charles Paperman[2]

[1] Wilhelm Schickard Institut, Universität Tübingen
[2] University of Warsaw

Abstract. We investigate the deterministic rational transductions computable by constant-depth, polysize circuits. To this end, we first propose a framework of independent interest to express functions of variable output length using circuits, and argue for its pertinence. We then provide a general characterization of the set of transductions realizable by such circuits, relying on a notion of continuity. We deduce that it is decidable whether a transduction is definable in AC^0 and, assuming a well-established conjecture, the same for ACC^0.

Introduction

Automata provide for an elegant and intuitive way to express regular languages, and even their intrinsic sequentiality is no obstacle to implementation, as they admit efficient parallelization. The connections between the theory of regular languages and that of circuits were first unveiled in the 1960s—it is now folklore that any regular language admits a (highly uniform) family of logdepth, polysize, and constant fan-in circuits. The natural further steps taken up by McNaughton and Papert [14], and then by Barrington *et al.* [2] led to the characterization of the regular languages in the class AC^0 of constant-depth, polysize, unbounded fan-in circuits using a decidable invariant. The property therein deviates sharply from the prevailing line of work at the time, which relied on the study of the syntactic monoid of regular languages. It was indeed known from the lower bound of Furst, Saxe and Sipser [11] that the language PARITY $= \{w \in \{0,1\}^* \mid |w|_1 \equiv 0 \mod 2\}$ is not in AC^0, while the language EVEN of even-length words over $\{0,1\}$, which has the same syntactic monoid, does belong to AC^0. Hence the class of regular languages in AC^0 does not admit a characterization solely in terms of the syntactic monoid.

We propose to take this study to the functional case, that is, to characterize the functions realized by rational transducers that are expressible by an AC^0 circuit family. Similarly to the context at the time of [2], we face a situation where, to the best of our knowledge, most characterizations focused on algebraic properties that would blur the line between PARITY and EVEN. This is the case, e.g., for the characterization of functional aperiodic nondeterministic transducers by FO-translations [13].

We rely on a property we abusively call *continuity* for a class $\mathcal{V}$, as borrowed from the field of topology: a transduction is $\mathcal{V}$-continuous if it preserves $\mathcal{V}$ by

inverse image. It is well known that any transduction τ is continuous for the regular languages; together with an additional property on the output length of τ, this even characterizes deterministic transductions [8]. Namely, with $d(u, v) = |u| + |v| - |u \wedge v|$, where $u \wedge v$ is the largest common prefix of u and v, the latter property is that $d(\tau(u), \tau(v)) \leq k \times d(u, v)$, a strong form of *uniform* continuity. Continuity thus appears as a natural invariant when characterizing transductions—the *forward* behaviors of τ, that is, its images, are less relevant, as any NP problem is the image of Σ^* under an AC^0 function [4].

Our contributions are three-fold:

- We propose a model of circuits that allows for functions of unrestricted output length: as opposed to previous models, e.g., [19], we do not impose the existence of a mapping between the input and output lengths.
- Relying on this model, we characterize the deterministic rational transductions in $\mathrm{AC}^0(\mathcal{V})$. Our characterization relies for one part on algebraic objects similar to the ones used in [2], through the use of the modern framework of *lm-varieties* [18]. For the other part, we rely on the notion of *continuity*. This bears a striking resemblance to the characterization of Reutenaeur and Schützenberger [16] of the transductions with a group as transition monoid.
- Our characterization then leads to the decidability of the membership of a deterministic rational transduction in AC^0 or in ACC^0. The procedure is effective, in the sense that an appropriate circuit can be produced realizing the transduction.

The paper is organized as follows. In Section 1, we cover in a succinct manner the automata- and circuit-theoretic notions necessary to our presentation. In Section 2, we introduce the aforementioned circuit model for functions and argue for its legitimacy by considering alternative frameworks. In Section 3, we show that studying the transition morphism of an automaton is equivalent to studying the languages accepted at each of its states; this enables us to keep to a minimum the algebraic references throughout our presentation. In Section 4, we show our characterization of $\mathrm{AC}^0(\mathcal{V})$ and delay until Section 5 its implications about AC^0 and ACC^0, in particular on decidability. We discuss our results and their limitations in Section 6.

1   Preliminaries

Monoid, morphisms, quotient. A monoid is a set equipped with a binary associative operation, denoted multiplicatively, with a unit element. For an alphabet Σ, the set Σ^* is the free monoid generated by Σ, its unit element being the empty word ε. A morphism is a map $\varphi \colon M \to N$ satisfying $\varphi(ab) = \varphi(a)\varphi(b)$ and $\varphi(1) = 1$, with $a, b \in M$ and 1 denoting the unit element of M and N. A morphism $\varphi \colon \Sigma^* \to T^*$ is an *lm-morphism*, where lm stands for *length-multiplying*, if there is a k such that $\varphi(\Sigma) \subseteq T^k$. Given a language L and a word u, the *left quotient* of L by u is the set $u^{-1}L = \{v \mid uv \in L\}$. The *right quotient* Lu^{-1} is defined symmetrically. For $w \in \Sigma^*$ and $a \in \Sigma$, we let $|w|_a$ be the number of a's in w, i.e., the image of w under the morphism $a \mapsto 1$, $\Sigma \setminus \{a\} \mapsto 0$ into $(\mathbb{N}, +)$.

Circuits. We use standard notations, as presented for instance in [17] and [19]. By AC^0, we denote the class of languages recognized by constant-depth, polysize circuit families with Boolean gates of unbounded fan-in. Such families recognize languages in $L \subseteq \{0,1\}^*$; to extend this to any alphabet Σ, we always assume there is a canonical map from Σ to $\{0,1\}^{|\Sigma|}$, that let us encode and decode words of Σ^* in binary. For a class of languages $\mathcal{V}$, we write $\mathrm{AC}^0(\mathcal{V})$ for languages recognized by AC^0 circuit families with access to gates for the languages in $\mathcal{V}$. We let $\mathrm{ACC}^0 = \mathrm{AC}^0(\mathrm{MOD})$ where MOD is the class of regular languages on $\{0,1\}^*$ of the form $\{|w|_0 \equiv |w|_1 \mod k\}$ for some k. Further, we define $\mathrm{TC}^0 = \mathrm{AC}^0(\mathrm{MAJ})$ where MAJ is the nonregular language $\{|w|_1 > |w|_0 \mid w \in \{0,1\}^*\}$. We will occasionally rely on the conjectured and widely-believed separation of $\mathrm{ACC}^0 \neq \mathrm{TC}^0$. Extending circuits to functions, a function f is in FAC^0 if there is a family of constant-depth, polysize circuits with multiple ordered output bits, such that $f(u)$ is the output of the circuit for input size $|u|$. We naturally extend the notation $\mathrm{AC}^0(\mathcal{V})$ to $\mathrm{FAC}^0(\mathcal{V})$.

Automata. A *deterministic automaton* is a tuple $A = (Q, \Sigma, \delta, q_0, F)$, where Q is the finite set of states, Σ the alphabet, $\delta \colon Q \times \Sigma \to Q$ is a partial transition function, q_0 is the initial state, and F is the set of final states. We naturally extend δ to words by letting $\delta(q, \varepsilon) = q$, and $\delta(q, aw) = \delta(\delta(q, a), w)$ when $\delta(q, a)$ is defined. We always assume that any state q is accessible and coaccessible, i.e., there is a word uv such that $\delta(q_0, u) = q$ and $\delta(q, v) \in F$. We write $L(A, q)$ for $\{w \mid \delta(q_0, w) = q\}$, and $L(A) = \cup_{f \in F} L(A, f)$ for the language of A. For two states q, q', we say that q can be *separated* from q' in $\mathcal{V}$ if there is a language L in $\mathcal{V}$ such that $L(A, q) \subseteq L \subseteq \overline{L(A, q')}$. An automaton is *all-separable in $\mathcal{V}$* if each pair of states can be separated in $\mathcal{V}$. It is *all-definable in $\mathcal{V}$* if every language $L(A, q)$ is in $\mathcal{V}$. We often use the shorter terms $\mathcal{V}$-all-separable and $\mathcal{V}$-all-definable, of self-explanatory meanings. We write REG for the class of regular languages.

Continuity, lm-varieties. A mapping $f \colon \Sigma^* \to T^*$ is *continuous for $\mathcal{V}$*, in short *$\mathcal{V}$-continuous*, if $L \in \mathcal{V}$ implies $f^{-1}(L) \in \mathcal{V}$—this name stems from the notion of continuity in topology. The sets of regular languages recognized by circuit families form a backbone of our work. It is thus natural to assume that these sets be closed under operations that AC^0 circuits can compute; this is formalized as follows. A class of languages $\mathcal{V}$ is an *lm-variety* if it is a Boolean algebra of languages closed under left and right quotient such that any lm-morphism is $\mathcal{V}$-continuous. It can be shown that if $\mathrm{AC}^0(\mathcal{V}) \cap \mathrm{REG} = \mathcal{V}$, then $\mathcal{V}$ is an lm-variety. As is customary, we write $\mathcal{QA}$ for $\mathrm{AC}^0 \cap \mathrm{REG}$ and $\mathcal{M}_{\mathrm{sol}}$ for $\mathrm{ACC}^0 \cap \mathrm{REG}$—these names stem from the algebraic classes recognizing the languages: quasi-aperiodic stamps and solvable monoids respectively, see Section 3 and [17] for more details. In particular, $\mathrm{ACC}^0 \neq \mathrm{TC}^0$ iff $\mathrm{AC}^0(\mathcal{M}_{\mathrm{sol}}) \cap \mathrm{REG} = \mathcal{M}_{\mathrm{sol}}$ [2]. In the sequel, the symbol $\mathcal{V}$ always denotes some lm-variety of languages.

Transducers. A *deterministic transducer* is a tuple $A = (Q, \Sigma, T, \delta, \nu, q_0, F)$ which is an automaton equipped with an additional alphabet T and a mapping

$\nu \colon Q \times \Sigma \to T^*$ of same domain as δ. We extend ν to words in Σ^* by letting $\nu(q, \varepsilon) = \varepsilon$ and $\nu(q, aw) = \nu(q, a)\nu(\delta(q, a), w)$, when $\delta(q, a)$ is defined. The partial function $\tau \colon \Sigma^* \to T^*$ mapping $w \in L(A)$ to $\nu(q_0, w)$ is called a *transduction*. A transducer is said to be *output-minimal* if for every pair of states q, q', there is a word w such that either only one of $\delta(q, w)$ or $\delta(q', w)$ is final, or both are and $\nu(q, w) \neq \nu(q', w)$. For any transduction τ, we fix an arbitrary output-minimal transducer $\mathsf{MinT}(\tau)$ realizing it. We will see that the choice of $\mathsf{MinT}(\tau)$ does not bear any impact on the results.

We freely use Q, Σ, T, etc. when an automaton or a transducer is under study, with the understanding that they are the relevant components of its defining tuple. Our focus being solely on automata, transducers, and transductions that are deterministic, we will omit mentioning determinism from now on.

## 2 Circuit frameworks for variable-length functions

In the literature, most of the work on functions computed by circuits focus on variants of the class FAC^0 (see, e.g., [1]). In these, multiple (ordered) output gates are provided, and there is thus an implicit mapping from input length to output length. Towards circumventing this limitation, we propose a few different frameworks, and establish some formal shortcomings in order to legitimize our final choice. Our main consideration is that a function defined using a constant-depth, polysize circuit family of some kind should be AC^0-continuous—this corresponds to a simple composition of the circuits. In particular, any FAC^0 function is readily seen to be AC^0-continuous.

2.1 Noninversability

To keep close to the Boolean case, we may want to equip circuits with a pair of inputs $\langle u, v \rangle$, and deem that the represented function is valued v on u if the circuit accepts the pair. By making no syntactic distinction between input and output, any function has the same complexity as its inverse if it is functional. We show that this blurs definability in the following sense:

Proposition 1. *There is an AC^0-continuous transduction in FAC^0 whose inverse is functional and not AC^0-continuous.*

Transducer for τ

Proof. Consider the minimal, two-state automaton for $L = 0^*(a0^*b0^*)^*$ and turn it into a transducer by letting $\nu(\cdot, 0) = 0$ and $\nu(\cdot, a) = \nu(\cdot, b) = 1$, and call τ the resulting transduction. The FAC^0 circuit for τ first checks that the input is in L. This can be done as $L \in \mathrm{AC}^0$, a fact that can be seen relying on the logical characterization of AC^0: a word is in L iff its first non-0 letter is an a, its last a b, and the closest non-0 letters to an a (resp. a b) are b's (resp. a's). Next, the circuit simply maps 0 to 0 and a, b to 1. The transduction being in FAC^0, it is AC^0-continuous. Now let $\sigma = \tau^{-1}$, it is clearly functional. But $\sigma^{-1}(L)$ is PARITY, hence σ is not AC^0-continuous. $\qquad\square$

Thus, much in the fashion of FAC^0, this implies that there should be distinguished input and output gates. We deal with how their lengths are specified next.

2.2 Output length as a parameter

Aiming for a natural and succinct model, we may want that the family of circuits be parametrized solely by the input length. In such a framework, the presented circuit for a given input length is equipped with a way to "deactivate" output gates, in order to allow for different output lengths. Formalizing this idea further, a *deactivating circuit* C with n inputs and m outputs is an usual circuit with an extra input valued $\mathbf{z}$, a new constant symbol. This new symbol behaves as follows: $1 \vee \mathbf{z} = \mathbf{z} \vee 1 = 1$, and any other combination of $\mathbf{z}, 0, 1, \vee, \wedge, \neg$ is valued $\mathbf{z}$. The output of C on a given input is its usual output stripped of the $\mathbf{z}$ symbol. The frameworks used in [13,10,5] are logic counterparts of this model. Then:

Proposition 2. *There is a transduction expressible as a constant-depth, polysize family of deactivating circuits which is not AC^0-continuous.*

Proof. The erasing morphism $0 \mapsto \varepsilon$, $1 \mapsto 1$ is a transduction τ that can be expressed as a family of circuits as in the statement of the Proposition. But $\tau^{-1}(1^{2\mathbb{N}})$ is PARITY $\notin AC^0$. □

Circuit for τ

We thus naturally reach the following definition, that will serve as a basis for our study:

Definition 1 (Functional circuits). *A function $\tau \colon \Sigma^* \to T^*$ is expressed as a circuit family $(C_m^n)_{n,m \geq 0}$, where C_m^n is a circuit with n inputs and $m{+}1$ outputs, if:*

$$(\forall u, v \in \Sigma^*) \qquad \tau(u) = v \Leftrightarrow C_{|v|}^{|u|}(u) = (v, 1) \ .$$

The size *of the family is the mapping from $\mathbb{N}$ to $\mathbb{N} \cup \{\infty\}$, defined by $n \mapsto \sup_{m \geq 0} |C_m^n|$. Similarly, the* depth *of the family is the mapping that associates n to the supremum of the depths of each C_m^n. The class FAC_v^0, standing for functions in AC^0 with variable output length, is the class of functions expressible as a family of constant-depth, polysize circuits. The class $FAC_v^0(\mathcal{V})$ is defined in the same fashion as $AC^0(\mathcal{V})$, and we let $FACC_v^0 = FAC_v^0(MOD)$.*

Remark 1. – Any function τ in $FAC_v^0(\mathcal{V})$ is such that $n \mapsto \max_{u \in \Sigma^n} |\tau(u)|$ is polynomially bounded. We show this implies that τ is $AC^0(\mathcal{V})$-continuous. Let $(C_m^n)_{n,m \geq 0}$ be the circuit family for τ. Given a language L in $AC^0(\mathcal{V})$ expressed by the circuit family $(D_n)_{n>0}$, $\tau^{-1}(L) \cap \Sigma^n$ is recognized by the circuit that applies a polynomial number of circuits C_m^n to the input, and checks that the only m such that C_m^n outputs $(v, 1)$ is such that $v \in D_m$.
  – If for any n there is an m such that $\tau(\Sigma^n) \subseteq \Sigma^m$, then $\tau \in FAC_v^0(\mathcal{V})$ is equivalent to $\tau \in FAC^0(\mathcal{V})$.

– We will be interested in functions from Σ^* to $\mathbb{N}$, and will speak of their circuit definability. In this context, the function is either seen as taking value in $\{1\}^*$, and dealt with using a variable-output-length circuit, or taking value in $\{0,1\}^*$ using an FAC^0-like circuit, the output value then corresponding to the position of the last 1 in the output. We note that these two views are equivalent, and hence we do not rely on a specific one. We note that (general) transductions from Σ^* to $\{1\}^*$ have been extensively studied in [7]; therein, Choffrut and Schützenberger show that such a function is a transduction iff it has a strong form of *uniform* continuity, akin to the one presented in the introduction, but for longest subwords rather than prefixes.

3 Separability, definability, and lm-varieties of stamps

Recall that the transition monoid of an automaton A is the monoid under composition consisting of the functions $f_w \colon Q \to Q$ defined by $f_w(q) = \delta(q, w)$. Historically, regular languages were studied through properties of the transition monoids of their minimal automata (the so-called *syntactic monoids*). As previously mentioned, the minimal automata for $\mathrm{EVEN} \in AC^0$ and $\mathrm{PARITY} \notin AC^0$ have the same transition monoid, hence the class $AC^0 \cap REG$ admits no syntactic monoid characterization. Starting with [2], there was thus a shift towards a study of the *transition morphisms* of automata, i.e., the surjective morphisms $\varphi \colon w \mapsto f_w$—it is shown therein that a REG language is in AC^0 iff $\varphi(\Sigma^s) \cup \{\varphi(\varepsilon)\}$ is an aperiodic monoid for some $s > 0$ and for the minimal automaton.

A *stamp* is a surjective morphism from a free monoid to a finite monoid. A systematic study of the classes of languages described by stamps turned out to be a particularly fruitful research endeavor of the past decade [18,9,15,6]. Our use of this theory will however be kept minimal, and we will strive to only appeal to it in this section. The goal of the forthcoming Lemma 1 is indeed to express algebraic properties in a language-theoretic framework only.

Given a stamp $\varphi \colon \Sigma^* \to M$, we say that L is *recognized* by φ if there is a set $E \subseteq M$ such that $L = \varphi^{-1}(E)$—in this case, we also say that L is recognized by M, which corresponds to the usual definition of recognition (e.g., [17]). We say that a stamp $\varphi \colon \Sigma^* \to M$ *lm-divides* a stamp $\psi \colon T^* \to N$ if $\varphi = \eta \circ \psi \circ h$, where $h \colon \Sigma^* \to T^*$ is an lm-morphism and $\eta \colon N \to M$ is a partial surjective morphism. The *product* of two stamps φ and ψ with the same domain Σ^* is the stamp mapping $a \in \Sigma$ to $(\varphi(a), \psi(a))$. Finally, an lm-variety of stamps is a class of stamps containing the stamps $\Sigma^* \to \{1\}$ and closed under lm-division and product. An Eilenberg theorem holds for lm-varieties: there is a one-to-one correspondence between lm-varieties of stamps and the lm-varieties of languages they recognize [18]. We show:

Lemma 1. *Let A be an automaton, $\mathbf{V}$ an lm-variety of stamps, and $\mathcal{V}$ its corresponding lm-variety of languages. The following are equivalent:*
  (i) The transition morphism of A is in $\mathbf{V}$;
  (ii) A is $\mathcal{V}$-all-definable;
 (iii) A is $\mathcal{V}$-all-separable.

Proof. (i) → (ii). Let φ be the transition morphism of A. Then $L(A, q) = \varphi^{-1}(E)$ where $E = \{f_w \mid f_w(q_0) = q\}$, hence $L(A, q) \in \mathcal{V}$.

(ii) → (iii). This is immediate, as $L(A, q)$ separates q from any other state.

(iii) → (i). Write $L_{q,q'}$ for the language separating q from q'. As each of these are recognized by stamps in $\mathbf{V}$ and $\mathbf{V}$ is closed under product, the language $L_q = \cap_{q' \neq q} L_{q,q'}$ is also recognized by a stamp in $\mathbf{V}$. Similarly, taking the product of the stamps recognizing the different L_q's, we see that all of the L_q's are recognized by the same stamp $\psi \colon \Sigma^* \to N$ in $\mathbf{V}$; let thus E_q be such that $L_q = \psi^{-1}(E_q)$. Let $\varphi \colon \Sigma^* \to M$ be the transition morphism of A. We claim that φ lm-divides ψ, concluding the proof as $\mathbf{V}$ is closed under lm-division. Define $\eta \colon N \to M$ by $\eta(\psi(w)) = \varphi(w)$. If η is well-defined, then it is a surjective morphism, and we are done as $\varphi = \eta \circ \psi$. Suppose $\varphi(u) \neq \varphi(v)$, then there is a $p \in Q$ such that $\delta(p, u) = q$ and $\delta(p, v)$ is either undefined or a state $q' \neq q$. Let w be a word such that $\delta(q_0, w) = p$, then $\psi(wu) \in E_q$ and $\psi(wv) \notin E_q$, hence $\psi(u) \neq \psi(v)$, showing that η is well-defined. $\square$

Remark 2. For $\mathcal{V} = \mathcal{QA}$ and $\mathcal{V} = \mathcal{M}_{\mathrm{sol}}$, the properties of Lemma 1 are decidable.

As advertised, the rest of this paper will now be free from (lm-varieties of) stamps except for a brief incursion when discussing our results in Section 6. Lemma 1 enables a study that stands in the algebraic tradition with no appeal to its tools.

4 The transductions in $\mathrm{FAC}^0_\mathrm{v}(\mathcal{V})$

In sharp contrast with the work of Reutenaeur and Schützenberger [16], we are especially interested in the shape of the outputs of the transduction. It turns out that most of its complexity is given by the following output-length function:

Definition 2 ($\tau_\#$). *Let τ be a transduction. The function $\tau_\# \colon \Sigma^* \to \mathbb{N}$ is the output-length function of $\mathsf{MinT}(\tau)$ with all the states deemed final. In symbols, $\tau_\#(w) = |\nu(q_0, w)|$, with $\mathsf{MinT}(\tau)$ as the underlying transducer.*

Theorem 1. *Let τ be a transduction and $\mathcal{V}$ be such that $\mathrm{AC}^0(\mathcal{V}) \cap \mathrm{REG} = \mathcal{V}$. The following constitutes a chain of implications:*

(i) $\tau \in \mathrm{FAC}^0_\mathrm{v}(\mathcal{V})$;
(ii) τ is $\mathrm{AC}^0(\mathcal{V})$-continuous;
(iii) τ is $\mathcal{V}$-continuous;
(iv) $\mathsf{MinT}(\tau)$ is $\mathcal{V}$-all-definable.
Moreover, if $\tau_\# \in \mathrm{FAC}^0_\mathrm{v}(\mathcal{V})$ then (iv) implies (i). Somewhat conversely, (i) implies $\tau_\# \in \mathrm{FAC}^0_\mathrm{v}(\mathcal{V})$.

Proof. (i) → (ii). This was alluded to in Remark 1.

(ii) → (iii). This follows from the closure under inverse transductions of REG and the hypothesis that the regular languages of $\mathrm{AC}^0(\mathcal{V})$ are in $\mathcal{V}$.

(iii) → (iv). Let q, q' be two states of $A = \mathsf{MinT}(\tau)$. We show that we can separate q from q'. We distinguish the following cases, that span all the possibilities thanks to the output-minimality of A:

Case 1: There is a w such that only one of $\delta(q, w)$ or $\delta(q', w)$ is in F. We suppose $\delta(q, w) \in F$, without loss of generality as $\mathcal{V}$ is closed under complement. Let $L = (\tau^{-1}(T^*))w^{-1}$, then a word x is in L iff $\delta(q_0, xw) \in F$. This is the case for all the words in $L(A, q)$ and for none of the words in $L(A, q')$, hence $L \in \mathcal{V}$ separates these two languages.

Case 2: There is a w such that both $\delta(q, w)$ and $\delta(q', w)$ are in F and $\nu(q, w) = u \neq u' = \nu(q', w)$. Then we have two possibilities:

Case 2.1: If $|u| = |u'|$. For a word $x \in \Sigma^*$, if $\tau(xw)$ ends with u, it cannot be that $\delta(q_0, x) = q'$. Hence $L = (\tau^{-1}(T^*u))w^{-1} \in \mathcal{V}$ separates $L(A, q)$ from $L(A, q')$.

Case 2.2: If $|u| \neq |u'|$. Define $k \in \mathbb{N}$ to be such that $|u| \not\equiv |u'| \bmod k$, and let $s \in \{0, 1, \dots, k-1\}$. Let $x \in \Sigma^*$ be such that $s \equiv |\nu(q_0, x)| \bmod k$. Then if $\delta(q_0, x) = q$, we have $|\tau(xw)| \equiv s + |u| \bmod k$, while $\delta(q_0, x) = q'$ makes this equation false. Hence the union over every s of the languages $(\tau^{-1}(T^{k\mathbb{N}+s+|u|}))w^{-1}$ separates $L(A, q)$ from $L(A, q')$.

(iv) → (i), assuming $\tau_\# \in \mathrm{FAC}^0_\mathrm{v}(\mathcal{V})$. We construct an $\mathrm{FAC}^0_\mathrm{v}(\mathcal{V})$ circuit family for τ. Fix an input size n and an output size m. Given an input $x = x_1 x_2 \cdots x_n$, we first check, using $\tau_\#$, that the output length of τ on x is indeed m, and wire this answer properly to the $(m+1)$-th output bit. Next, the j-th output bit, $1 \leq j \leq m$, is computed as follows. We apply $\tau_\#$ to every prefix of x, until we find an i such that $\tau_\#(x_{<i}) < j \leq \tau_\#(x_{\leq i})$, where $x_{<i} = x_1 x_2 \cdots x_{i-1}$ and similarly for $x_{\leq i}$. Relying on the languages $L(A, q)$, we find the state q in $\mathsf{MinT}(\tau)$ reached by $x_{<i}$, and let $u = \nu(q, x_i)$. The j-th output bit then corresponds to the $(j - \tau_\#(x_{<i}))$-th letter of u.

(i) → ($\tau_\# \in \mathrm{FAC}^0_\mathrm{v}(\mathcal{V})$). Suppose *(i)*, this implies *(iv)*. We construct an $\mathrm{FAC}^0(\mathcal{V})$ circuit family for $\tau_\#$. Fix the input size n, and let $x = x_1 x_2 \cdots x_n$ be the input. We can check, using the languages $L(\mathsf{MinT}(\tau), q)$, in which state q the word x ends when read. Let w_q be a fixed word such that $\delta(q, w_q) \in F$, and let $r = |\nu(q, w_q)|$. It suffices now to plug the word xw_q in the circuit for τ; the value of $\tau_\#(x)$ is then the length of $\tau(xw_q)$ minus r. $\qquad\square$

Remark 3. It is clear from the proof of Theorem 1 that $\mathsf{MinT}(\tau)$, and hence $\tau_\#$, can be arbitrarily chosen as long as the output-minimality property is respected. We discuss the role of $\tau_\#$ at greater length in Section 6.

5   An application to AC^0 and ACC^0

Our primary focus is on the *decidability* of the membership of transductions in small-complexity classes. Theorem 1, while providing a characterization of these transductions, does not come with a decidable property in the general case—even when some conjectured separations are presupposed. In the specific case of AC^0 and ACC^0, we are able to meaningfully characterize the functions $\tau_\#$ that can be expressed with circuits. To this end, we define:

Definition 3 (Constant ratio). *A transducer has* constant ratio *if every two words of the same length looping on a state produce outputs of the same length from this state. In symbols, for any state q and any words u, v of the same length, $\delta(q, u) = \delta(q, v) = q$ implies $|\nu(q, u)| = |\nu(q, v)|$.*

Remark 4. The name of the latter property stems from the fact that in such a transducer, for any state q, there is a ratio θ such that if $\delta(q, u) = q$, then $|\nu(q, u)| = \theta|u|$. Indeed, suppose a transducer has constant ratio, and let u and v be words with $\delta(q, u) = \delta(q, v) = q$ for some q. Write $|\nu(q, u)| = \theta_1|u|$ and $|\nu(q, v)| = \theta_2|v|$. Then $x = u^{|v|}$ and $y = v^{|u|}$ are of the same length, and $\theta_1|u| \times |v| = |\nu(q, x)| = |\nu(q, y)| = \theta_2|v| \times |u|$, hence $\theta_1 = \theta_2$.

Lemma 2. (Assuming $\mathrm{ACC}^0 \neq \mathrm{TC}^0$.) *Let τ be a transduction. If $\tau_\#$ is in FACC_v^0, then $\mathsf{MinT}(\tau)$ has constant ratio.*

Proof. Suppose that $A = \mathsf{MinT}(\tau)$ does not have constant ratio. We give a circuit family in $\mathrm{AC}^0(\tau_\#)$ for the language $L = \{w \in \{0, 1\}^* \mid |w|_0 = |w|_1\}$, which is complete for TC^0. This shows that $\tau_\#$ cannot admit an FACC_v^0 circuit family.

As A does not have constant ratio, there are a state q in A and two words $u, v \in \Sigma^*$ of the same length, such that $\delta(q, u) = \delta(q, v) = q$ and $\ell_u = |\nu(q, u)|$ is different from $\ell_v = |\nu(q, v)|$. Further, let w_{in} (resp. w_{out}) be such that $\delta(q_0, w_{\mathrm{in}}) = q$ (resp. $\delta(q, w_{\mathrm{out}}) \in F$), and let $\ell_{\mathrm{in}} = |\nu(q_0, w_{\mathrm{in}})|$ (resp. $\ell_{\mathrm{out}} = |\nu(q, w_{\mathrm{out}})|$).

We describe the circuit for L for input size n. Let x denote the input. First, the circuit transforms each 0 into u, and each 1 into v—this can be done as $|u| = |v|$. Then w_{in} is prepended and w_{out} appended to it, and the resulting word x' is fed to $\tau_\#$. The output is $\ell_{\mathrm{in}} + |x|_0 \times \ell_u + |x|_1 \times \ell_v + \ell_{\mathrm{out}}$. This is equal to:

$$\tau_\#(x') = \ell_{\mathrm{in}} + \frac{1}{2}(|x|(\ell_u + \ell_v) + (|x|_0 - |x|_1)(\ell_u - \ell_v)) + \ell_{\mathrm{out}} \ .$$

Now $(|x|_0 - |x|_1)(\ell_u - \ell_v)$ cancels out if and only if x has as many 0's as 1's. Hence $x \in L$ iff the output of $\tau_\#$ is $\ell_{\mathrm{in}} + \frac{1}{2}|x|(\ell_u + \ell_v) + \ell_{\mathrm{out}}$, and this is verifiable with an AC^0 circuit. □

Having a constant ratio provides an easy way to compute the length function:

Lemma 3. *Let τ be a transduction. If $\mathsf{MinT}(\tau)$ has constant ratio and is $\mathcal{V}$-all-definable, then $\tau_\#$ is in $\mathrm{FAC}_v^0(\mathcal{V})$.*

Proof. We rely on the strongly connected components (SCC) of $A = \mathsf{MinT}(\tau)$, as every state in an SCC has the same ratio. Let us first make this remark clear. Let q, q' be two states in the same SCC of A. Suppose q has ratio θ and q' has ratio θ'. Then let u be a word from q to q', x a word from q' to q', and v a word from q' back to q. The output produced when reading uxv starting from q is of length $\theta|uv| + \theta'|x|$ and also $\theta|uvx|$, hence $\theta = \theta'$.

Now the circuit for $\tau_\#$ simply guesses the (constant number of) passing points between SCCs, and applies the correct ratio in-between. The rest of this proof is a more precise formulation of that idea. Consider the directed acyclic graph (DAG) of SCCs of A; one goes down the tree using a set of transitions Δ that cannot be taken twice in a run. Consider a branch in the DAG of SCCs from the initial state, and let $D \subseteq \Delta$ be the set of transitions in Δ in the branch. The circuit for $\tau_\#$ checks that the input word follows this branch: for each transition $(\delta(q, a) = q')$ in D, there should be a prefix of the input x that is in $L(A, q)$ and such that the next letter is an a. Once the positions of all the transitions in D have been found in the input word, the output length is entirely determined. □

Corollary 1. *Let τ be a transduction. The following are equivalent, where the the "resp." part assumes $\mathrm{ACC}^0 \neq \mathrm{TC}^0$:*

(i) $\tau \in \mathrm{FAC}_v^0$ (resp. $\in \mathrm{FACC}_v^0$);

(ii) τ is continuous for AC^0 (resp. for ACC^0) and $\mathsf{MinT}(\tau)$ has constant ratio;

(iii) τ is continuous for $\mathcal{QA}$ (resp. for $\mathcal{M}_{sol}$) and $\mathsf{MinT}(\tau)$ has constant ratio;

(iv) $\mathsf{MinT}(\tau)$ is all-definable for $\mathcal{QA}$ (resp. for $\mathcal{M}_{sol}$) and has constant ratio.

Remark 5. It should be noted that the choice of $\mathsf{MinT}(\tau)$ is again irrelevant. Either all the output-minimal transducers for τ are constant ratio, or none are.

Theorem 2. *It is decidable whether a transducer realizes an FAC_v^0 function. If it does, then a circuit family can be constructed. The same holds for FACC_v^0 assuming $\mathrm{ACC}^0 \neq \mathrm{TC}^0$.*

Proof. This is a direct consequence of Corollary 1, together with the minimization algorithm of [8], the fact that $\mathcal{V}$-all-definability is decidable, and the fact that it can be checked that a transducer has constant ratio: it is indeed enough to check the property on cycles that do not go twice in the same state except for the first. The constructions of the circuits in Theorem 1 and Lemma 3 are then effective. □

6 Discussion and limitations

1   Corollary 1 can be slightly strengthened for AC^0, as in this case:

**Proposition 3.** *If τ is $\mathcal{QA}$-continuous, then $\mathsf{MinT}(\tau)$ has constant ratio.*

Proof. This is a variant of Lemma 2, where we only rely on the inverse image of τ instead of a full circuit construction.

Suppose that $A = \mathsf{MinT}(\tau)$ does not have constant ratio. There are a state q in $A = \mathsf{MinT}(\tau)$ and two words $u, v \in \Sigma^*$ of the same length, such that reading u (resp. v) from q produces an output of length ℓ_u (resp. ℓ_v), and $\ell_u < \ell_v$. Now the words $y = u^{\ell_v}$ and $z = v^{2\ell_u}$ are such that y produces an output of size $\ell_y = \ell_u \times \ell_v$, and z produces an output of size $\ell_z = \ell_v \times 2\ell_u = 2\ell_y$.

Now if τ is a $\mathcal{QA}$-continuous transduction, so is the function τ' mapping $x \in \{0,1\}^*$ to a word on $\{a\}^*$ with $\ell_y \times |x|_1 + \ell_z \times |x|_0$ letters a—it is simply a matter of replacing 1 with y, 0 with z, and correctly reaching the state q. But $\tau'^{-1}(a^{2\ell_y \mathbb{N}})$ is PARITY: indeed, x has an odd number of 1 iff $\tau'(x)$ contains an odd number of blocks a^{ℓ_y}. Hence τ' is not $\mathcal{QA}$-continuous, and neither is τ. □

We note that this fails in the case of ACC^0, as the following example shows. Consider the morphism $h: a \mapsto a, b \mapsto aa$. As the regular languages of ACC^0, $\mathcal{M}_{sol}$, are closed under inverse morphism (this is a consequence of $\mathcal{M}_{sol}$ being a variety), h is $\mathcal{M}_{sol}$-continuous. However, $\mathsf{MinT}(h)$ does not have constant ratio. This was already noted in a different setting by Lange and McKenzie [12].

2   The major role that $\tau_\#$ plays in Theorem 1 raises several questions. First, is it the case that all the complexity of a transduction is characterized by its

length function? In symbols, is it true that $\tau_\# \in \mathrm{FAC}_v^0(\mathcal{V}) \Rightarrow \tau \in \mathrm{FAC}_v^0(\mathcal{V})$? The following example shows that it is not. Consider the transduction from $\{0,1\}^*$ to $\{a,b\}^*$ that outputs a if the word read so far is in PARITY, and b otherwise. Then $\tau_\#$ is total and maps every word to its length, it is thus in FAC_v^0. However, w is in PARITY iff the last letter of $\tau(w)$ is an a, that is, $\tau^{-1}(\{a,b\}^*a)$ is PARITY, hence τ is not $\mathcal{QA}$-continuous, thus cannot be in FAC_v^0.

Transducer for τ

Next, going down two levels in the statement of Theorem 1, we may wonder whether the $\mathcal{V}$-continuity of τ is equivalent to that of $\tau_\#$. One direction is true, but its converse fails, as the previous example shows:

Proposition 4. *If τ is $\mathcal{V}$-continuous, then so is $\tau_\#$.*

Proof. Suppose τ is $\mathcal{V}$-continuous. Let E be a set of integers, and write Σ^E for the words of lengths in E. Suppose $\{a\}^E$ is in $\mathcal{V}$; we show $\tau_\#^{-1}(E) \in \mathcal{V}$.

From Theorem 1, $A = \mathsf{MinT}(\tau)$ is $\mathcal{V}$-all-definable. Let q be a state of A, and w a word mapping q to a final state while outputing u. Then $(\tau^{-1}(\Sigma^E.u))w^{-1} \cap L(A, q)$ is in $\mathcal{V}$, as τ is $\mathcal{V}$-continuous and $\{a\}^E \in \mathcal{V}$. Now the union of all these sets for all states q is precisely $\tau_\#^{-1}(E)$, hence it is in $\mathcal{V}$. $\qquad \square$

  3   Our interest in circuits obscured an equally interesting problem: characterizing the $\mathcal{V}$-continuous transductions. A general question raised by our characterization is:

Question 1. Which lm-varieties $\mathcal{V}$ verify the following statement? A transduction τ is $\mathcal{V}$-continuous iff $\mathsf{MinT}(\tau)$ is $\mathcal{V}$-all-definable and $\tau_\#$ is $\mathcal{V}$-continuous.

A direct consequence of Proposition 3 is that $\mathcal{V} = \mathcal{QA}$ verifies Question 1. Another such class is given in [16]; therein, Reutenaer and Schützenberger show that the property holds for $\mathcal{V} = \mathcal{G}$, the (lm-)variety of group languages, that is, languages with a group as syntactic monoid. More precisely, they show that τ is $\mathcal{G}$-continuous iff the transition monoid of $\mathsf{MinT}(\tau)$ is a group; this latter property is equivalent to: the transition morphism of $\mathsf{MinT}(\tau)$ is a stamp $\Sigma^* \to G$, for G a group. The set of such stamps is an lm-variety of stamps (see [6]), thus by Lemma 1, their characterization is indeed of the form of Question 1.

  4   It is interesting to note that the property on $\tau_\#$ of Question 1 vanishes for groups: this can be seen as a consequence of Reutenaer and Schützenberger's characterization itself, as $\tau_\#$ has the same transition monoid as $\mathsf{MinT}(\tau)$. On the other hand it is shown in the same article that there are transductions with an aperiodic monoid that are not continuous for aperiodic languages. This raises the question:

Question 2. Which lm-varieties $\mathcal{V}$ verify the following statement? If a transduction τ is such that $\mathsf{MinT}(\tau)$ is $\mathcal{V}$-all-definable, then $\tau_\#$ is $\mathcal{V}$-continuous.

  5   Recall that a nondeterministic transduction is functional iff it is realized by an unambiguous transduction (see, e.g., [3]). As circuits can read the input multiple times and in any direction, it seems that they can handle deterministic

and unambiguous transductions in the same fashion. Hence a generalization of Theorem 1 to the unrestricted case of rational functions should hold.

References

1. Agrawal, M., Allender, E., Datta, S.: On TC^0, AC^0, and arithmetic circuits. J. Comput. Syst. Sci 60(2), 395–421 (2000)
2. Barrington, D.A.M., Compton, K., Straubing, H., Thérien, D.: Regular languages in NC^1. J. Computer and System Sciences 44(3), 478–499 (1992)
3. Berstel, J.: Transductions and Context-Free Languages, Leitfäden der angewandten Mathematik und Mechanik LAMM, vol. 38. Teubner (1979)
4. Beyersdorff, O., Datta, S., Krebs, A., Mahajan, M., Scharfenberger-Fabian, G., Sreenivasaiah, K., Thomas, M., Vollmer, H.: Verifying proofs in constant depth. TOCT 5(1), 2 (2013)
5. Bojańczyk, M.: Transducers with origin information. In: ICALP. pp. 26–37. LNCS, Springer (2014)
6. Chaubard, L., Pin, J.É., Straubing, H.: First-order formulas with modular predicates. In: LICS. pp. 211–220. IEEE (2006)
7. Choffrut, C., Schutzenberger, M.P.: Counting with rational functions. Theoretical Computer Science 58(1–3), 81–101 (Jun 1988)
8. Choffrut, C.: A generalization of Ginsburg and Rose's characterization of G-S-M mappings. In: ICALP. pp. 88–103. Springer-Verlag, London, UK, UK (1979)
9. Esik, Z., Ito, M.: Temporal logic with cyclic counting and the degree of aperiodicity of finite automata. Acta Cybern. 16(1), 1–28 (2003)
10. Filiot, E., Krishna, S.N., Trivedi, A.: First-order definable string transformations. In: Raman, V., Suresh, S.P. (eds.) FSTTCS. LIPIcs, vol. 29, pp. 147–159 (2014)
11. Furst, M., Saxe, J.B., Sipser, M.: Parity, circuits, and the polynomial-time hierarchy. Theory of Computing Systems 17, 13–27 (1984)
12. Lange, K.J., McKenzie, P.: On the complexity of free monoid morphisms. In: Algorithms and Computation. pp. 247–255. LNCS, Springer Berlin Heidelberg (1998)
13. Lautemann, C., McKenzie, P., Schwentick, T., Vollmer, H.: The descriptive complexity approach to LOGCFL. J. Computer and Systems Sciences 62(4), 629–652 (2001)
14. McNaughton, R., Papert, S.: Counter-Free Automata. The MIT Press, Cambridge, Mass. (1971)
15. Pin, J.É., Straubing, H.: Some results on C-varieties. RAIRO-Theoretical Informatics and Applications 39(01), 239–262 (2005)
16. Reutenaeur, C., Schützenberger, M.P.: Variétés et fonctions rationnelles. Theoretical Computer Science 145(1–2), 229–240 (Jul 1995)
17. Straubing, H.: Finite Automata, Formal Logic, and Circuit Complexity. Birkhäuser, Boston (1994)
18. Straubing, H.: On logical descriptions of regular languages. In: Rajsbaum, S. (ed.) LATIN. LNCS, vol. 2286, pp. 528–538. Springer (2002)
19. Vollmer, H.: Introduction to circuit complexity. Springer-Verlag, Berlin-Heidelberg-New York-Barcelona-Hong Kong-London-Milan-Paris-Singapur-Tokyo (1999)