Cardinality and counting quantifiers on omega-automatic structures

Lukasz Kaiser, Sasha Rubin, Vince Bárány

Symposium on Theoretical Aspects of Computer Science 2008 (Bordeaux), pp. 385-396
www.stacs-conf.org

CARDINALITY AND COUNTING QUANTIFIERS ON OMEGA-AUTOMATIC STRUCTURES

ŁUKASZ KAISER [1], SASHA RUBIN [2], AND VINCE BÁRÁNY [1]

[1] Mathematische Grundlagen der Informatik, RWTH Aachen
E-mail address: `{kaiser,vbarany}@informatik.rwth-aachen.de`

[2] Department of Computer Science, University of Auckland
E-mail address: `rubin@cs.auckland.ac.nz`

Abstract. We investigate structures that can be represented by omega-automata, so called omega-automatic structures, and prove that relations defined over such structures in first-order logic expanded by the first-order quantifiers 'there exist at most $\aleph_0$ many', 'there exist finitely many' and 'there exist k modulo m many' are omega-regular. The proof identifies certain algebraic properties of omega-semigroups.

As a consequence an omega-regular equivalence relation of countable index has an omega-regular set of representatives. This implies Blumensath's conjecture that a countable structure with an ω-automatic presentation can be represented using automata on finite words. This also complements a very recent result of Hjörth, Khoussainov, Montalban and Nies showing that there is an omega-automatic structure which has no injective presentation.

1. Introduction

Automatic structures were introduced in [5] and later again in [6, 2] along the lines of the Büchi-Rabin equivalence of automata and monadic second-order logic. The idea is to encode elements of a structure $\mathfrak{A}$ via words or labelled trees (the codes need not be unique) and to represent the relations of $\mathfrak{A}$ via synchronised automata. This way we reduce the first-order theory of $\mathfrak{A}$ to the monadic second-order theory of one or two successors. In particular, the encoding of relations defined in $\mathfrak{A}$ by first order formulas are also regular, and automata for them can be computed from the original automata. Thus we have the fundamental fact that the first-order theory of an automatic structure is decidable.

Depending on the type of elements encoding the structure, the following natural classes of structures appear: automatic (finite words), ω-automatic (infinite words), tree-automatic (finite trees), and ω-tree automatic (infinite trees). Besides the obvious inclusions, for instance that automatic structures are also ω-automatic, there are still some some outstanding problems. For instance, a presentation over finite words or over finite trees can be transformed into one where each element has a unique representative.

Key words and phrases: ω-automatic presentations, ω-semigroups, ω-automata.

SYMPOSIUM ON THEORETICAL ASPECTS OF COMPUTER SCIENCE

Kuske and Lohrey [9] point out an ω-regular equivalence relation (namely $\sim_{\mathrm{e}}$ stating that two infinite words are position-wise eventually equal) with no ω-regular set of representatives. Thus, unlike the finite-word case, injectivity can not generally be achieved by selecting a regular set of representatives from a given presentation. In fact, using topological methods it has recently been shown [4] that there are omega-automatic structures having no injective presentation. However, we are able to prove that every omega-regular equivalence relation having only countably many classes does allow to select an omega-regular set of unique representants. Therefore, every countable omega-automatic structure does have an injective presentation.

A related question raised by Blumensath [1] is whether every countable ω-automatic structure is also automatic. In Corollary 2.8 we confirm this by transforming the given presentation into an injective one, and then noting that an injective ω-automatic presentation of a countable structure can be "packed" into one over finite words.

All these results rest on our main contribution: a characterisation of when there exist countably many words x satisfying a given formula with parameters in a given ω-automatic structure $\mathfrak{A}$ (with no restriction on the cardinality of the domain of $\mathfrak{A}$ or the injectivity of the presentation). The characterisation is first-order expressible in an ω-automatic presentation of an extension of $\mathfrak{A}$ by $\sim_{\mathrm{e}}$. Hence we obtain an extension of the fundamental fact for ω-automatic structures to include cardinality and counting quantifiers such as 'there exists (un)countably many', 'there exists finitely many', and 'there exists k modulo m many'. This generalises results of Kuske and Lohrey [9] who achieve this for structures with *injective* ω-automatic presentations.

2. Preliminaries

By countable we mean finite or countably infinite. Let Σ be a finite alphabet. With Σ^* and Σ^ω we denote the set of finite, respectively ω-words over Σ. The length of a word $w \in \Sigma^*$ is denoted by $|w|$, the empty word by ε, and for each $0 \le i < |w|$ the ith symbol of w is written as $w[i]$. Similarly $w[n,m]$ is the factor $w[n]w[n+1]\cdots w[m]$ and $w[n,m)$ is defined by $w[n,m-1]$. Note that we start indexing with 0 and that for $u \in \Sigma^*$ we denote by u^n the concatenation of n number of us, in particular $u^\omega \in \Sigma^\omega$.

We consider relations on finite and ω-words recognised by multi-tape finite automata operating in a synchronised letter-to-letter fashion. Formally, *ω-regular relations* are those accepted by some finite non-deterministic automaton $\mathcal{A}$ with Büchi, parity or Muller acceptance conditions, collectively known as ω-automata, and having transitions labelled by m-tuples of symbols of Σ. Equivalently, $\mathcal{A}$ is a usual one-tape ω-automaton over the alphabet Σ^m accepting the *convolution* $\otimes\vec{w}$ of ω-words $w_1, \dots, w_m$ defined by $\otimes\vec{w}[i] = (w_1[i], \dots, w_m[i])$ for all i.

Words $u, v \in \Sigma^\omega$ have *equal ends*, written $u \sim_{\mathrm{e}} v$, if for almost all $n \in \mathbb{N}$, $u[n] = v[n]$. This is an important ω-regular equivalence relation. We overload notation so that for $S, T \subset \mathbb{N}$ we write $S \sim_{\mathrm{e}} T$ to mean for almost all $n \in \mathbb{N}$, $n \in S \iff n \in T$.

Example 2.1. The non-deterministic Büchi automaton depicted in Fig. 1 accepts the equal-ends relation on alphabet $\{0, 1\}$.

In the case of finite words one needs to introduce a padding end-of-word symbol $\Box \notin \Sigma$ to formally define convolution of words of different length. For simplicity, we shall identify each finite word $w \in \Sigma^*$ with its infinite padding $w^{\Box} = w\Box^\omega \in \Sigma_{\Box}^\omega$ where $\Sigma_{\Box} = \Sigma \cup \{\Box\}$.

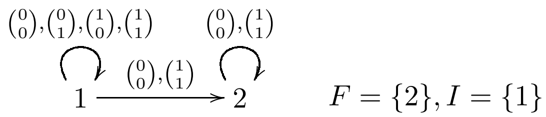

Figure 1: An automaton for the equal ends relation $\sim_{\mathrm{e}}$.

To avoid repeating the definition of automata for finite words, we say that a m-ary relation $R \subseteq (\Sigma^*)^m$ is *regular* (*synchronised rational*) whenever it is ω-regular over $\Sigma_\Box$.

2.1. Automatic structures

We now define what it means for a relational structure (we implicitly replace any structure with its relational counterpart) to have an (ω-)automatic presentation.

Definition 2.2 ((ω-)Automatic presentations)**.**
Consider a relational structure $\mathfrak{A} = (A, \{R_i\}_i)$ with universe $\mathsf{dom}(\mathfrak{A}) = A$ and relations R_i. A tuple of ω-automata $\mathfrak{d} = (\mathcal{A}, \mathcal{A}_\approx, \{\mathcal{A}_i\}_i)$ together with a surjective naming function $f : L(\mathcal{A}) \to A$ constitutes an *(ω-)automatic presentation* of $\mathcal{A}$ if the following criteria are met:

(i) the equivalence, denoted $\approx$, and defined by $\{(u, w) \in L(\mathcal{A})^2 \mid f(u) = f(w)\}$ is recognised by $\mathcal{A}_\approx$,
(ii) every $L(\mathcal{A}_i)$ has the same arity as R_i,
(iii) f is an isomorphism between $\mathfrak{A}_\mathfrak{d} = (L(\mathcal{A}), \{L(\mathcal{A}_i)\}_i)/_\approx$ and $\mathfrak{A}$.

The presentation is said to be *injective* whenever f is, in which case $\mathcal{A}_\approx$ can be omitted.

The relation $\approx$ needs to be a congruence of the structure $(L(\mathcal{A}), \{L(\mathcal{A}_i)\}_i)$ for item (iii) to make sense. In case $L(\mathcal{A})$ only consists of words of the form $w^\Box$ where $w \in \Sigma^*$, we say that the presentation is *automatic*. Call a structure ($\omega-$)*automatic* if it has an (ω-)automatic presentation.

The advantage of having an (ω-)automatic presentation of a structure lies in the fact that first-order (FO) formulas can be effectively evaluated using classical automata constructions. This is expressed by the following fundamental theorem.

Theorem 2.3. *(Cf.* [5], [6], [3].*)*

(i) There is an effective procedure that given an ($\omega-$)automatic presentation $\mathfrak{d}, f$ of a structure $\mathfrak{A}$, and given a FO*-formula $\varphi(\vec{a}, \vec{x})$ with parameters $\vec{a}$ from $\mathfrak{A}$ (defining a k-ary relation R over $\mathfrak{A}$), constructs a k-tape synchronous ($\omega-$)automaton recognising $f^{-1}(R)$.*
(ii) The FO*-theory of every ($\omega-$)automatic structure is decidable.*
(iii) The class of ($\omega-$)automatic structures is closed under FO*-interpretations*

Let FOC denote the extension of first-order logic with all quantifiers of the form

- $\exists^{(r \bmod m)} x \,.\, \varphi$ meaning that the number of x satisfying φ is finite and is congruent to r mod m;
- $\exists^\infty x \,.\, \varphi$ meaning that there are infinitely many x satisfying φ;
- $\exists^{\leq \aleph_0} x \,.\, \varphi$ and $\exists^{> \aleph_0} x \,.\, \varphi$ meaning that the cardinality of the set of all x satisfying φ is countable, or uncountable, respectively.

It has been observed that for *injective* (ω-)automatic presentations Theorem 2.3 can be extended from FO to FOC [8, 9]. Moreover, Kuske and Lohrey show that the cardinality of any set definable in FOC is either countable or equal to that of the continuum. Our main contribution is the following generalisation of their result.

Theorem 2.4. *The statements of Theorem 2.3 hold true for* FOC *over all (not necessarily injective)* ω*-automatic presentations.*

It is easily seen that finite-word automatic presentations can be assumed to be injective. This is achieved by restricting the domain of the presentation to a regular set of representatives of the equivalence involved. This can be done effectively, e.g. by selecting the length-lexicographically least word of every class.

This brings us to the question which ω-automatic structures allow an injective ω-automatic presentation. In [9] Kuske and Lohrey have pointed out that not every ω-regular equivalence has an ω-regular set of representatives. In particular, the following Lemma shows that the *equal-ends relation* $\sim_e$ of Example 2.1 is a counterexample.

Lemma 2.5 ([9, Lemma 2.4])**.** *Let* $\mathcal{A}$ *be a Büchi automaton with* n *states over* $\Sigma \times \Gamma$ *and let* $u \in \Sigma^\omega$ *be given. Consider the set* $V = \{v \in \Gamma^\omega \mid u \otimes v \in L(\mathcal{A})\}$*. Then* V *is uncountable if and only if* $|V/\sim_e| > n$*, otherwise it is finite or countable.*

The lemma implies that an ω-regular set is countable if and only if it meets only finitely many equal-ends-classes. In this case each of its members is ultimately periodic with one of finitely many periods.

Corollary 2.6. *An* ω*-regular set is countable iff it can be written as a finite union of sets of the form* $U_j \cdot (w_j)^\omega$ *with each* U_j *a regular set of finite words and each* w_j *a finite word.*

A related question raised by Blumensath [1] is whether every *countable* ω-automatic structure is also automatic. It is easy to see that every *injective* ω-automatic presentation of a countable structure can be "packed" into an automatic presentation.

Proposition 2.7. *(*[1, Theorem 5.32]*) Let* $\mathfrak{d}$ *be an injective* ω*-automatic presentation of a countable structure* $\mathcal{A}$*. Then, an (injective) automatic presentation* $\mathfrak{d}'$ *of* $\mathcal{A}$ *can be effectively constructed.*

In our proof of Theorem 2.4 we identify a property of finite semigroups that recognise transitive relations (Lemma 3.3 item (3)) that allows us to drop the assumption of injectivity in the previous statement. We are thus able to answer the question of Blumensath.

Corollary 2.8. *A countable structure is* ω*-automatic if and only if it is automatic. Transforming a presentation of one type into the other can be done effectively.*

2.2. ω-Semigroups

The fundamental correspondence between recognisability by finite automata and by finite semigroups has been extended to ω-regular sets. This is based on the notion of ω-*semigroups*. Rudimentary facts on ω-semigroups are well presented in [10]. We only mention what is most necessary.

An ω-semigroup $S = (S_f, S_\omega, \cdot, *, \pi)$ is a two-sorted algebra, where $(S_f, \cdot)$ is a semigroup, $* : S_f \times S_\omega \mapsto S_\omega$ is the *mixed product* satisfying for every $s, t \in S_f$ and every $\alpha \in S_\omega$ the equality

$$s \cdot (t * \alpha) = (s \cdot t) * \alpha$$

and where $\pi : S_f^\omega \mapsto S_\omega$ is the *infinite product* satisfying

$$s_0 \cdot \pi(s_1, s_2, \ldots) = \pi(s_0, s_1, s_2, \ldots)$$

as well as the associativity rule

$$\pi(s_0, s_1, s_2, \ldots) = \pi(s_0 s_1 \cdots s_{k_1}, s_{k_1+1} s_{k_1+2} \cdots s_{k_2}, \ldots)$$

for every sequence $(s_i)_{i \geq 0}$ of elements of S_f and every strictly increasing sequence $(k_i)_{i \geq 0}$ of indices. For $s \in S_f$ we denote $s^\omega = \pi(s, s, \ldots)$.

Morphisms of ω-semigroups are defined to preserve all three products as expected. There is a natural way to extend finite semigroups and their morphisms to ω-semigroups. As in semigroup theory, idempotents play a central role in this extension. An *idempotent* is a semigroup element $e \in S$ satisfying $ee = e$. For every element s in a finite semigroup the sub-semigroup generated by s contains a unique idempotent s^k. The least $k > 0$ such that s^k is idempotent for every $s \in S_f$ is called the *exponent* of the semigroup S_f and is denoted by π. Another useful notion is absorption of semigroup elements: say that s *absorbs* t *(on the right)* if $st = s$.

There is also a natural extension of the free semigroup Σ^+ to the ω-semigroup $(\Sigma^+, \Sigma^\omega)$ with $*$ and π determined by concatenation. An ω-semigroup $S = (S_f, S_\omega)$ *recognises* a language $L \subseteq \Sigma^\omega$ via a morphism $\phi : (\Sigma^+, \Sigma^\omega) \to (S_f, S_\omega)$ if $\phi^{-1}(\phi(L)) = L$. This notion of recognisability coincides, as for finite words, with that by non-deterministic Büchi automata. In [10] constructions from Büchi automata to ω-semigroups and back are also presented.

Theorem 2.9 ([10]).
A language $L \subseteq \Sigma^\omega$ is ω-regular iff it is recognised by a finite ω-semigroup.

We note that this correspondence allows one to engage in an algebraic study of varieties of ω-regular languages, and also has the advantage of hiding complications of cutting apart and stitching together runs of Büchi automata as we shall do. This is precisely the reason that we use this algebraic framework. Most remarkably, one does not need to understand the exact relationship between automata and ω-semigroups and the technical details of the constructions behind Theorem 2.9 to comprehend our proof. An alternative approach, though likely less advantageous, would be to use the composition method, which is closer in spirit to ω-semigroups than to automata. [1]

3. Cardinality and modulo counting quantifiers

This section is devoted to establishing the key to Theorem 2.4 announced earlier.

We characterise when there exist countably many words x satisfying a given formula *with parameters* $\varphi(x, \vec{z})$ in some ω-automatic structure $\mathfrak{A}$. The characterisation is first-order expressible in an ω-automatic extension of $\mathfrak{A}$ by the equal-ends relation $\sim_e$.

So, fix an ω-automatic presentation of some $\mathfrak{A}$ with congruence $\approx$, and a first-order formula $\varphi(x, \vec{z})$ in the language of $\mathfrak{A}$ with x and $\vec{z}$ free variables.

Proposition 3.1. *There is a constant C, computable from the presentation $\mathfrak{d}$, so that for all tuples $\vec{z}$ of infinite words the following are equivalent:*

(1) $\varphi(-, \vec{z})$ is satisfiable and $\approx$ restricted to the domain $\varphi(-, \vec{z})$ has countably many equivalence classes.

[1] Define T_f resp. T_ω as the sets of bounded (in terms of quantifier rank) theories of finite, respectively, of ω-words. The composition theorem ensures that $\cdot$, $*$, π can naturally be defined on bounded theories.

(2) *there exist C-many words $x_1, \cdots, x_C$ each satisfying $\varphi(-, \vec{z})$, so that every x satisfying $\varphi(-, \vec{z})$ is $\approx$-equivalent to some $y \sim_e x_i$. Formally, the structure $(\mathfrak{A}, \approx, \sim_e)$ models the sentence below.*

$$\forall \vec{z} \left(\exists^{\leq \aleph_0} w \,.\, \varphi(w, \vec{z}) \longleftrightarrow \exists x_1 \ldots x_C \left(\bigwedge_i \varphi(x_i, \vec{z}) \wedge \forall x \varphi(x, \vec{z}) \rightarrow \exists y (x \approx y \wedge \bigvee_i y \sim_e x_i) \right) \right)$$

Proof. Suppose $\mathfrak{d}$, $\mathfrak{A}$, and φ are given. Define C to be c^2, where c is the size of the largest ω-semigroup corresponding to any of the given automata (from the presentation or corresponding to φ). Now fix parameters $\vec{z}$. From now on, $\approx$ denotes the equivalence relation $\approx$ restricted to domain $\varphi(-, \vec{z})$.

$2 \to 1$: Condition 2 and the fact that every $\sim_e$-class is countable imply that all words satisfying $\varphi(-, \vec{z})$ are contained in a countable number of $\approx$-classes.

$1 \to 2$: We prove the contra-positive in three steps.

If $\varphi(-, \vec{z})$ is satisfiable then the negation of condition 2 implies that there are $C + 1$ many words $x_0, \ldots, x_C$ each satisfying $\varphi(-, \vec{z})$, and so that for $i, j \leq C$, $i \neq j$, the $\approx$-class of x_j does not meet the $\sim_e$-class of x_i. In particular, the x_is are pairwise $\not\sim_e$.

The plan is to produce uncountably many pairwise non-$\approx$ words that satisfy $\varphi(-, \vec{z})$. In the first 'Ramsey step', similar to what is done in [9], we find two words from the given C many, say $x_1, x_2 \in \Sigma^*$, and a factorisation $H \subset \mathbb{N}$ so that both words behave the same way along the factored sub-words with respect to the $\approx$- and φ-semigroups. In the second 'Coarsening step' we identify a technical property of finite semigroups recognising transitive relations. This allows us to produce an altered factorisation G and new, well-behaving words y_1, y_2. In the final step, the new words are 'shuffled along G' to produce continuum many pairwise non-$\approx$ words, each satisfying $\varphi(-, \vec{z})$.

3.1. Ramsey step

This step effectively allows us to discard the parameters $\vec{z}$. Before we use Ramsey's theorem, we introduce a convenient notation to talk about factorisations of words.

Definition 3.2. Let $A = a_1 < a_2 < \cdots$ be any subset of $\mathbb{N}$ and $h : \Sigma^* \to S$ be a morphism into a finite semigroup S. For an ω-word $\alpha \in \Sigma^\omega$, and element $e \in S$, say that A *is an h, e-homogeneous factorisation of* α if for all $n \in \mathbb{N}^+$, $h\big(\alpha[a_n, a_{n+1})\big) = e$.

Observe that

(1) if A is an h, s-homogeneous factorisation of α and $k \in \mathbb{N}^+$ then the set $\{a_{ki}\}_{i \in \mathbb{N}^+}$ is an h, s^k-homogeneous factorisation of α.

(2) if A is an h, e-homogeneous factorisation of α and e is idempotent, then every infinite $B \subset A$ is also an h, e-homogeneous factorisation of α.

In the following we write w^φ and $w^\approx$ to denote the image of w under the semigroup morphism into the finite semigroup associated to φ and $\approx$, respectively, as determined by the presentation. Accordingly, we will speak of e.g. φ, s_i-homogeneous factorisations.

Let us now colour every $\{n, m\} \in [\mathbb{N}]^2$, say $n < m$, by the tuple of ω-semigroup elements

$$\langle \ \big(\otimes (x_i, \vec{z})[n, m)^\varphi \big)_{0 \leq i \leq C} \ , \ \big(\otimes (x_i, x_j)[n, m)^\approx \big)_{0 \leq i \leq j \leq C} \ \rangle.$$

By Ramsey's theorem there exists infinite $H \subset \mathbb{N}$ and a tuple of ω-semigroup elements

$$\big\langle (s_i)_{1 \leq i \leq C}, (t_{(i,j)})_{1 \leq i \leq j \leq C} \big\rangle$$

so that for all $0 \leq i \leq j \leq C$,

- H is a φ, s_i-homogeneous factorisation of the word $\otimes(x_i, \vec{z})$,
- H is $\approx, t_{(i,j)}$-homogeneous factorisation of the word $\otimes(x_i, x_j)$.

Note that by virtue of additivity of our colouring and Ramsey's theorem each of the s_i and $t_{(i,j)}$ above are idempotents. Note that since there are at most c-many s_is and c-many $t_{(i,i)}$s there are at most c^2 many pairs $(s_i, t_{(i,i)})$ and so there must be two indices, we may suppose 1 and 2, with $s_1 = s_2$ and $t_{(1,1)} = t_{(2,2)}$.

3.2. Coarsening step

For technical reasons we now refine H and alter x_1, x_2 so that the semigroup elements have certain additional properties.

To start with, using the fact that $x_1 \not\sim_{\mathrm{e}} x_2$ and our observation on coarsenings, we assume without loss of generality that H is coarse enough so that $x_1[h_n, h_{n+1}) \neq x_2[h_n, h_{n+1})$ for all $n \in \mathbb{N}$.

Lemma 3.3. *There exists a subset $G \subset H$, listed as $g_1 < g_2 < \cdots$, and ω-words y_1, y_2 with the following properties:*

(1) The words y_1 and y_2 are neither $\approx$-equivalent nor $\sim_e$-equivalent, and each satisfies $\varphi(-, \vec{z})$.
(2) There exists an idempotent φ-semigroup element s such that G is a φ, s-homogeneous factorisation for each of $\otimes(y_1, \vec{z})$ and $\otimes(y_2, \vec{z})$.
(3) There exist idempotent $\approx$-semigroup elements $t, t^{\uparrow}, t^{\downarrow}$ so that for $y_j \in \{y_1, y_2\}$
   - *both $t^{\uparrow}$ and $t^{\downarrow}$ absorb t*
   - *$\otimes(y_j, y_j)[0, g_1)^{\approx}$ absorbs t*
   - *G is an $\approx, t$-homogeneous factorisation of $\otimes(y_j, y_j)$*
   - *G is an $\approx, t^{\uparrow}$-homogeneous factorisation of $\otimes(y_1, y_2)$*
   - *G is an $\approx, t^{\downarrow}$-homogeneous factorisation of $\otimes(y_2, y_1)$.*

Proof. Define ω-words $y_1 := x_2[0, h_2)x_1[h_2, \infty)$, and y_2 by

$$\begin{aligned} y_2[0, h_2) &:= x_2[0, h_2) \text{ and} \\ y_2[h_{2n}, h_{2n+2}) &:= x_2[h_{2n}, h_{2n+1})x_1[h_{2n+1}, h_{2n+2}) \text{ for } n > 0. \end{aligned}$$

Item 1. Clearly, $y_1 \not\sim_{\mathrm{e}} y_2$ and each $y_j \in \{y_1, y_2\}$ satisfies $\varphi(y_j, \vec{z})$ since by homogeneity and $s_1 = s_2$

$$\begin{aligned} \otimes(y_1, \vec{z})^{\varphi} &= \otimes(x_2, \vec{z})[0, h_2)^{\varphi} s_1^{\omega} \\ &= \otimes(x_2, \vec{z})[0, h_2)^{\varphi} s_2^{\omega} \\ &= \otimes(x_2, \vec{z})^{\varphi} \end{aligned}$$

and similarly

$$\begin{aligned} \otimes(y_2, \vec{z})^{\varphi} &= \otimes(x_2, \vec{z})[0, h_2)^{\varphi} (s_2 s_1)^{\omega} \\ &= \otimes(x_2, \vec{z})[0, h_2)^{\varphi} s_2^{\omega} \\ &= \otimes(x_2, \vec{z})^{\varphi} \end{aligned}$$

Next we check that $y_1 \not\approx y_2$.

$$\begin{aligned}
\otimes(y_1,y_2)^{\approx} &= \pi_{\approx}\big(\otimes(x_2,x_2)[0,h_2)^{\approx},\ \big(\otimes(x_1,x_2)[h_{2n},h_{2n+1})^{\approx},\ \otimes(x_1,x_1)[h_{2n+1},h_{2n+2})^{\approx}\big)_{n\in\mathbb{N}^+}\big)\\
&= \otimes(x_2,x_2)[0,h_1)^{\approx}\, t_{(2,2)}\ (t_{(1,2)}t_{(1,1)})^{\omega}\\
&= \otimes(x_2,x_2)[0,h_1)^{\approx}\, t_{(2,2)}t_{(2,2)}\ (t_{(1,2)}t_{(1,1)})^{\omega}\\
&= \otimes(x_2,x_2)[0,h_1)^{\approx}\, t_{(2,2)}t_{(2,2)}\ (t_{(1,2)}t_{(2,2)})^{\omega}\\
&= \otimes(x_2,x_2)[0,h_1)^{\approx}\, t_{(2,2)}\ (t_{(2,2)}t_{(1,2)})^{\omega}\\
&= \pi_{\approx}\big(\otimes(x_2,x_2)[0,h_2)^{\approx},\ \big(\otimes(x_2,x_2)[h_{2n},h_{2n+1})^{\approx},\ \otimes(x_1,x_2)[h_{2n+1},h_{2n+2})^{\approx}\big)_{n\in\mathbb{N}^+}\big)\\
&= \otimes(y_2,x_2)^{\approx}
\end{aligned}$$

Thus, if $y_1 \approx y_2$ then also $y_2 \approx x_2$ and so **by transitivity** $y_1 \approx x_2$. But since $y_1 \sim_{\mathrm{e}} x_1$, the $\approx$-class of x_2 meets the $\sim_{\mathrm{e}}$-class of x_1, contradicting the initial choice of the x_is.

Items 2 and 3. Define intermediate semigroup elements $q := s_1$, $r := t_{(1,1)}$, $r^{\uparrow} := t_{(1,2)}t_{(1,1)}$ and $r^{\downarrow} := t_{(2,1)}t_{(1,1)}$. Then

(1) both $r^{\uparrow}$ and $r^{\downarrow}$ absorb r, since $t_{(1,1)}$ is idempotent;
(2) $\otimes(y_j,y_j)[0,h_2)^{\approx} = \otimes(y_j,y_j)[0,h_1)^{\approx}t_{(2,2)}$ and thus absorbs r (for $y_j \in \{y_1,y_2\}$).

In this notation, for all $i \in \mathbb{N}^+$ and $y_j \in \{y_1,y_2\}$,

- $\otimes(y_j,\vec{z})[h_{2i},h_{2i+2})^{\varphi}$ is $qq = q$,
- $\otimes(y_j,y_j)[h_{2i},h_{2i+2})^{\approx}$ is $rr = r$,
- $\otimes(y_1,y_2)[h_{2i},h_{2i+2})^{\approx}$ is $t_{(1,2)}t_{(1,1)} = r^{\uparrow}$,
- $\otimes(y_2,y_1)[h_{2i},h_{2i+2})^{\approx}$ is $t_{(2,1)}t_{(1,1)} = r^{\downarrow}$.

Finally, define the set $G := \{h_{2ki}\}_{i>1}$, i.e. $g_i = h_{2k(i+1)}$, and the semigroup elements $t := r^k$, $t^{\uparrow} := (r^{\uparrow})^k$, $t^{\downarrow} := (r^{\downarrow})^k$ and $s := q^k$. The extra multiple of k (defined as the product of the exponents of the give semigroups for $\sim_{\mathrm{e}}$ and $\approx$) ensures all these semigroup elements (in particular $t^{\uparrow}$ and $t^{\downarrow}$) are idempotent. We now verify the absorption properties:

$$t^{\uparrow}t = r^{\uparrow k}r^k = r^{\uparrow k} = t^{\uparrow} \qquad \text{because } r^{\uparrow} \text{ absorbs } r$$

Similarly, $t^{\downarrow}t$ absorbs t. Further, since $g_1 = h_{4k}$, we have

$$\begin{aligned}
\otimes(y_j,y_j)[0,g_1)^{\approx} &= \otimes(y_j,y_j)[0,h_2)^{\approx} \otimes (y_j,y_j)[h_2,h_{4k})^{\approx}\\
&= \otimes(y_j,y_j)[0,h_2)^{\approx} r^{4k-2}\\
&= \otimes(y_j,y_j)[0,h_2)^{\approx} r^{3k-2}t
\end{aligned}$$

and thus absorbs t.

Finally we verify the homogeneity properties: G is an $\approx, t^{\downarrow}$-homogeneous factorisation of $\otimes(y_2,y_1)$ since for $i \in \mathbb{N}^+$

$$\otimes(y_2,y_1)[g_i,g_{i+1})^{\approx} = \otimes(y_2,y_1)[h_{2k(i+1)},h_{2k(i+2)})^{\approx} = (r^{\downarrow})^k = t^{\downarrow}.$$

The other cases are similar. ∎

3.3. Shuffling step

We continue the proof of Proposition 3.1 by 'shuffling' the words y_1 and y_2 along G resulting in continuum many pairwise distinct words that are pairwise not $\approx$-equivalent, each satisfying $\varphi(-, \vec{z})$. To this end, define for $S \subset \mathbb{N}^+$ the 'characteristic word' χ_S by

$$\begin{aligned} \chi_S[0, g_1) &:= y_2[0, g_1) \text{ , and} \\ \chi_S[g_n, g_{n+1}) &:= \begin{cases} y_2[g_n, g_{n+1}) & \text{if } n \in S \\ y_1[g_n, g_{n+1}) & \text{otherwise} \end{cases} \end{aligned}$$

First note that $\mathfrak{A} \models \varphi(\chi_S, \vec{z})$. Indeed, by Lemma 3.3 item 2

$$\begin{aligned} \otimes(\chi_S, \vec{z})^\varphi &= \otimes(y_2, \vec{z})[0, g_1)^\varphi s^\omega \\ &= \otimes(y_2, \vec{z})^\varphi \end{aligned}$$

and $\mathfrak{A} \models \varphi(y_2, \vec{z})$ by Lemma 3.3 item 1. Moreover, for $S \not\sim_\mathrm{e} T$ the construction gives that $\chi_S \not\sim_\mathrm{e} \chi_T$. This is due our initial choice of $x_1 \not\sim_\mathrm{e} x_2$ and the assumption that the factorisation $(h_n)_n$ is coarse enough so that $x_1[h_n, h_{n+1}) \neq x_2[h_n, h_{n+1})$ and therefore also $y_1[g_n, g_{n+1}) \neq y_2[g_n, g_{n+1})$ for all n.

The following two lemmas establish that if $S \not\sim_\mathrm{e} T$ then $\chi_S \not\approx \chi_T$.

Write $x_{\circ\bullet}$ for the word $\chi_{2\mathbb{N}^+}$, and $x_{\bullet\circ}$ for $\chi_{2\mathbb{N}^+-1}$ and let p denote $\otimes(y_2, y_2)[0, g_1)^\approx$.

Lemma 3.4. *For all $S \not\sim_e T$,*

$$\otimes(\chi_S, \chi_T)^\approx = \begin{cases} \otimes(x_{\circ\bullet}, x_{\bullet\circ})^\approx & \text{or} \\ \otimes(x_{\bullet\circ}, x_{\circ\bullet})^\approx \end{cases}$$

Proof. Define semigroup-elements p_n for $n \in \mathbb{N}$ by

$$p_n := \begin{cases} t^\downarrow & \text{if } n \in S \setminus T \\ t^\uparrow & \text{if } n \in T \setminus S \\ t & \text{otherwise} \end{cases}$$

Let m be the smallest number in $S \triangle T$. Suppose that $m \in S \setminus T$. Because both $t^\uparrow$ and $t^\downarrow$ are idempotent and since t is absorbed by both p, $t^\uparrow$ and $t^\downarrow$ we have

$$\begin{aligned} \otimes(\chi_S, \chi_T)^\approx &= \pi_\approx(p, (p_n)_{n\in\mathbb{N}}) = p(t^\downarrow t^\uparrow)^\omega \\ &= \otimes(x_{\bullet\circ}, x_{\circ\bullet})^\approx \end{aligned}$$

and the case that $m \in T \setminus S$ similarly results in $\otimes(x_{\circ\bullet}, x_{\bullet\circ})^\approx$. ∎

Lemma 3.5. *$x_{\circ\bullet} \not\approx x_{\bullet\circ}$.*

Proof. Define an intermediate word $x_{\circ\bullet\circ\circ} := \chi_{4\mathbb{N}^+-2}$. By computations similar to the above we find that

$$\begin{aligned} \otimes(x_{\bullet\circ}, x_{\circ\bullet\circ\circ})^\approx &= p(t^\downarrow t^\uparrow t^\downarrow t)^\omega = p(t^\downarrow t^\uparrow t^\downarrow)^\omega = p(t^\downarrow t^\uparrow)^\omega \\ &= \otimes(x_{\bullet\circ}, x_{\circ\bullet})^\approx \end{aligned}$$

and

$$\begin{aligned} \otimes(x_{\circ\bullet}, x_{\circ\bullet\circ\circ})^\approx &= p(tttt^\downarrow)^\omega = p(t^\downarrow)^\omega \\ &= \otimes(y_2, y_1)^\approx \end{aligned}$$

Therefore, if $x_{\bullet\circ} \approx x_{\circ\bullet}$ then also $x_{\bullet\circ} \approx x_{\circ\bullet\circ\circ}$ and so **by symmetry** and **by transitivity** $x_{\circ\bullet} \approx x_{\circ\bullet\circ\circ}$. But in this case also $y_2 \approx y_1$, contradicting Lemma 3.3 item 1. ∎

There are continuum many classes in $\mathcal{P}(\mathbb{N})/\sim_e$, thus there is a continuum of pairwise not $\approx$-equivalent words χ_S each satisfying $\varphi(-, \vec{z})$. This completes the proof of Proposition 3.1. ∎

4. Consequences

Theorem 2.4 *The statements of Theorem 2.3 hold true for* FOC *over all (not necessarily injective) ω-automatic presentations.*

Proof. We prove item (i) from which the rest of the theorem follows immediately. We inductively eliminate occurrences of cardinality and modulo-counting quantifiers in the following way.

The countability quantifier $\exists^{\leq\aleph_0}$ and uncountability quantifier $\exists^{>\aleph_0}$ can be eliminated (in an extension of the presentation by $\sim_{\mathrm{e}}$) by the formula given in Proposition 3.1.

For the remaining quantifiers we further expand the presentation with the ω-regular relations

- $\pi(a, b, c)$ saying that $a \sim_{\mathrm{e}} b \sim_{\mathrm{e}} c$ and the last position where a differs from c is no larger than the last position where b differs from c, and
- $\lambda(a, b, c)$ saying that $\pi(a, b, c)$ and $\pi(b, a, c)$ and, writing k for this common position, the word $a[0, k]$ is lexicographically smaller than the word $b[0, k]$.

Now $\exists^{<\infty}. \varphi(x, \vec{z})$ is equivalent to

$$\exists x_1 \cdots x_C \, \Psi(x_1, \cdots, x_C, \vec{z})$$

where Ψ expresses that $x_1, \cdots x_C$ satisfy $\varphi(-, \vec{z})$ and there exists a position, say $k \in \mathbb{N}$, so that every $\approx$-class contains a word satisfying $\varphi(-, \vec{z})$ that coincides with one of the x_i from position k onwards. This additional condition can be expressed by

$$\exists y_1 \cdots y_C \forall x \exists y \left(\varphi(x, \vec{z}) \rightarrow x \approx y \wedge \bigvee_i \pi(y, y_i, x_i) \right)$$

Consequently, $\exists^{(r \bmod m)} x \,.\, \varphi(x, \vec{z})$ can be eliminated since we can pick out unique representatives of the $\approx$-classes as those x so that, writing $i(w)$ for the smallest index i for which $w \sim_{\mathrm{e}} x_i$, for every $y \neq x$ in the same $\approx$-class as x, either

- $i(x) < i(y)$, or
- $i(x) = i(y)$ and $\lambda(x, y, x_{i(x)})$.

Now we can apply the construction of [9] or [8] for elimination of the $\exists^{(r \bmod m)}$ quantifier. ∎

As a corollary of Proposition 3.1 we obtain that for every omega-regular equivalence with countably many classes a set of unique representants is definable.

Corollary 4.1. *Let $\approx$ be an ω-automatic equivalence relation on Σ^ω. There is a constant C, depending on the presentation, so that the following are equivalent:*

(1) $\approx$ has countably many equivalence classes.

(2) there exist C many $\sim_e$-classes so that every $\approx$-class has non-empty intersection with at least one of these C.

In this case there is an ω-regular set of representatives of $\approx$. Moreover an automaton for this set can be effectively found given an automaton for $\approx$.

Proof. The first two items are simply a specialisation of Proposition 3.1. We get the representatives as follows.

Write A for the domain of $\approx$ and consider the formula $\psi(x_1, \cdots, x_C)$ with free variables $x_1, \cdots, x_C$:

$$\bigwedge_i x_i \in A \wedge (\forall x \in A)(\exists y)\, [x \approx y \wedge \bigvee_i y \sim_{\mathrm{e}} x_i]$$

The relation defined by ψ is ω-regular since it is a first order formula over ω-regular relations. By assumption it is non-empty. Thus it contains an ultimately periodic word of the form $\otimes(a_1, \cdots, a_C)$. Thus each of these a_is is ultimately periodic; say $a_i = v_i(u_i)^\omega$.

Then every x has an $\approx$-representative in $B := \bigcup_i \Sigma^*(u_i)^\omega$. It remains to prune B to select unique representatives for each $\approx$-class.

It is easy to construct an ω-regular well-founded linear order on B. For every $w \in B$, let $p(w) \in \Sigma^*$ be the length-lexicographically smallest word such that w has period $p(w)$. Also let $t(w) \in \Sigma^*$ be the length-lexicographically smallest word so that $w = t(w) \cdot p(w)^\omega$. Define an order $\prec$ on B by $w \prec w'$ if $p(w)$ is length-lexicographically smaller than $p(w')$, or otherwise if $p(w) = p(w')$ and $t(w)$ is length-lexicographically smaller than $t(w')$. The ordering $\prec$ is ω-regular since it is FO-definable in terms of ω-regular relations. Finally, the required set of representatives may be defined as the set of $\prec$-minimal elements of every $\approx$-class; and an automaton for this set can be constructed from an automaton for $\approx$. ∎

This immediately yields an *injective* ω-automatic presentation from a given ω-automatic presentation which by Proposition 2.7 can be transformed into an automatic presentation of the structure. Thus we conclude that every countable ω-automatic structure is already automatic.

Corollary 2.8 *A countable structure is ω-automatic if and only if it is automatic. Transforming a presentation of one type into the other can be done effectively.*

Note that some of our technical results, in particular Lemmas 3.3 and 3.4, only require transitivity of the relation $\approx$ and do not use symmetry. Applying them to an ω-automatic linear order $\prec$ we get an interesting uncountable set of words of the form $\chi_S, S \subseteq \mathbb{N}$. For any two such words with $S \not\sim_{\mathrm{e}} T$, whether $\chi_S \prec \chi_T$ or not depends only on the first position $m \in S \triangle T$. Thus, $\prec$ behaves like the lexicographic order on such words.

4.1. Failure of Löwenheim-Skolem theorem for ω-automatic structures

While so far the area of automatic structures has mainly focused on individual structures, it is interesting to look at their theories as well. We note a consequence of our work for 'automatic model theory'.

An automatic version of the Downward Löwenheim-Skolem Theorem would say that every uncountable ω-automatic structure has a countable elementary substructure that is also ω-automatic. Unfortunately this is false since there is a first-order theory with an ω-automatic model but no countable ω-automatic model. Indeed, consider the first-order theory of atomless Boolean Algebras. Kuske and Lohrey [9] have observed that it has an uncountable ω-automatic model, namely $(\mathcal{P}(\mathbb{N}), \cap, \cup, \neg)/\sim_{\mathrm{e}}$. However, Khoussainov et al. [7] show that the countable atomless Boolean algebra is not automatic and so, by Corollary 2.8, neither ω-automatic.

Here is the closest we can get to an automatic Downward Löwenheim-Skolem Theorem for ω-automatic structures.

Proposition 4.2. *Let $(D, \approx, \{R_i\}_{i\le\omega})$ be an omega-automatic presentation of $\mathfrak{A}$ and let $\mathfrak{A}_{up}$ be its restriction to the ultimately periodic words of D. Then $\mathfrak{A}_{up}$ is a countable elementary substructure of $\mathfrak{A}$.*

Proof. Relying on the Tarski-Vaught criterion for elementary substructures we only need to show that for all first-order formulas $\varphi(\vec{x}, y)$ and elements $\vec{b}$ of $\mathfrak{A}_{\mathrm{up}}$

$$\mathfrak{A} \models \exists y \varphi(\vec{b}, y) \quad \Rightarrow \quad \mathfrak{A}_{\mathrm{up}} \models \exists y \varphi(\vec{b}, y) \ .$$

By Theorem 2.3 $\varphi(\vec{x}, y)$ defines an omega-regular relation and, similarly, since the parameters $\vec{b}$ are all ultimately periodic the set defined by $\varphi(\vec{b}, y)$ is omega-regular. Therefore, if it is non-empty, then it also contains an ultimately periodic word, which is precisely what we needed. ∎

This proof can be viewed as a model construction akin to a classical compactness proof. Indeed, starting with ultimately constant words and throwing in witnesses for all existential formulas satisfied in $\mathfrak{A}$ in each round one constructs an increasing sequence of substructures comprising ultimately periodic words of increasing period lengths. The union of these is closed under witnesses by construction. The argument is valid for relational structures with constants assuming that every constant is represented by an ultimately periodic word.

Future work It remains to be seen whether statements analogous to Theorem 2.4 and Corollary 2.8 also hold for automatic presentations over infinite trees.

Acknowledgment We thank the referees for detailed technical remarks and corrections.

References

[1] A. Blumensath. Automatic structures. Diploma thesis, RWTH-Aachen, 1999.

[2] A. Blumensath and E. Grädel. Automatic Structures. In *Proceedings of 15th IEEE Symposium on Logic in Computer Science LICS 2000*, pages 51–62, 2000.

[3] A. Blumensath and E. Grädel. Finite presentations of infinite structures: Automata and interpretations. *Theory of Comp. Sys.*, 37:641 – 674, 2004.

[4] G. Hjörth, B. Khoussainov, A. Montalban, and A. Nies. Borel structures. Manuscript, 2007.

[5] B.R. Hodgson. Décidabilité par automate fini. *Ann. sc. math. Québec*, 7(1):39–57, 1983.

[6] B. Khoussainov and A. Nerode. Automatic presentations of structures. In *LCC '94*, volume 960 of *LNCS*, pages 367–392. Springer-Verlag, 1995.

[7] B. Khoussainov, A. Nies, S. Rubin, and F. Stephan. Automatic structures: Richness and limitations. In *LICS*, pages 44–53. IEEE Comp. Soc., 2004.

[8] B. Khoussainov, S. Rubin, and F. Stephan. Definability and regularity in automatic structures. In *STACS '04*, volume 2996 of *LNCS*, pages 440–451, 2004.

[9] D. Kuske and M. Lohrey. First-order and counting theories of ω-automatic structures. In *FoSSaCS*, pages 322–336, 2006.

[10] D. Perrin and J.-E. Pin. Semigroups and automata on infinite words. In J. Fountain, editor, *Semigroups, Formal Languages and Groups*, NATO Advanced Study Institute, pages 49–72. Kluwer, 1995.