On a Fragment of AMSO and Tiling Systems

Achim Blumensath*[1], Thomas Colcombet†[2], and Paweł Parys‡[3]

1 **Faculty of Informatics, Masaryk University, Czech Republic**
blumensath@fi.muni.cz
2 **CNRS, LIAFA, Université Paris Diderot, Paris 7, France**
thomas.colcombet@liafa.univ-paris-diderot.fr
3 **Institute of Informatics, University of Warsaw, Poland**
parys@mimuw.edu.pl

Abstract

We prove that satisfiability over infinite words is decidable for a fragment of asymptotic monadic second-order logic. In this fragment we only allow formulae of the form $\exists t \forall s \exists r\, \varphi(r, s, t)$, where φ does not use quantifiers over number variables, and variables r and s can be only used simultaneously, in subformulae of the form $s < f(x) \leq r$.

1998 ACM Subject Classification F.4.1 Mathematical Logic

Keywords and phrases monadic second-order logic, boundedness, tiling problems

1 Introduction

This paper continues a line of research trying to find logics where satisfiability is decidable over infinite words (and infinite trees). The most well-known logic of this kind is monadic second-order logic (MSO) considered in the seminal work of Büchi [8]. Extending MSO by the ability of comparing some quantities quickly leads to undecidability. The idea behind the logic MSO+U and a more recently introduced logic called *asymptotic monadic second-order logic* (AMSO) is to extend MSO by the ability to express boundedness properties of some sequences of numbers. In MSO+U this is realized by an additional quantifier U stating that there are arbitrarily large finite sets satisfying the given formula. AMSO does not have a built in ability to refer to the size of sets. Instead, it describes weighted structures (in particular weighted infinite words), which are structures in which the elements are labelled by natural numbers, called their *weights.* More precisely, AMSO extends MSO by quantifiers over variables of a new kind, ranging over natural numbers. These variables can be compared with weights in the word, but only under a certain positivity requirement: existentially quantified numbers can only serve as upper bounds, while universally quantified numbers can only serve as lower bounds. The two logics MSO+U and AMSO happen to be inter-reducible as far as the decidability of satisfiability is concerned [1], and, unfortunately, this means that both are undecidable over infinite words [5]. Nevertheless, some natural fragments of these logics remain decidable.

In [2] the satisfiability problem for MSO+U is solved over infinite trees for formulae where the quantifier U is at the outermost position. A significantly more powerful fragment of the

* Work partially supported by the German Science Foundation, grant No. BL 1127/2-2, and the Czech Science Foundation, grant No. P202/12/G061.
† The research leading to these results has received funding from the European Union's Seventh Framework Programme (FP7/2007-2013) under grant agreement n° 259454.
‡ Work supported by the fellowship of the Foundation for Polish Science, during the author's post-doc stay at Université Paris Diderot.

logic, although over infinite words, was shown to be decidable in [4] using automata with counters. These automata were further developed into the theory of regular cost functions [11]. Another possibility is to consider the weak fragment of the logic (WMSO+U), where set quantification is restricted to finite sets. Satisfiability for this logic was shown to be decidable over infinite words [3] and infinite trees [6]. Note that the mentioned decidability results can be used to solve, via reductions, several seemingly unrelated problems, among others: the star height problem [15], the finite power property problem [18], deciding properties of CTL* [9], the realizability problem for prompt LTL [16], deciding the winner in cost parity games [13], or deciding certain properties of energy games [7].

Concerning AMSO, which was more recently introduced [1], so far no fragments are known to be decidable (except trivial ones). Such fragments should, at least, circumvent the arguments of undecidability of AMSO, that involve complicated number quantifiers nested inside complicated quantification over infinite sets. There are two ways to avoid this: either to consider the weak fragment (WAMSO), where set quantification is restricted to finite sets, or to consider the number-prenex fragment ($\mathrm{AMSO}^{\mathrm{np}}$), where number quantifiers are required to be placed only at the head of the formula. It turns out that these two fragments are inter-reducible (Theorem 5 in [1]). It is conjectured that these two fragments have a decidable satisfiability problem over infinite words. Under a topological point of view, it is known that MSO+U and AMSO inhabit all finite levels of the projective hierarchy [14, 1], while WAMSO is much simpler since it only inhabits the finite levels of the Borel hierarchy.

et us emphasize the fact that WAMSO is not related at all to WMSO+U, even though AMSO and MSO+U are highly related. This is due to the fact that, since AMSO and MSO+U have significantly different syntax, the restriction to finite set quantifiers has dramatically different consequences. In particular languages definable in WAMSO inhabit all finite levels of the Borel hierarchy, while WMSO+U is confined in the third level.

In [1], the satisfiability problem for $\mathrm{AMSO}^{\mathrm{np}}$/WAMSO was reduced to a certain form of tiling problem. The main contribution of this paper is to solve a special case of this tiling problem. In consequence we can solve the satisfiability problem for a fragment of $\mathrm{AMSO}^{\mathrm{np}}$, which we denote $\mathrm{AMSO}^{\mathrm{np}}_{2s}$. In this fragment we only allow formulae of the form $\exists t \forall s \exists r\, \varphi(r, s, t)$, where φ does not use quantifiers over number variables, and the variables r and s can be only used simultaneously, in subformulae of the form $s < f(x) \leq r$. For the proof, we develop a new generalization of the Simon's theorem about factorization forests [17].

2 Preliminaries

Asymptotic monadic second-order logic (AMSO for short) extends MSO by the ability to describe asymptotic properties of quantities. It refers to *weighted structures* $\langle \mathfrak{A}, \overline{f} \rangle$ consisting of a relational structure $\mathfrak{A}$ and a tuple of functions $f_i \colon \mathit{dom}(\mathfrak{A}) \to \mathbb{N}$ (the *weight functions*). We only consider the case when $\mathfrak{A}$ is an infinite word (ω-word). Syntactically AMSO extends MSO by the following constructions:

- quantifiers over *number variables* that range over natural numbers, and
- atomic formulae $f(x) \leq r$, where f is a weight function, x a first-order variable, and r a number variable; such formulae are restricted to appear positively inside the existential quantifier binding r (and dually: negatively inside a universal quantifier).

We will usually reserve the letters $x, y, z, \ldots$ for first-order variables and the letters $r, s, t, \ldots$ for number variables.

The main theorem of this paper is about a fragment of AMSO, denoted $\mathrm{AMSO}^{\mathrm{np}}_{2s}$, where the formulae are of the form $\exists t\forall s\exists r\,\varphi(r,s,t)$ where φ does not use number quantifiers, and the variables r and s can be only used simultaneously, in subformulae of the form $s < f(x) \leq r$ (formally: $(f(x) \leq r) \wedge \neg(f(x) \leq s)$).

▶ **Example 2.1.** The following are formulae of $\mathrm{AMSO}^{\mathrm{np}}_{2s}$:
- $\exists t\forall x\,(f(x) \leq t)$ says that the weights are bounded;
- $\forall s\exists r\forall x\exists y\,(y > x \wedge s < f(y) \leq r)$ says that infinitely many weights occur infinitely often in the weighted infinite word;
- the disjunction of the above two (we can move the quantifiers to the front).

▶ Remark. It is easy to see that a formula of the form

$$\exists t_1 \ldots \exists t_k \forall s_1 \ldots \forall s_l \exists r_1 \ldots r_m\,\varphi(r_1,\ldots,r_m,s_1,\ldots,s_l,t_1,\ldots,t_k)$$

is equivalent to $\exists t\forall s\exists r\,\varphi(r,\ldots,r,s,\ldots,s,t,\ldots,t)$.[1] For this reason we allow in $\mathrm{AMSO}^{\mathrm{np}}_{2s}$ only formulae with single quantifiers $\exists t\forall s\exists r$, having in mind that decidability immediately extends to formulae with blocks of such quantifiers.

The following is the main result of this paper.

▶ **Theorem 2.2.** *Given a formula $\psi \in \mathrm{AMSO}^{\mathrm{np}}_{2s}$, it is decidable whether there exists a weighted infinite word in which ψ is satisfied.*

The Commutative Lossy Tiling Problem

Theorem 9 of [1] reduces satisfiability of $\mathrm{AMSO}^{\mathrm{np}}$ to a certain (multidimensional) *lossy tiling problem*. In this paper we solve a commutative variant of this problem, in dimension one.

A *picture* $p\colon \{1,\ldots,h\} \times \{1,\ldots,w\} \to \Sigma$ is a rectangle labelled by letters from a finite alphabet Σ, where h and w are *height* and *width* of the picture. For $i \in \{1,\ldots,w\}$, the i-th *column* of the picture is the word $p(1,i)p(2,i)\ldots p(h,i)$; similarly we define the j-th row for $j \in \{1,\ldots,h\}$. A language $K \subseteq \Sigma^*$ is *commutative* (*lossy*) if it is closed under reordering (respectively: removing) of letters. In the *commutative lossy tiling problem* we are given regular languages $K, L \subseteq \Sigma^*$ (the *column language* and the *row language*), where the column language K is commutative and lossy. A *solution* of the tiling system (K, L) is a picture p such that all columns in p belong to K and all rows in p belong to L. We ask whether, for all $h \in \mathbb{N}$, there exists a solution of height h. Notice that since K is commutative and lossy, we can reorder rows in a solution and again obtain a solution; we can also remove some rows and obtain a solution of smaller height. Consequently demanding solutions of each height $h \in \mathbb{N}$ amounts to demanding solutions of arbitrarily large height $h \in \mathbb{N}$.

3 From the Logic to Tiling Systems

The reduction from satisfiability of $\mathrm{AMSO}^{\mathrm{np}}$ to the multidimensional lossy tiling problem is given in [1], but we need to observe that the restriction to $\mathrm{AMSO}^{\mathrm{np}}_{2s}$ yields the commutative lossy tiling problem. Let us concentrate on the case where the formulae contain only one weight function; satisfiability of the general case easily reduces to this situation.

Before starting, we eliminate the outermost existential quantifier. Suppose that we have a formula $\psi = \exists t\forall s\exists r\,\varphi(r,s,t) \in \mathrm{AMSO}^{\mathrm{np}}_{2s}$. We create a formula $\psi' = \forall s\exists r\,\varphi'(r,s) \in \mathrm{AMSO}^{\mathrm{np}}_{2s}$

[1] See Proposition 14 in the appendix to [1], available at the authors' webpages.

using an additional unary predicate $\mathsf{small}(x)$: φ' is obtained from φ by replacing each atom $f(x) \leq t$ by $\mathsf{small}(x)$, and by replacing each subformula $s < f(x) \leq r$ by $s < f(x) \leq r \wedge \neg\mathsf{small}(x)$. It is easy to see that ψ is satisfiable if and only if ψ' is satisfiable. The idea is that small marks those positions on which the weight function f "is small".

Next, we apply the reduction of [1] to the formula ψ'. Let us explain briefly that the resulting tiling system is indeed a commutative lossy tiling system. The reduction is realized in three steps.

In the first step, the satisfiability of AMSO$^{\mathrm{np}}$ is reduced to the *limit satisfiability problem*. The idea is to chop an infinite word into infinitely many finite pieces that have the same theory (making repeated use of the Theorem of Ramsey). Originally, this is a theory with respect to all AMSO$^{\mathrm{np}}$ formulae up to some quantifier rank. We should replace it by the theory with respect to formulae where r and s are only used simultaneously, in subformulae of the form $s < f(x) \leq r$. Such theories have all compositionality properties needed for the proof which, thus, still goes through after this modification. The resulting formulae in the limit satisfiability problem test only for the theory of the finite words. So again r and s are only used simultaneously, in subformulae of the form $s < f(x) \leq r$.

In the second step, it is argued that a formula of the form $\forall s \exists r \varphi(r, s)$ is equivalent to $\forall s \varphi(s+1, s)$. This step is not affected by our modification.

In the third step, the limit satisfiability problem is reduced to the lossy tiling problem. First, we observe that, because there is just one universally quantified variable s, the resulting tiling system has dimension one. Then we have to slightly change the resulting tiling system to make it commutative. The alphabet of the system was $\Sigma \times \{<, =, >\}$, and the column language was $K = \bigcup_{a \in \Sigma} (a, <)^* ((a, =) \cup \varepsilon)(a, >)^*$. Intuitively, the meaning of a letter $(a, <)$ (or $(a, =)$, $(a, >)$) is that the row number is smaller (respectively: equal, greater) than the value of the weight function on this position (thus in each column initial rows contain $(a, <)$, then there is at most one $(a, =)$ marking the value of the weight function, and then we have $(a, >)$). Now in our formulae we cannot distinguish small values from big values, we can only test whether $s < f(x) \leq s+1$ holds. For this reason $(a, <)$ and $(a, >)$ become indistinguishable and can be replaced by one letter, call it $(a, \neq)$. The row language now becomes $K = \bigcup_{a \in \Sigma} (a, \neq)^* ((a, =) \cup \varepsilon)(a, \neq)^*$, which is commutative.

4 Monoids

In this section we slightly rephrase the problem of deciding the commutative lossy tiling problem using algebraic methods. Recall that every regular language (in particular the row language L) can be recognized by a morphism into a finite monoid. This means that there exists a morphism $\varphi \colon \Sigma^* \to M$ into a finite monoid M, and a set $F \subseteq M$ such that $L = \varphi^{-1}(F)$. It will be more convenient to label the picture directly with elements of M instead of Σ (using $\varphi(a)$ instead of a). The row language then becomes $\pi^{-1}(F)$, where the *evaluation* map $\pi \colon M^* \to M$ is the morphism defined by $\pi(s_1 \dots s_k) = s_1 \cdot \dots \cdot s_k$. The column language changes into $K' = \{\varphi(a_1) \dots \varphi(a_h) \mid a_1 \dots a_h \in K\}$, which again is commutative and lossy.

Next, we observe that we can restrict our considerations to sets F that are singletons. Namely, the tiling system $(K', \pi^{-1}(F))$ has arbitrarily high solutions if and only if for some $s \in F$ the system $(K', \pi^{-1}(s))$ has arbitrarily high solutions. Indeed, every solution of the latter system is a solution of the former. On the other hand, from a solution of $(K', \pi^{-1}(F))$ of height h we can choose rows evaluating to the most popular element $s_h \in F$ and obtain a solution of $(K', \pi^{-1}(s_h))$ of height at least $\frac{h}{|F|}$. Although elements s_h depend on h, some of them has to be used for infinitely many h (that is, for arbitrarily large h).

As a final simplification, let us analyze the column language. For a language L, let $L^{\downarrow}$ be the closure of L under removing letters (we add to L all words obtained by removing letters in words from L), and $L^{\circlearrowleft}$ the closure of L under reordering letters (we add to L all words obtained by reordering letters in words from L). A language (over M) is called a *base language* if it is of the form $(wA^*)^{\downarrow\circlearrowleft}$, where $A \subseteq M$ and $w \in (M \setminus A)^*$ (words in $(wA^*)^{\downarrow\circlearrowleft}$ can use letters from A arbitrarily many times, and letters from w at most as many times as they appear in w). Base languages play an important role in our proof. We use the letter ρ to denote base languages. Notice that the content of a base language $(wA^*)^{\downarrow\circlearrowleft}$ determines A uniquely, and w up to the order of its letters (with the assumption that w does not contain letters from A). The set A is called the *global part* of $\rho = (wA^*)^{\downarrow\circlearrowleft}$. We denoted it by $gl(\rho)$. The *norm* $\|\rho\|$ of a base language ρ is the length $|w|$.

It is a consequence of Higman's Lemma that every lossy language (over M) can be written as a finite union of languages of the form $(A_0^*b_1A_1^*\dots b_kA_k^*)^{\downarrow}$, where $A_0, \dots, A_k \subseteq M$ and $b_1, \dots, b_k \in M$. Our column language K is lossy and commutative, so it is a finite union of base languages. Summing up, we can restate our problem as follows:

Input: a finite monoid M, a finite set B of base languages over M, an element $s \in M$;

Question: does there exist for every $h \in \mathbb{N}$ a picture of height h each column of which belongs to $\bigcup B$ and every row of which to $\pi^{-1}(s)$?

For a picture p we define the *evaluation* of p, denoted $\pi(p)$, as the word of the same length as the height of p, whose i-th letter equals the evaluation of the i-th row of p. Then, instead of requesting that every row of p belongs to $\pi^{-1}(s)$, we can say that $\pi(p) \in s^*$.

5 The Decision Procedure

Our decision procedure maintains a set of base languages such that for every word from some of these languages there is a picture evaluating to this word where each column belongs to $\bigcup B$. New base languages are added following two kinds of schemas, the *product schema* and *diagonal schema.* These schemas are just ways of describing pictures of arbitrarily large size, evaluating to all words in some base language. The main difficulty is to prove completeness, i.e., showing that using some other fancy pictures one cannot obtain more base languages than we obtain using pictures generated from our schemas.

Let us now define the two kinds of schemas we use to generate new base languages. Let ρ_1 and ρ_2 be base languages. A *product schema* for ρ_1, ρ_2 is given by a picture q whose rows are divided into *special rows* and *global rows* such that (for $j \in \{1, 2\}$)

1. q has width 2 and the j-th column belongs to ρ_j, and
2. the height of q is at most $\|\rho_1\| + \|\rho_2\| + |M|^2$, and
3. the j-th letter of each global row belongs to $gl(\rho_j)$.

The base language *generated* by q is $(wA^*)^{\downarrow\circlearrowleft}$, where w consists of the letters of $\pi(q)$ corresponding to the special rows and A contains the letters of $\pi(q)$ corresponding to the global rows. We only allow schemas q for which w does not contain letters from A.

While defining a diagonal schema we need to use the power-set monoid. The set $\mathcal{P}(M)$ of subsets of M has a natural monoid structure: $C \cdot D = \{c \cdot d \mid c \in C, d \in D\}$. We say that a set of base languages B is *uniform* when it is nonempty, for all $\rho_1, \rho_2 \in B$ we have $gl(\rho_1) = gl(\rho_2)$, and this set is idempotent. For a uniform B we write $gl(B)$ for the set $gl(\rho)$ with $\rho \in B$. The set of all finite uniform sets of base languages over M is denoted by $UBL(M)$.

x	a	c	z	x
a	b	a	c	c
b	c	y	b	a

y	z	x	x	z	y	z	x	x	a	c	z	x
z	x	z	x	y	a	c	z	z	x	z	y	x
x	a	c	x	x	z	z	y	x	z	x	z	x
a	b	a	c	a	b	a	c	a	b	a	c	c
b	c	z	y	z	y	z	x	y	x	x	b	a

Figure 1 On the left we have an example diagonal schema. Elements of $gl(B)$ are shaded in gray. The first row is a global row, and the other two are special rows (we suppose that $a \cdot b \cdot a \cdot c$ is idempotent). The double line divides the schema horizontally into two pictures. On the right there is a picture created out of the schema for $n = 3$. Here double lines are introduced only for readability. Gray cells are stretched into longer areas evaluating to the same value (e.g. $x = z \cdot x \cdot z \cdot x \cdot y$).

Let B be a uniform set of base languages. A *diagonal schema* for B is given by a picture q whose rows are divided into *special rows* and *global rows* and which is divided horizontally into pictures $q_1, \dots, q_k$ (which means that $q_1, \dots, q_k$ have as many rows as q, and the i-th row of q is the concatenation of the i-th rows of $q_1, \dots, q_k$) such that:

1. each column of q belongs to $\bigcup B$, and
2. each special row of each q_j either has length 1, or evaluates to an idempotent, or it contains a letter belonging to $gl(B)$, and
3. the first and the last letter of each global row of each q_j belongs to $gl(B)$.

The base language *generated* by q is $(wA^*)^{\downarrow\circlearrowright}$ where w consists of the letters of $\pi(q)$ corresponding to the special rows and A contains the letters of $\pi(q)$ corresponding to the global rows. Again, we only allow schemas q for which w does not contain letters from A. An example diagonal schema is depicted in Figure 1 on the left.

The following theorem states soundness and completeness of our schemas.

▶ **Theorem 5.1.** *Let B_0 be a finite set of base languages over a monoid M. For a function η: $UBL(M) \to \mathbb{N}$ let $B_0^{\leq\eta} = B_0$ and for each $i > 0$, inductively, let $B_i^{\leq\eta}$ be the set of all base languages ρ such that*

- *$\rho \in B_{i-1}^{\leq\eta}$, or*
- *ρ is generated by some product schema for some base languages $\rho_1, \rho_2 \in B_{i-1}^{\leq\eta}$, or*
- *ρ is generated by some diagonal schema for a uniform set of base languages $B \subseteq B_{i-1}^{\leq\eta}$, of width and height at most $\eta(B)$.*

Then there is a computable function η: $UBL(M) \to \mathbb{N}$ such that for every $s \in M$ the following two statements are equivalent.

- *For each $h \in \mathbb{N}$, there exists a picture p of height h such that $\pi(p) \in s^*$ and each column belongs to $\bigcup B_0$.*
- *For $x = 3 \cdot (2^{|M|} + 1)^2$, there exists a base language $\rho \in B_x^{\leq\eta}$ with $s \in gl(\rho)$.*

Notice that this theorem implies the decidability of the commutative lossy tiling problem. Indeed, given $B_{i-1}^{\leq\eta}$ we can calculate $B_i^{\leq\eta}$ because the number of product and diagonal schemas to consider is finite (the size of product schemas is bounded by definition, and the size of diagonal schemas is bounded by the function η).

6 Soundness

In this section we prove the easier direction of Theorem 5.1: the implication from the second to the first statement. The proof is based on the following two lemmas.

▶ **Lemma 6.1.** *Let ρ be a base language generated by some product schema for ρ_1, ρ_2 and let $u \in \rho$. Then there exists a picture p each column of which belongs to $\rho_1 \cup \rho_2$ and such that $\pi(p) = u$.*

▶ **Lemma 6.2.** *Let ρ be a base language generated by some diagonal schema for a uniform set of base languages B and let $u \in \rho$. Then there exists a picture p each column of which belongs to $\bigcup B$ and such that $\pi(p) = u$.*

Using these the lemmas we can prove the soundness implication of Theorem 5.1 as follows. Let $B_i^{\leq\eta}$ be the sets from Theorem 5.1. The function η bounding the sizes of diagonal schemas does not matter for this implication. We will prove by induction on i that if $u \in \bigcup B_i^{\leq\eta}$, then there exists a picture p each column of which belongs to $\bigcup B_0$ and such that $\pi(p) = u$. Then the statement of the lemma follows by taking $u = s^h$ since we have $s \in gl(\rho)$ for some $\rho \in B_x^{\leq\eta}$, which implies that $u \in \bigcup B_x^{\leq\eta}$.

For $i = 0$ the claim is trivial: we can take a picture p containing u as the only column. For $i > 0$, let $u \in \bigcup B_i^{\leq\eta}$. Then $u \in \rho$ for some $\rho \in B_i^{\leq\eta}$. If $\rho \in B_{i-1}^{\leq\eta}$ the claim follows by inductive hypothesis. Otherwise, we are in the second or the third case of definition of $B_i^{\leq\eta}$. Hence, we can apply Lemma 6.1 or 6.2 to obtain a picture p' each column of which belongs to $\bigcup B_{i-1}^{\leq\eta}$ and such that $\pi(p') = u$. Moreover, by inductive hypothesis there exists, for each column u_j of p', a picture p_j each column of which belongs to $\bigcup B_0$ and such that $\pi(p_j) = u_j$. To obtain p we replace in p' the j-th column u_j by p_j, for every j. Then $\pi(p) = \pi(p')$ and p has the desired properties.

In the remaining part of this section we prove Lemmas 6.1 and 6.2.

Proof of Lemma 6.1. The proof follows immediately from the definitions. Let q be a product schema for ρ_1, ρ_2 which generates ρ. Since the global rows of q contain only letters from the global parts of ρ_1, ρ_2, we can duplicate in q any global row without destroying the property that the j-th column belongs to ρ_j. We can also remove any row and reorder the rows. By performing such operations we can obtain a picture p such that $\pi(p) = u$. ◀

Proof of Lemma 6.2. Let $\rho = (wA^*)^{\downarrow\circlearrowleft}$, let q be a diagonal schema for B generating ρ, and let $q_1, \ldots, q_k$ be the pictures into which q is divided. W.l.o.g. we assume that each global row of q evaluates to a different element of A (otherwise we remove redundant rows). Note that, if the lemma holds for some word u, then it holds also for any u' obtained from u by removing and reordering letters (because we can remove and reorder the rows of the resulting picture). Thus it is enough to consider, for each $n \in \mathbb{N}$, a column u which begins by w and then has each letter of A repeated n times.

The idea of constructing a picture p out of the diagonal schema q is depicted in Figure 1. For each $j \in \{1, \ldots, k\}$ we create a picture p_j by modifying q_j as follows. p_j will have $|A| \cdot (n-1)$ more rows than q_j; more precisely, each global row of q_j will produce n rows of p_j, while each special row of q_j will produce only one row of p_j. Fix some j and let m be the width of q_j. If $m = 1$, we just replace each global row by n copies. Assume now that $m > 1$. Then the width of p_j will be nm.

We start by considering a special row v. If $\pi(v)$ is idempotent, we can just repeat the content of the row n times without changing the value of the product. Otherwise, by definition there exists an index i such that the i-th letter of v belongs to $gl(B)$. As the first $i-1$ letters of the new row we take the first $i-1$ letters of v. As the last $m-i$ letters of the new row we take the last $m-i$ letters of v. On the remaining $mn-m+1$ positions we place letters from $gl(B)$ in such a way that their product is equal to the i-th letter of v (this is possible since $gl(B)$ is idempotent by uniformity of B). Again, the value of the product remains unchanged.

Finally, we consider a global row v of q_j. We will produce n rows in p_j; the i-th of them, for $i \in \{1, \dots, n\}$, is created in the following way. On the first $(i-1)m+1$ positions of the new row we place letters from $gl(B)$ in such a way that their product is equal to the first letter of v (recall that by definition the first and the last letter of v are in $gl(B)$). On the last $(n-i)m+1$ positions of the new row we place letters from $gl(B)$ in such a way that their product is equal to the last letter of v. On the remaining $m-2$ positions we put the middle $m-2$ letters of v, without the first and the last letter.

For the picture p we take the concatenation of $p_1, \dots, p_k$ (which means that the i-th row of p is obtained by concatenating the i-th rows of $p_1, \dots, p_k$). We observe that the evaluation of p is u (the rows created out of special rows evaluate to w, and the rows created out of global rows evaluate to elements of A, each n times). It remains to observe that each column of p (so of each p_j) belongs to $\bigcup B$. When p_j has only one column, this is clear, because it is obtained by duplicating some letters from $gl(B)$ in a column from $\bigcup B$. Otherwise (with m as above), the column with number $i + i'm$ of p_j (for $i \in \{1, \dots, m\}$) is obtained from the column number i of q_j (which is in $\bigcup B$): the letters which are not in $gl(B)$ are taken at most once, on the other positions we take some letters from $gl(B)$. Thus the new column is also in $\bigcup B$. ◀

7 Completeness

In this section we prove the remaining direction of Theorem 5.1: the implication from the first to the second statement. The strategy is as follows. First we consider special cases that can be described by a single schema. In Section 7.1 we analyze pictures of width 2. For these one can extract a product schema. In Section 7.2 we analyze pictures whose columns come from a union of a uniform set of base languages. These can be turned into a diagonal schema. As a technical tool we introduce in Section 7.3 a new version of the Factorization Trees Theorem [17]. This theorem is used in Section 7.4 to decompose arbitrary picture into simple fragments corresponding to single schemas, which allows us to conclude the proof. During the whole section we consider the monoid M as fixed.

7.1 Products

We start by analyzing width 2 pictures in order to turn them into product schemas.

▶ **Lemma 7.1.** *Let ρ_1, ρ_2 be two base languages and let p be a picture of width 2 such that the first column belongs to ρ_1 and the second one to ρ_2. Then there exists a product schema for ρ_1, ρ_2 which generates a base language ρ such that $\pi(p) \in \rho$ and $gl(\rho) = gl(\rho_1) \cdot gl(\rho_2)$.*

Proof. We take $\rho = (wA^*)^{\downarrow\circlearrowleft}$ where $A = gl(\rho_1) \cdot gl(\rho_2)$ and w consists of those letters of $\pi(p)$ which are not in A (taken as many times as they appear in $\pi(p)$). Obviously $\pi(p) \in \rho$. In q we include all rows of p that do not evaluate to an element of A. These will be the special rows. Note that in each of these rows either the first letter does not belong to $gl(\rho_1)$, or the second letter does not belong to $gl(\rho_2)$. Thus we have at most $\|\rho_1\| + \|\rho_2\|$ of such rows. Moreover, for each $r \in gl(\rho_1)$ and each $s \in gl(\rho_2)$, we add to q a row with r in the first column and s in the second one. These will be the global rows. We have $|gl(\rho_1)| \cdot |gl(\rho_2)| \leq |M|^2$ of them. We see that q is a product schema for ρ_1, ρ_2 that generates ρ. ◀

7.2 Uniform Case

Next, we consider a special case when the set of base languages allowed in columns is uniform, and we show that such a picture can be transformed into a single diagonal schema.

▶ **Lemma 7.2.** *There is a computable function $\eta\colon UBL(M) \to \mathbb{N}$ such that, for every finite uniform set of base languages B and every picture p each column of which belongs to $\bigcup B$, there exists a diagonal schema for B of width and height at most $\eta(B)$ that generates a base language ρ such that*
- *$\pi(p) \in \rho$ and*
- *$E = gl(B)$ and $A = gl(\rho)$ satisfy $E \subseteq A = E \cdot A \cdot E$.*

Let us comment on the second condition ($E \subseteq A = E \cdot A \cdot E$). It enforces that the base language ρ (and hence also the diagonal schema) is more robust. This will be useful later. Namely, the global part of ρ contains not only the letters that appear many times in $\pi(p)$, but also all letters from $gl(B)$ (since $E \subseteq A$) and all results of surrounding the former letters by letters from $gl(B)$ (since $E \cdot A \cdot E \subseteq A$). Note that we always have $A \subseteq E \cdot A \cdot E$, as each global row begins and ends by a letter from $gl(B)$.

The proof of the lemma is based the following fact saying that each word can be chopped into a small number of idempotents and single letters. To simplify notation, we write $\exp(x)$ for 2^x.

▶ **Fact 7.3.** *Let M' be a finite monoid and w a word over M'. Then we can divide w into fragments $w = w_1 \dots w_k$ for $k \leq \exp(3|M'|)$ such that, for every i, either $|w_i| = 1$, or $\pi(w_i)$ is idempotent.*

This fact is applied to a picture, in order to split it horizontally as in a diagonal schema. While reading the next lemma have in mind that E will be used for $gl(B)$.

▶ **Lemma 7.4.** *Let p be a picture and $E \subseteq M$. Let x be the number of rows of p which contain only letters from $M \setminus E$ and let y be the smallest number such that in each column of p there are at most y positions containing a letter from $M \setminus E$. Then, for some $k \leq \exp(3(y-x+1)|M|^y)$, we can divide p horizontally into pictures $p_1, \dots, p_k$ in such a way that each row of each p_j either has length 1, or evaluates to an idempotent, or contains a letter from E.*

Proof. We prove the claim by induction on $y - x$ (note that $x \leq y$). Consider the monoid $M' = M^x$ with coordinatewise multiplication. Let I be the set of (numbers of) those rows which contain only letters from E (by definition $|I| = x$). Let $w \in (M')^*$ be the word consisting of the rows of p which are in I (each letter contains the elements of M appearing in the x rows of a column). Applying Fact 7.3 to w, we obtain a factorisation $w = w_1 \dots w_m$ for $m \leq \exp(3|M|^x) \leq \exp(3|M|^y)$ where each w_j either has length 1, or evaluates to an idempotent. We divide p into $p'_1, \dots, p'_m$ in the same way: the width of p'_j is the same as the length of w_j. Then every row of each p'_j which is in I either has length 1, or evaluates to an idempotent. For each p'_j we proceed in one of two ways.
- If each row of p'_j which is not in I contains a letter from E, this p'_j satisfies the statement of the lemma.
- Otherwise, there exists a row of p'_j not in I which contains only letters from $M \setminus E$. Then $x' \geq x + 1$ and $y' \leq y$, where x' is the number of rows of p'_j which contain only letters from $M \setminus E$ and y' is the smallest number such that in each column of p'_j there are at most y' positions containing a letter from $M \setminus E$. We use the inductive hypothesis for p'_j to obtain a subdivision of p'_j as required by the statement of the lemma.

Since each of the subdivisions returns at most $\exp(3(y'-x'+1)|M|^{y'}) \le \exp(3(y-x)|M|^y)$ pictures, in total we have at most $m\cdot\exp(3(y-x)|M|^y) \le \exp(3(y-x+1)|M|^y)$ pictures. ◀

Proof of Lemma 7.2. Set $E = gl(B)$. First, we divide p into pictures $p_1,\dots,p_k$ by applying Lemma 7.4 to the picture p and to the set E. Note that the number y in the statement of the lemma is equal to the maximal norm of a base language in B, and that $x \ge 0$. We have $k \le \exp(3(y-x+1)|M|^y) \le \exp(3(y+1)|M|^y)$. Let I_1 be the set of all those numbers i of rows of p such that the first or the last letter of the i-th row of some p_j is in $M \setminus E$. Note that $|I_1| \le 2ky$ (where y is again the maximal norm of a base language in B): we look for letters from $M \setminus E$ only in $2k$ columns (the first and the last column of each p_j), and in each of these columns we have at most y letters from $M \setminus E$. The picture p with this division is almost a diagonal schema as needed (when the rows from I_1 are treated as the special rows). However we still need to reduce its size and ensure that $E \subseteq A = E \cdot A \cdot E$.

For each i, we denote by s_i the evaluation of the i-th row without the first and the last letter (so the value of the i-th row can be obtained by multiplying its first letter by s_i and by its last letter). Let I_2 be the set of numbers $i \notin I_1$ of rows of p such that there are less than $|E|^2$ numbers $j \notin I_1$ for which $s_i = s_j$. Notice that $|I_2| \le |M|^3$ (we have at most $|E|^2 - 1 \le |M|^2$ rows for each of $|M|$ possible values of s_i). Set $I = I_1 \cup I_2$.

Next, let A' be the set of s_i for all $i \notin I$. Let $A = (E \cdot A' \cdot E) \cup E$ and let w contain those letters of $\pi(p)$ which are not in A (as many times as they appear in $\pi(p)$); we take $\rho = (wA^*)^{\downarrow\circlearrowleft}$. As E is idempotent, it follows that $\pi(p) \in \rho$ and $E \subseteq A = E \cdot A \cdot E$. It remains to construct a diagonal schema q for B that generates ρ.

The width of q will be the same as of p. We also divide q into $q_1,\dots,q_k$ of the same widths as $p_1,\dots,p_k$. We include in q all those rows of p which do not evaluate to an element of A. These will be the special rows. Note that by the statement of Lemma 7.4, any row of p can be taken as a special row: inside each p_j it either has length 1, or evaluates to an idempotent, or it contains a letter belonging to E. Moreover, all these rows are in I; indeed, any other row $i \notin I$ evaluates to $r \cdot s_i \cdot r'$, where $s_i \in A'$ and r, r' are the first and the last letter of the row, which are in E by definition of I_1. Consequently, there are at most $|I|$ such rows.

Then, for each $s \in A'$ we consider $|E|^2$ rows $i \notin I$ for which $s_i = s$ (we have at least $|E|^2$ such rows by definition of I_2) and we modify them as follows. For each pair $r, r' \in E$ we add to q one such row in which we replace the first letter by r and the last letter by r'. These will be global rows. This works as the first and the last letter of each such row inside each p_j belong to E and the replaced letters are also in E. Additionally, for each $s \in E$, we add to q a row containing only letters from E, which evaluates to s (as E is idempotent, we can find such rows of every desired length). These will also be global rows. This works since all letters of these rows are in E.

We see that every column of q belongs to $\bigcup B$: it is a column of p with some letters removed and some letters from E added. The special rows evaluate exactly to the letters of w. The global rows of the first kind evaluate to all elements of $E \cdot A' \cdot E$, and the global rows of the second kind to all elements of E. Thus q generates the base language ρ.

It remains to bound the size. The number of rows in q is at most

$$|I| + |E|\cdot|A'|\cdot|E| + |E| \le 2ky + 2|M|^3 + |M| \le 2y\cdot\exp(3(y+1)|M|^y) + 3|M|^3,$$

where y is the maximal norm of a base language in B. We denote the last number by $\theta(B)$ (it depends only on B and $|M|$).

We also have to restrict the width of q. Since we have started from an arbitrary picture p, the width can be arbitrary; so we have to remove some columns. Fix some q_j that has more than one column. In each special row whose value is not idempotent there is some letter from E. In each such row we choose one of these letters and we mark the column containing it (we don't want to remove this column). We also mark the first and the last column of q_j; they contain letters from E in global rows, so we also don't want to remove them. We have marked at most $\theta(B)+2$ columns. We want to remove some not-marked columns, so that the resulting picture evaluates to the same word. For each number of columns i, consider the picture consisting of the first i columns of q_j; let w_i be the evaluation of this picture (w_i is a word in M^h, where $h \leq \theta(B)$ is the height of q_j). Whenever $w_i = w_l$ for some $i < l$, we can remove the columns number $i+1, \dots, l$, and the whole new picture will still evaluate to $\pi(q_j)$; we do this only when none of these columns is marked. We repeat this removal procedure as long as such pair of indices i, l exists. By the Pigeon Hole Principle, among any $|M|^h + 1$ numbers we can find two i, l for which $w_i = w_l$. Thus, after the removal, we have at most $(\theta(B)+1) \cdot (|M|^h + 1) + 1$ columns in q_j. Because we do not remove marked columns, the properties of a diagonal schema are preserved. In total we have at most $k \cdot ((\theta(B)+1) \cdot (|M|^h+1)+1) \leq \exp(3(y+1)|M|^y) \cdot ((\theta(B)+1) \cdot (|M|^h+1)+1)$ columns. We denote the last number by $\eta(B)$. Note that $\theta(B) \leq \eta(B)$. Thus, not only the width but also the height of q is bounded by $\eta(B)$. ◀

7.3 Factorization Trees

In this subsection we present a new generalization of the Factorization Trees Theorem [17]. In this generalization the result in an "idempotent" node depends on some additional data in the arguments. This theorem will be used in Section 7.4 to decompose an arbitrary picture into pictures of the special form considered in Sections 7.1 and 7.2.

The nodes of our factorization trees will be labelled by elements of some set D, possibly infinite. We also have a finite monoid M' and a projection $\sigma\colon D \to M'$. The construction is parameterized by two functions. The function $pr\colon D^2 \to D$ describes a product. The other function

$$st\colon \{d_1 \dots d_c \in D^+ \mid \sigma(d_1) = \dots = \sigma(d_c) \text{ is idempotent}\} \to D$$

describes an operation which will be used in idempotent nodes. We require that these functions satisfy the following axioms:

$$(*) \quad \sigma(pr(a,b)) = \sigma(a) \cdot \sigma(b)\,, \qquad \text{for all } a, b \in D,$$

$$(**) \quad \sigma(st(d_1 \dots d_c)) = \sigma(d_1) \quad \text{or} \quad \sigma(st(d_1 \dots d_c)) <_{\mathcal{J}} \sigma(d_1)\,, \qquad \text{for all } d_1 \dots d_c \in dom(st).$$

The preorder $\leq_{\mathcal{J}}$ in the second axiom is defined by $r \leq_{\mathcal{J}} s$ if there are u_1, u_2 such that $r = u_1 \cdot s \cdot u_2$ (recall that each monoid contains an identity element, that is allowed as u_1 and u_2). Two elements are $\mathcal{J}$-equivalent, denoted $r \sim_{\mathcal{J}} s$, when $r \leq_{\mathcal{J}} s$ and $s \leq_{\mathcal{J}} r$. Equivalence classes of this relation are called $\mathcal{J}$-classes. We write $r <_{\mathcal{J}} s$ when $r \leq_{\mathcal{J}} s$, but $r \not\sim_{\mathcal{J}} s$. A *factorization tree* is a tree labelled by elements of D whose nodes are of one of three forms:

- a *leaf*,
- a *binary node* with exactly two children; it is labelled by $pr(d_1, d_2)$, where d_1, d_2 are the labels of its children,
- an *idempotent node* with at least three children labelled by $d_1, \dots, d_c$ such that $\sigma(d_1) = \dots = \sigma(d_c)$ is idempotent; the node itself is labelled by $st(d_1 \dots d_c)$.

The word (in D^+) read from the leaves of a factorization tree t (from left to right) is called the *input* of t, and the label of the root of t is called its *output*.

Note that standard factorization trees as in [17] can be obtained by taking $D = M'$ and $st(e \dots e) = e$. In computation trees for a stabilization monoid [12], we again have $D = M'$, but $st(e \dots e)$ now depends on the number of arguments: it is e for short sequences $e \dots e$, and $e^\sharp$ for longer $e \dots e$. The key result is the existence of factorization trees of constant height as described in the following theorem.

▶ **Theorem 7.5.** *For every $v \in D^+$, there exists a factorization tree with input v and height at most*[2] $3(|M'|+1)^2$.

This theorem can be proved basically in the same way as its version for stabilization monoids ([12], Theorem 3.3): the tree is constructed in a bottom-up way, so it is not a problem that the result in an idempotent node depends in some way on the subtree constructed below.

7.4 The Final Argument

In this subsection we conclude our proof of the missing implication of Theorem 5.1. The function η is taken from Lemma 7.2. Let $B_i^{\leq\eta}$ be sets of base languages as in Theorem 5.1, for some finite set of base languages B_0. Each $B_i^{\leq\eta}$ is finite. Let h be the smallest number greater than the norm of each base language in $B_x^{\leq\eta}$, where $x = 3 \cdot (2^{|M|}+1)^2$. Take some picture p of height h each column of which belongs to $\bigcup B_0$ and for which $\pi(p) \in s^*$. Our goal is to find $\rho \in B_x^{\leq\eta}$ such that $s \in gl(\rho)$.

We use the theorem about factorization trees from the previous subsection. For D we take the set of pairs (w, ρ), where $w \in M^h$ and ρ is a base language containing w. We set $M' = \mathcal{P}(M)$ and $\sigma((w,\rho)) = gl(\rho)$. It remains to define the functions pr and st.

To define pr, consider two letters (w_1, ρ_1) and (w_2, ρ_2) from D. Let p be the picture with two columns w_1 and w_2. By Lemma 7.1, there exists a base language ρ such that $\pi(p) \in \rho$, $gl(\rho) = gl(\rho_1) \cdot gl(\rho_2)$, and there exists a product schema for ρ_1, ρ_2 generating ρ. We define $pr((w_1,\rho_1),(w_2,\rho_2)) = (\pi(p), \rho)$. Then Axiom $(*)$ is satisfied because $gl(\rho) = gl(\rho_1) \cdot gl(\rho_2)$. Observe also that when $\rho_1, \rho_2 \in B_j^{\leq\eta}$, for some j, then $\rho \in B_{j+1}^{\leq\eta}$.

To define st, consider elements $(w_1, \rho_1) \dots (w_k, \rho_k) \in D^+$ such that $gl(\rho_1) = \dots = gl(\rho_k)$ is idempotent. Let p be the picture with k columns $w_1, \dots, w_k$, set $B = \{\rho_1, \dots, \rho_k\}$, and let $E = gl(B)$. Then B is a uniform set of base languages and each column of p belongs to $\bigcup B$. By Lemma 7.2, there exists a base language ρ such that $\pi(p) \in \rho$, $E \subseteq gl(\rho) = E \cdot gl(\rho) \cdot E$, and there exists a diagonal schema for B of width and height at most $\eta(B)$ generating ρ. We define $st((w_1,\rho_1)\dots(w_k,\rho_k)) = (\pi(p), \rho)$. Observe that when $\rho_i \in B_j^{\leq\eta}$, for some j and all i, then $\rho \in B_{j+1}^{\leq\eta}$. Axiom $(**)$ is satisfied due to the following fact.

▶ **Fact 7.6.** *Let $E, A \subseteq M$ where E is idempotent and $E \subseteq A = E \cdot A \cdot E$. Then either $A = E$ or $A <_{\mathcal{J}} E$.*

To conclude the proof, recall that p is a picture of height h each column of which belongs to $\bigcup B_0$ and such that $\pi(p) \in s^*$. We want to find a base language $\rho \in B_x^{\leq\eta}$ with $s \in gl(\rho)$. Consider a word $w = (d_1, \rho_1) \dots (d_m, \rho_m) \in D^+$, where d_i is the i-th column of p and $\rho_i \in B_0$ is some base language with $d_i \in \rho_i$. By Theorem 7.5 there exists a factorization tree t with height at most x and input w. Let (d, ρ) be its output. Note that $d = \pi(p) = s^h$ (by definition of pr and st), and $d \in \rho$ (by definition of D). Moreover, $\rho \in B_x^{\leq\eta}$ (more generally, when a root of a subtree of height at most i is labelled by some (d', ρ'), then $\rho' \in B_i^{\leq\eta}$). As h is greater than the size of ρ, we have $s \in gl(\rho)$, which is what we wanted to prove.

---

[2] One can obtain a bound of $3|M'|$, but this requires a more complicated proof.

References

1. Achim Blumensath, Olivier Carton, and Thomas Colcombet. Asymptotic monadic second-order logic. In Erzsébet Csuhaj-Varjú, Martin Dietzfelbinger, and Zoltán Ésik, editors, *Mathematical Foundations of Computer Science 2014 - 39th International Symposium, MFCS 2014, Budapest, Hungary, August 25-29, 2014. Proceedings, Part I*, volume 8634 of *Lecture Notes in Computer Science*, pages 87–98. Springer, 2014.
2. Mikołaj Bojańczyk. The finite graph problem for two-way alternating automata. *Theor. Comput. Sci.*, 3(298):511–528, 2003.
3. Mikołaj Bojańczyk. Weak MSO with the unbounding quantifier. *Theory Comput. Syst.*, 48(3):554–576, 2011.
4. Mikołaj Bojańczyk and Thomas Colcombet. Bounds in w-regularity. In *21th IEEE Symposium on Logic in Computer Science (LICS 2006), 12-15 August 2006, Seattle, WA, USA, Proceedings*, pages 285–296. IEEE Computer Society, 2006.
5. Mikołaj Bojańczyk, Paweł Parys, and Szymon Toruńczyk. The MSO+U theory of (N, <) is undecidable. *CoRR*, abs/1502.04578, 2015. Accepted to STACS 2016.
6. Mikołaj Bojańczyk and Szymon Toruńczyk. Weak MSO+U over infinite trees. In Christoph Dürr and Thomas Wilke, editors, *29th International Symposium on Theoretical Aspects of Computer Science, STACS 2012, February 29th - March 3rd, 2012, Paris, France*, volume 14 of *LIPIcs*, pages 648–660. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2012.
7. Tomás Brázdil, Krishnendu Chatterjee, Antonín Kucera, and Petr Novotný. Efficient controller synthesis for consumption games with multiple resource types. In P. Madhusudan and Sanjit A. Seshia, editors, *Computer Aided Verification - 24th International Conference, CAV 2012, Berkeley, CA, USA, July 7-13, 2012 Proceedings*, volume 7358 of *Lecture Notes in Computer Science*, pages 23–38. Springer, 2012.
8. J. Richard Büchi. On a decision method in a restricted second order arithmetic. In *Logic, Methodology and Philosophy of Science (Proc. 1960 Internat. Congr.)*, pages 1–11, Stanford, Calif., 1962. Stanford Univ. Press.
9. Claudia Carapelle, Alexander Kartzow, and Markus Lohrey. Satisfiability of CTL* with constraints. In Pedro R. D'Argenio and Hernán C. Melgratti, editors, *CONCUR 2013 - Concurrency Theory - 24th International Conference, CONCUR 2013, Buenos Aires, Argentina, August 27-30, 2013. Proceedings*, volume 8052 of *Lecture Notes in Computer Science*, pages 455–469. Springer, 2013.
10. Thomas Colcombet. Factorisation forests for infinite words. In Erzsébet Csuhaj-Varjú and Zoltán Ésik, editors, *FCT*, volume 4639 of *Lecture Notes in Computer Science*, pages 226–237. Springer, 2007.
11. Thomas Colcombet. The theory of stabilisation monoids and regular cost functions. In Susanne Albers, Alberto Marchetti-Spaccamela, Yossi Matias, Sotiris E. Nikoletseas, and Wolfgang Thomas, editors, *Automata, Languages and Programming, 36th Internatilonal Collogquium, ICALP 2009, Rhodes, greece, July 5-12, 2009, Proceedings, Part II*, volume 5556 of *Lecture Notes in Computer Science*, pages 139–150. Springer, 2009.
12. Thomas Colcombet. Regular cost functions, part I: logic and algebra over words. *Logical Methods in Computer Science*, 9(3), 2013.
13. Nathanaël Fijalkow and Martin Zimmermann. Cost-parity and cost-streett games. *Logical Methods in Computer Science*, 10(2), 2014.
14. Szczepan Hummel and Michal Skrzypczak. The topological complexity of MSO+U and related automata models. *Fundam. Inform.*, 119(1):87–111, 2012.
15. Daniel Kirsten. Distance desert automata and the star height problem. *ITA*, 39(3):455–509, 2005.

16 Orna Kupferman, Nir Piterman, and Moshe Y. Vardi. From liveness to promptness. *Formal Methods in System Design*, 34(2):83–103, 2009.
17 Imre Simon. Factorization forests of finite height. *Theor. Comput. Sci.*, 72(1):65–94, 1990.
18 Szymon Toruńczyk. *Languages of profinite words and the limitedness problem.* PhD thesis, Warsaw University, 2011.

A Appendix to Section 3

Let us explain in more detail the fact stated in Section 3 saying that ψ is satisfiable if and only if ψ' is satisfiable. Recall that $\psi = \exists t \forall s \exists r\, \varphi(r,s,t)$ and $\psi' = \forall s \exists r\, \varphi'(r,s)$, where φ' is obtained from φ by replacing each atom $f(x) \leq t$ by $\mathsf{small}(x)$, and by replacing each subformula $s < f(x) \leq r$ by $s < f(x) \leq r \wedge \neg\mathsf{small}(x)$.

Suppose that we have a weighted infinite word $\langle w, f\rangle$ that is a model for ψ. This gives some value of t for which $\forall s \exists r\, \varphi(r,s,t)$ is true in $\langle w, f\rangle$. To obtain a model $\langle w', f\rangle$ for ψ', it is enough to mark by small those positions x where $f(x) \leq t$. Then clearly for every $s \geq t$ the formula $\exists r\, \varphi'(r,s)$ holds in $\langle w', f\rangle$, since $f(x) \leq t$ in $\langle w, f\rangle$ implies $\mathsf{small}(x)$ in $\langle w', f\rangle$ (for the same position x), and $s < f(x) \leq r$ in $\langle w, f\rangle$ implies $s < f(x) \leq r \wedge \neg\mathsf{small}(x)$ in $\langle w', f\rangle$ (recall that these subformulae appear only positively). But since all comparisons with s are $s < f(x)$ appearing positively, the formula $\exists r\, \varphi'(r,s)$ holds even more for smaller s, thus $\psi' = \forall s \exists r\, \varphi'(r,s)$ holds in $\langle w', f\rangle$.

Conversely, suppose that $\langle w', f'\rangle$ is a model for ψ'. In a model $\langle w, f\rangle$ for ψ we take $f(x) = 0$ if $\mathsf{small}(x)$ holds, and $f(x) = f'(x)$ otherwise (and we remove the predicate small). For $t = 0$ we have that $\mathsf{small}(x)$ in $\langle w', f'\rangle$ implies $f(x) \leq t$ in $\langle w, f\rangle$ and $s < f(x) \leq r \wedge \neg\mathsf{small}(x)$ in $\langle w', f'\rangle$ implies $s < f(x) \leq r$ in $\langle w, f\rangle$. Thus ψ holds in $\langle w, f\rangle$.

B Factorization Trees

In this section we prove Theorem 7.5. As we have said, a proof of this theorem can be obtained by minor modifications in the proof for the stabilization monoid case ([12], Theorem 3.3). Here, instead of repeating that proof, we base on the standard factorization trees theorem (see e.g. [10], Theorem 1). This theorem only deals with the case when $D = M'$ and $\sigma(s) = s$. However a factorization tree for this case remains correct (after relabeling its nodes) for any D and σ such that $\sigma(st(d_1 \dots d_c)) = \sigma(d_1)$, as stated below.

▶ **Theorem B.1** ([10]). *Assume that $\sigma(st(d_1 \dots d_c)) = \sigma(d_1)$ for each $d_1 \dots d_c$ in the domain of st. Let $v \in D^+$. Then there exists a factorization tree with input v and height at most $3|M'|$.*

Next, we show how to repair the factorization tree obtained in the above theorem when the operation st changes. The first auxiliary lemma deals with a single $\mathcal{J}$-class.

▶ **Lemma B.2.** *Let J be a $\mathcal{J}$-class of M', and let $v \in D^+$. Then there exist factorization trees $t_1, \dots, t_k$ with height at most $3|M'|$, such that the concatenation of their inputs gives v, and whenever some t_i for $i \in \{1, \dots, k-1\}$ has output in $\sigma^{-1}(J)$, then t_{i+1} has output outside $\sigma^{-1}(J)$.*

Proof. The proof is by induction on the length of v. One case is that there exists some infix w (where $v = uwu'$) for which there exists a factorization tree t with input w, height at most $3|M'|$, and output outside $\sigma^{-1}(J)$. Then we use the induction assumption for the shorter words u and u' (if nonempty); the trees over these words together with t give the thesis.

The remaining case is that for no infix w of v there exists a factorization tree with input w, height at most $3|M'|$, and output outside $\sigma^{-1}(J)$. This in particular means that each letter of v is in $\sigma^{-1}(J)$ (otherwise we can construct a one-node factorization tree with this letter as input and with output outside $\sigma^{-1}(J)$). Consider the operation st' defined by

$$st'(d_1 \dots d_k) = \begin{cases} st(d_1 \dots d_k) & \text{when } \sigma(st(d_1 \dots d_k)) = \sigma(d_1), \\ d_1 & \text{otherwise.} \end{cases}$$

We construct a factorization tree t with input v using Theorem B.1 for the operation st' instead of st. We will prove that t is a correct factorization tree also for the original st function (that is, we always use only the first case in the definition of st'); this will finish the proof: we take $k = 1$ and $t_1 = t$. Assume the contrary: fix some idempotent node x of t, for which $\sigma(st(d_1 \dots d_c)) \neq \sigma(d_1)$, where $d_1, \dots, d_c$ are the labels of the children of x, and such that no descendant of x has this property. Notice that $\sigma(st(d_1 \dots d_c)) <_{\mathcal{J}} \sigma(d_1) \leq_{\mathcal{J}} J$: the first inequality is true due to axiom (**), since $\sigma(st(d_1 \dots d_c)) \neq \sigma(d_1)$, and the second because $\sigma(d_1)$ is the product of the letters in the leaf nodes below x, which are all in J. Consider the subtree of t rooted in x, in which we change the label of x into $st(d_1 \dots d_c)$. It is a factorization tree for the st function (recall that in descendants of x the functions st and st' return the same values) with height at most $3|M'|$, output outside $\sigma^{-1}(J)$, and its input is an infix of v. This contradicts with our assumption about v. ◀

The next lemma constructs a factorization tree for sets A consisting of multiple $\mathcal{J}$-classes, by composing factorization trees for single $\mathcal{J}$-classes obtained from the previous lemma.

▶ **Lemma B.3.** *Let $A \subseteq M'$ be such that when $s \in A$ and $r \geq_{\mathcal{J}} s$ then $r \in A$.*[3] *Let $v \in D^+$. Then there exist factorization trees $t_1, \dots, t_k$ with height at most $(3|M'|+2)|A|$, such that the concatenation of their inputs gives v, and either $k = 1$, or all these trees have output outside $\sigma^{-1}(A)$.*

Proof. The proof is by induction on the size of A. The base case is that A is empty. Then for each letter of v we construct a one-node tree with this letter as input. These trees are of height 0, and they have outputs outside $\sigma^{-1}(A)$.

Next, assume that A is nonempty. Let J be some $\leq_{\mathcal{J}}$-minimal $\mathcal{J}$-class in A; denote $A' = A \setminus J$. We apply the induction assumption for v and A'. We obtain factorization trees $t^0_1, \dots, t^0_m$ of height at most $(3|M'|+2)|A'|$, such that the concatenation of their inputs gives v; we either have $m = 1$, or each t^0_i has output outside $\sigma^{-1}(A')$. When $m = 1$, this already concludes the thesis of the lemma; below we assume that $m > 1$.

We apply Lemma B.2 to w and J. We obtain factorization trees $t^1_1, \dots, t^1_n$ with height at most $3|M'|$, such that the concatenation of their inputs gives w; whenever some t^1_i for $i \in \{1, \dots, n-1\}$ has output in $\sigma^{-1}(J)$, then t^1_{i+1} has output outside $\sigma^{-1}(J)$. Notice additionally that the projection of the output of a factorization tree is $\leq_{\mathcal{J}}$ than the projection of any letter in its input (we have $\sigma(pr(d_1, d_2)) = \sigma(d_1) \cdot \sigma(d_2) \leq_{\mathcal{J}} \sigma(d_i)$ and $\sigma(st(d_1 \dots d_k)) \leq_{\mathcal{J}} \sigma(st(d_1))$ by axioms (*) and (**)). Thus, since the letters of w are outside $\sigma^{-1}(A')$, also the output of each t^1_i is outside $\sigma^{-1}(A')$. So we can strengthen the statement above: whenever some t^1_i for $i \in \{1, \dots, n-1\}$ has output in $\sigma^{-1}(A)$, then t^1_{i+1} has output outside $\sigma^{-1}(A)$.

Next, in the place of the i-th leaf in the sequence of trees $t^1_1, \dots, t^1_n$ we substitute the tree t^0_i (notice that the label of this leaf and of the root of t^0_i is the same: it is the i-th letter of w). In this way we obtain factorization trees $t^2_1, \dots, t^2_n$ of height at most $(3|M'|+2)|A'| + 3|M'|$.

[3] That is, $M' \setminus A$ is an ideal.

The concatenation of their inputs gives v, and whenever some t_i^2 for $i \in \{1, \ldots, n-1\}$ has output in $\sigma^{-1}(A)$, then t_{i+1}^2 has output outside $\sigma^{-1}(A)$.

Finally, when some t_i^2 for $i \in \{1, \ldots, n-1\}$ has output in $\sigma^{-1}(A)$, we merge it with t_{i+1}^2 using a binary node. The output of this new tree is outside $\sigma^{-1}(A)$ (notice that $t \notin A$ implies $s \cdot t \notin A$, since $t \geq_{\mathcal{J}} s \cdot t$). Similarly, if the last tree has output in $\sigma^{-1}(A)$, we merge it with its predecessor (which is possibly already merged with its predecessor). After this merging we obtain factorization trees $t_1, \ldots, t_k$ with height at most $(3|M'|+2)|A'| + 3|M'| + 2 \leq (3|M'|+2)|A|$; the concatenation of their inputs is v. If we had $n > 1$, the output of each of these trees is outside $\sigma^{-1}(A)$ (however it is possible that $n = 1$ and the only tree has output in $\sigma^{-1}(A)$). ◀

Notice that this lemma for $A = M'$ implies immediately Theorem 7.5.

C Proof of Facts 7.3 and 7.6

Proof of Fact 7.3. Recall that we want to divide an arbitrary word w over a finite monoid M' into fragments $w = w_1 \ldots w_k$ for $k \leq \exp(3|M'|)$ such that for each i either $|w_i| = 1$, or $\pi(w_i)$ is idempotent. We apply the standard factorization tree theorem (Theorem B.1, where $D = M'$ and $st(e \ldots e) = e$) to w: we obtain a factorization tree with input w. In this tree we identify those leaves and idempotent nodes which do not have idempotent nodes as ancestors. They give a division of w into fragments $w = w_1 \ldots w_k$. The fragments corresponding to leaves have length 1; the fragments corresponding to idempotent nodes evaluate to idempotents. Notice that above the considered nodes there are only binary nodes, and the tree has height at most $3|M'|$, so there are at most $\exp(3|M'|)$ fragments. ◀

Proof of Fact 7.6. Because $A = E \cdot A \cdot E$, we have $A \leq_{\mathcal{J}} E$. If $A <_{\mathcal{J}} E$ we are done, so assume that $A \sim_{\mathcal{J}} E$. Because E is idempotent, we have $A = E \cdot A \cdot E = E \cdot E \cdot A \cdot E = E \cdot A$, and similarly $A = A \cdot E$.

We have to define more relations. For elements r, s of a monoid, we write $r \sim_{\mathcal{R}} s$ when there exist u_1, u_2 such that $r = s \cdot u_1$ and $s = r \cdot u_2$. Symmetrically, we write $r \sim_{\mathcal{L}} s$ when there exist u_1, u_2 such that $r = u_1 \cdot s$ and $s = u_2 \cdot r$. We also define $r \sim_{\mathcal{H}} s$ when $r \sim_{\mathcal{R}} s$ and $r \sim_{\mathcal{L}} s$. Lemma 3.5 of [18] says that $r \sim_{\mathcal{J}} r \cdot s$ implies $r \sim_{\mathcal{R}} r \cdot s$; symmetrically, $r \sim_{\mathcal{J}} s \cdot r$ implies $r \sim_{\mathcal{L}} s \cdot r$. Moreover, Lemma 3.8 of [18] says that if H is an $\mathcal{H}$-class such that for some $r, s \in H$ we have $r \cdot s \in H$, then H is a group.

We apply the above facts to our case. Since $E \sim_{\mathcal{J}} A = E \cdot A$, we have $E \sim_{\mathcal{R}} A$, and since $E \sim_{\mathcal{J}} A = A \cdot E$, we have $E \sim_{\mathcal{L}} A$; thus $E \sim_{\mathcal{H}} A$. Because $A = E \cdot A$, the $\mathcal{H}$-class of E and A is a group. Notice that E is the neutral element of the group (the neutral element is the only idempotent in a group). Since the group is finite, for some $k > 1$ we have $A^k = E$. Because $E \subseteq A$, we have $A = A \cdot E^{k-1} \subseteq A \cdot A^{k-1} = E$, so $E = A$. ◀