Satisfiability is Decidable for a Fragment of AMSO Logic on Infinite Words

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

1 **Department of Mathematics, Technische Universität Darmstadt, Germany**
`blumensath@mathematik.tu-darmstadt.de`

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 $s < f(x) \leq r$.

1 Introduction

This paper continues a line of research trying to find logics that have decidable satisfiability over infinite words (and infinite trees). The most known such logic is the 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, introduced recently, asymptotic monadic second-order logic (AMSO) is to extend MSO by the ability of expressing boundedness properties of some sequences of numbers. In MSO+U this is realized by an additional quantifier U: a formula $\mathsf{U}X\varphi$ says that φ is satisfied for arbitrarily large finite sets X. AMSO does not have, at least built in, the ability to refer to the size of sets. Instead, it describes weighted structures (in particular weighted infinite words), which are structures in which elements are labeled by natural numbers called 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 under some 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 happens to be equivalent as far as decidability of satisfiability is concerned [1], and, unfortunately, this means that both are undecidable over infinite words [5]. Nevertheless, some fragments of these logics are be decidable.

Indeed, in [2] the satisfiability problem of MSO+U is solved over infinite trees for formulae where the U quantifier is at the outermost position. A significantly more powerful fragment of the logic, although over infinite words, was shown 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 decidable over infinite words [3] and infinite trees [6].

Notice 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

* Work supported by the fellowship of the Foundation for Polish Science, during the author's post-doc stay at Université Paris Diderot

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], no fragments are known to be decidable so far (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 equivalent (Theorem 5 in [1]). It is conjectured that these two fragments have decidable satisfiability 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 extremely simpler since it only inhabits the finite levels of the Borel hierarchy.

Let 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.

Contributions

In [1], the satisfiability problem for $\mathrm{AMSO}^{\mathrm{np}}$/WAMSO was reduced to a form of tiling systems. The main contribution of this paper is to solve a special case of this tiling problem. In consequence we can solve the satisfiability problem over infinite words 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 variables r and s can be only used simultaneously, in subformulae of the form $s < f(x) \leq r$. As a tool, 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 the MSO logic by the ability to describe asymptotic properties over quantities. It refers to *weighted structures*, that are pairs $\langle \mathfrak{A}, \overline{f} \rangle$ consisting of a relational structure $\mathfrak{A}$ and a tuple of functions $f_i\colon dom(\mathfrak{A}) \to \mathbb{N}$ (*weight functions*). We only consider the case when $\mathfrak{A}$ is an infinite word (ω-word). 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 is a first-order variable, and r is a number variable; such formulae are restricted to appear positively inside the existential quantifier (dually: negatively inside the universal quantifier) binding r.

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

▶ **Example 2.1.** The following are correct formulae of $\mathrm{AMSO}^{\mathrm{np}}_{2s}$:

- $\exists t \forall x\, (f(x) \leq t)$, saying that the weights are bounded,

- $\forall s \exists r \forall x \exists y\, (y > x \wedge s < f(y) \leq r)$, saying that infinitely many weights occur infinitely often in the weighted infinite word,
- the disjunction (or conjunction) of the above two (we can move the quantifiers before the disjunction).

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

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

is equivalent to $\exists t \forall s \exists r\, \varphi(r, \dots, r, s, \dots, s, t, \dots, t)$.[1] For this reason we only allow in $\mathrm{AMSO}^{\mathrm{np}}_{2s}$ 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.*

Commutative Lossy Tiling Problem

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

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

3 From Logic to Tilings

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 situation when there is exactly 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 $\psi' = \forall s \exists r\, \varphi'(r, s) \in \mathrm{AMSO}^{\mathrm{np}}_{2s}$ 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".

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

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 (using 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 as well all needed compositionality properties, and the proof can be repeated smoothly 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 $\forall s \exists r\, \varphi(r, s)$ is equivalent to $\forall s\, \varphi(s+1, s)$. This step is not affected.

In the third step, the limit satisfiability problem is reduced to the lossy tiling problem. First, we observe that, because of just one variable s quantified universally, the resulting tiling system is of dimension one. Then, we have to change slightly the resulting tiling system so that it becomes 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 becomes $K = \bigcup_{a\in\Sigma}(a, \neq)^*((a, =) \cup \varepsilon)(a, \neq)^*$, which is a commutative language.

4 Monoids

In this section we slightly rephrase the problem of deciding commutative lossy tiling problems using explicitly monoids. In our solution we use algebra, in particular monoids. 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 is more convenient to write in the picture directly elements of M ($\varphi(a)$ instead of a). The row language becomes $\pi^{-1}(F)$, where $\pi\colon M^* \to M$, called *evaluation*, 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 is some commutative lossy language.

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 $\rho = (wA^*)^{\downarrow\circlearrowleft}$, and denoted $gl(\rho)$. The *norm* of such ρ, denoted $\|\rho\|$, is defined as $|w|$.

It is a consequence of the Highman's lemma that every lossy language (over M) is a finite union of languages of the form $(A_0^* b_1 A_1^* \dots b_k A_k^*)^{\downarrow}$, 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 whose each column belongs to $\bigcup B$, and every row 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 to the evaluation of the i-th row of p, for each i. Then, instead of requesting that every row of p belongs to $\pi^{-1}(s)$, we can say that $\pi(p) \in s^*$.

5 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, such that each column of this picture belongs to $\bigcup B$. New base languages are added following two kind schemas, called the *product schemas* and *diagonal schemas*. 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, saying 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 generating new base languages: product schemas and diagonal schemas.

Let ρ_1, 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 is of 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 special rows, and A contains the letters of $\pi(q)$ corresponding to 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 powerset monoid. The set $\mathcal{P}(M)$ of subsets of M has a natural monoid structure: $A \cdot B = \{a \cdot b \mid a \in A, b \in B\}$. We say that a set of base languages B is *uniform*, when it is nonempty, and for all $\rho_1, \rho_2 \in B$ it holds $gl(\rho_1) = gl(\rho_2)$, and this set is idempotent. For a uniform B we denote $gl(B)$ for $gl(\rho)$ where $\rho \in B$. The set of all finite uniform sets of base languages over M is denoted by $UBL(M)$.

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

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. $x = z \cdot x \cdot z \cdot x \cdot y$).

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 $(wA^*)^{\downarrow\circlearrowleft}$, where w consists of the letters of $\pi(q)$ corresponding to special rows, and A contains the letters of $\pi(q)$ corresponding to 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 (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 $\eta\colon UBL(M) \to \mathbb{N}$ 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 $\eta(B)$.*

There is a computable function $\eta\colon UBL(M) \to \mathbb{N}$ 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, whose each column belongs to $\bigcup B_0$, and for which $\pi(p) \in s^*$, and*
- *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 gives 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, that is from right to left. This implication is based on the following two lemmas.

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

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

Using these lemmas we now conclude with the soundness implication of Theorem 5.1. Let $B_i^{\leq\eta}$ be the sets from Theorem 5.1. The function η bounding sizes of diagonal schemas does not matter in this implication. We will prove by induction on i that if $c \in \bigcup B_i^{\leq\eta}$, then there exists a picture p whose each column belongs to $\bigcup B_0$, and such that $\pi(p) = c$ (this concludes the proof: we take $c = s^h$; since $s \in gl(\rho)$ for some $\rho \in B_x^{\leq\eta}$, we have $c \in \bigcup B_x^{\leq\eta}$). This is immediate for $i = 0$: we can take p containing c as the only column. Take some $c \in \bigcup B_i^{\leq\eta}$ for $i > 0$. We have $c \in \rho$ for some $\rho \in B_i^{\leq\eta}$. If $\rho \in B_{i-1}^{\leq\eta}$ we are done. Otherwise we are in the second or the third case of definition of $B_i^{\leq\eta}$, and then we use Lemma 6.1 or 6.2. We obtain a picture p' whose each column belongs to $\bigcup B_{i-1}^{\leq\eta}$, and such that $\pi(p') = c$. Moreover, by induction assumption, for each column c_j of p' there exists a picture p_j whose each column belongs to $\bigcup B_0$ and such that $\pi(p_j) = c_j$. To obtain p, in p' we replace, for each j, the j-th column c_j by p_j. Notice that $\pi(p) = \pi(p')$, so p is as required.

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

Proof of Lemma 6.1. The proof is immediate. We start from a product schema q for ρ_1, ρ_2 which generates ρ. Since global rows of q contain only letters from the global parts of ρ_1, ρ_2, in q we can duplicate any global row, and still 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 $\pi(p) = c$. ◀

Proof of Lemma 6.2. Let $\rho = (wA^*)^{\downarrow\circlearrowleft}$, let q be a diagonal schema for B generating ρ, and let $q_1, \dots, q_k$ 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). Notice also that if the lemma holds for some word c, then it holds also for any c' obtained from c by removing and reordering letters (because we can remove and reorder rows of the resulting picture p). Thus it is enough to consider, for each $n \in \mathbb{N}$, a column c 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 $j \in \{1, \dots, k\}$ we create p_j by modifying q_j. In p_j we will have $|A| \cdot (n-1)$ more rows than in q_j; more precisely, each global row of q_j will evolve into n rows of p_j, and each special row of q_j will evolve into one row of p_j. Fix some j. Let m be the width of q_j. If $m = 1$, we just replace each global row by its n copies. Assume now that $m > 1$. Then the width of p_j will be nm. Consider a special row v. One case is that $\pi(v)$ is idempotent. Then we just repeat the content of the row n times. After the repetition the value remains the same. Otherwise, by definition there exists an index i such that the i-th letter of v belongs to $gl(B)$. Then, as the first $i-1$ letters of the new row we take the first $i-1$ letters of v. Also 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 (it is possible since $gl(B)$ is idempotent thanks to uniformity of B). Again, the value of the row remains unchanged. Finally, consider a global row v of q_j. Out of it we create 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 $gl(B)$). Also 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.

As 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 c (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), a column 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 opposite direction of Theorem 5.1, that is from left to right. 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, out of which one can extract product schemas. In Section 7.2 we analyze pictures whose columns come from a union of a uniform set of base languages; they can be turned into diagonal schemas. Next, in Section 7.3 we introduce a tool: 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 to finish the proof. For the scope of the whole section we assume that the monoid M is fixed.

7.1 Products

We start by analyzing pictures of width 2, and we say that they can be turned into product schemas.

▶ **Lemma 7.1.** *Let ρ_1, ρ_2 be two base languages. Let p be a picture of width 2 such that the first column belongs to ρ_1 and the second 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 $\pi(p) \in \rho$. To q we take all rows of p which do not evaluate to an element of A. That will be special rows. Notice that in each of these rows either its first letter does not belong to $gl(\rho_1)$, or its second letter does not belong to $gl(\rho_2)$. Thus we have at most $\|\rho_1\| + \|\rho_2\|$ such rows. Moreover, for each $r \in gl(\rho_1)$ and each $s \in gl(\rho_2)$, to q we add a row having r in the first column, and s in the second column. That will be 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 say that then 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 whose each column 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*
- *for $E = gl(B)$ and $A = gl(\rho)$ it holds $E \subseteq A = E \cdot A \cdot E$.*

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, what will be useful later. Namely, the global part of ρ contains not only the letters that appear many times in $\pi(p)$, but also ($E \subseteq A$) all letters from $gl(B)$, and ($E \cdot A \cdot E \subseteq A$) all results of surrounding the former letters by letters from $gl(B)$. Notice that always $A \subseteq E \cdot A \cdot E$, since each global row begins and ends by a letter from $gl(B)$.

Below we prove the above lemma. We base on the following fact saying that each word can be chopped into a small number of idempotents and single letters. To eliminate towers of exponents, we write $p_2(x)$ for 2^x.

▶ **Fact 7.3.** *Let M' be a finite monoid, and let w be a word over M'. Then we can divide w into fragments $w = w_1 \dots w_k$ for $k \leq p_2(3|M'|)$ such that for each 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 let $E \subseteq M$. 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 $k \leq p_2(3(y - x + 1)|M|^y)$, 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. This is induction on $y - x$ (notice that always $x \leq y$). 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 its letter contains the elements of M appearing in the x rows of a column). We apply Fact 7.3 to w. It gives us a division $w = w_1 \dots w_m$ for $m \leq p_2(3|M|^x) \leq p_2(3|M|^y)$ such that 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 each row of each p'_j which is in I either has length 1, or evaluates to an idempotent. Next, for each p'_j we proceed in one of two ways.

- The first case is that each row of p'_j which is not in I contains a letter from E. Then this p'_j satisfies the thesis of the lemma.
- 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 induction assumption for p'_j; it gives us a subdivision of p'_j as required by the statement of the lemma.

Since each of the subdivisions returns at most $p_2(3(y' - x' + 1)|M|^{y'}) \leq p_2(3(y - x)|M|^y)$ pictures, in total we have at most $m \cdot p_2(3(y - x)|M|^y) \leq p_2(3(y - x + 1)|M|^y)$ pictures. ◀

Proof of Lemma 7.2. Denote $E = gl(B)$. First, we apply Lemma 7.4 to the picture p and to the set E. It divides p into some pictures $p_1, \dots, p_k$. Notice that the number y in the statement of the lemma is equal to the maximal norm of a base language in B, and $x \geq 0$;

we have $k \le p_2(3(y-x+1)|M|^y) \le p_2(3(y+1)|M|^y)$. We identify a set I_1 of numbers of rows of p: we have $i \in I_1$ when the first or the last letter of the i-th row of some p_j is in $M \setminus E$. Notice 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 rows from I_1 are treated as special rows). However we still need to reduce its size, and ensure the condition $E \subseteq A = E \cdot A \cdot E$.

For each i, by s_i we denote 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). Denote $I = I_1 \cup I_2$.

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 $\rho = (wA^*)^{\downarrow\circlearrowleft}$. We easily see that $\pi(p) \in \rho$ and $E \subseteq A = E \cdot A \cdot E$, because E is idempotent. 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 $p_1, \dots, p_k$. To q we take all those rows of p which do not evaluate to an element of A. That will be special rows. Notice that by the thesis 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 r, r' are the first and the last letter of the row, that are in E by definition of I_1, and $s_i \in A'$. In consequence, 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: for each pair $r, r' \in E$ we take to q one such row, in which we replace the first letter by r, and the last letter by r'. That will be global rows. This is allowed: recall that the first and the last letter of each such row inside each p_j belongs to E, also the replaced letters are in E. Additionally, for each $s \in E$, we add to q a row containing only letters from E, which evaluates to s (for any length such row exists, because E is idempotent). That will be global rows as well. This is allowed, 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

$$|I| + |E| \cdot |A'| \cdot |E| + |E| \le 2ky + 2|M|^3 + |M| \le 2y \cdot p_2(3(y+1)|M|^y) + 3|M|^3,$$

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

We also have to restrict the width of q. Since we have started from any picture p, the width can be arbitrary; 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 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 $i+1, \ldots, l$, 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 removing as long as such pair of indices i, l exists. And, by pigeonhole principle, among any $|M|^h + 1$ numbers we can find two i, l for which $w_i = w_l$. Thus, after such 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 $k \cdot ((\theta(B)+1) \cdot (|M|^h+1)+1) \leq p_2(3(y+1)|M|^y) \cdot ((\theta(B)+1) \cdot (|M|^h+1)+1)$ columns. We denote the last number as $\eta(B)$. Notice that $\theta(B) \leq \eta(B)$, so 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 described in Sections 7.1 and 7.2.

The nodes of our factorization trees will be labeled by elements of any 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

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

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

- (*) for each $a, b \in D$ it holds $\sigma(pr(a,b)) = \sigma(a) \cdot \sigma(b)$, and
- (**) for each $d_1 \ldots d_c \in dom(st)$ it holds $\sigma(st(d_1 \ldots d_c)) = \sigma(d_1)$ or $\sigma(st(d_1 \ldots d_c)) <_{\mathcal{J}} \sigma(d_1)$.

In the second axiom above we use the $\leq_{\mathcal{J}}$ preorder, which is defined by $r \leq_{\mathcal{J}} s$ when there exist 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 labeled by elements of D, whose nodes are of one of three forms:

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

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*.

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

▶ **Theorem 7.5.** *Let $v \in D^+$. Then 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 stabilization monoid case ([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. Details are given in Appendix B.

7.4 Final Argument

In this subsection we conclude our proof of the left-to-right implication of Theorem 5.1. The function η in its statement 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, whose each column 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 $s \in \rho$.

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

Consider two letters (w_1, ρ_1) and (w_2, ρ_2) from D for which we want to define pr. Let p be the picture with two columns: w_1 and w_2. We fix some base language ρ such that $\pi(p) \in \rho$, and $gl(\rho) = gl(\rho_1) \cdot gl(\rho_2)$, and there exists a product schema for ρ_1, ρ_2 which generates ρ; it exists by Lemma 7.1. We return $pr((w_1,\rho_1),(w_2,\rho_2)) = (\pi(p), \rho)$. 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 $\rho \in B_{j+1}^{\leq\eta}$.

Consider now $(w_1,\rho_1)\dots(w_k,\rho_k) \in D^+$ such that $gl(\rho_1) = \dots = gl(\rho_k)$ is idempotent. Let p be the picture with k columns: the i-th column is w_i. Denote $B = \{\rho_1, \dots, \rho_k\}$; by definition it is a uniform set of base languages, and each column of p belongs to $\bigcup B$. Let $E = gl(B)$. We fix some base language ρ such that $\pi(p) \in \rho$, and $E \subseteq gl(\rho) = E \cdot gl(\rho) \cdot E$, and there exists a diagonal schema for B of width and height at most $\eta(B)$ which generates ρ; it exists by Lemma 7.2. We return $st((w_1,\rho_1)\dots(w_k,\rho_k)) = (\pi(p), \rho)$. 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 $E, A \subseteq M$. Assume that E is idempotent, and $E \subseteq A = E \cdot A \cdot E$. Then either $A = E$ or $A <_{\mathcal{J}} E$.*

Recall that p is a picture of height h, whose each column belongs to $\bigcup B_0$, and for which $\pi(p) \in s^*$. We want to find a base language $\rho \in B_x^{\leq\eta}$ for which $s \in gl(\rho)$. 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 such that $d_i \in \rho_i$. Consider a factorization tree t with height at most x and input w; it exists by Theorem 7.5. Denote its output as (d, ρ). Notice that $d = \pi(p) = s^h$ (by definition of the pr and st functions), 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 labeled by some (d', ρ'), then $\rho' \in B_i^{\leq\eta}$). Because h is by definition greater than the size of ρ, necessarily $s \in gl(\rho)$, which is what we wanted to prove.

---

[2] One can obtain a bound $3|M'|$, but it requires enhancing the proof.

References

1 A. Blumensath, O. Carton, and T. Colcombet. Asymptotic monadic second-order logic. In E. Csuhaj-Varjú, M. Dietzfelbinger, and Z. 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 M. Bojańczyk. The finite graph problem for two-way alternating automata. *Theor. Comput. Sci.*, 3(298):511–528, 2003.

3 M. Bojańczyk. Weak MSO with the unbounding quantifier. *Theory Comput. Syst.*, 48(3):554–576, 2011.

4 M. Bojańczyk and T. 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 M. Bojańczyk, P. Parys, and S. Toruńczyk. The MSO+U theory of (N, <) is undecidable. *CoRR*, abs/1502.04578, 2015. Submitted to STACS 2016.

6 M. Bojańczyk and S. Toruńczyk. Weak MSO+U over infinite trees. In C. Dürr and T. 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 T. Brázdil, K. Chatterjee, A. Kucera, and P. Novotný. Efficient controller synthesis for consumption games with multiple resource types. In P. Madhusudan and S. 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. R. 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 C. Carapelle, A. Kartzow, and M. Lohrey. Satisfiability of CTL* with constraints. In P. R. D'Argenio and H. 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 T. Colcombet. Factorisation forests for infinite words. In E. Csuhaj-Varjú and Z. Ésik, editors, *FCT*, volume 4639 of *Lecture Notes in Computer Science*, pages 226–237. Springer, 2007.

11 T. Colcombet. The theory of stabilisation monoids and regular cost functions. In S. Albers, A. Marchetti-Spaccamela, Y. Matias, S. E. Nikoletseas, and W. 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 T. Colcombet. Regular cost functions, part I: logic and algebra over words. *Logical Methods in Computer Science*, 9(3), 2013.

13 N. Fijalkow and M. Zimmermann. Cost-parity and cost-streett games. *Logical Methods in Computer Science*, 10(2), 2014.

14 S. Hummel and M. Skrzypczak. The topological complexity of MSO+U and related automata models. *Fundam. Inform.*, 119(1):87–111, 2012.

15 D. Kirsten. Distance desert automata and the star height problem. *ITA*, 39(3):455–509, 2005.

16 O. Kupferman, N. Piterman, and M. Y. Vardi. From liveness to promptness. *Formal Methods in System Design*, 34(2):83–103, 2009.

17 I. Simon. Factorization forests of finite height. *Theor. Comput. Sci.*, 72(1):65–94, 1990.
18 S. 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 $(3|M'|+2)|A'|+3|M'|$. The concatenation of their inputs gives v, and whenever some t^2_i for $i \in \{1, \dots, n-1\}$ has output in $\sigma^{-1}(A)$, then t^2_{i+1} has output outside $\sigma^{-1}(A)$.

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

Finally, when some t_i^2 for $i \in \{1, \dots, n-1\}$ 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 $t_1, \dots, t_k$ 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 $w = w_1 \dots w_k$ for $k \leq p_2(3|M'|)$ 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 $st(e \dots e) = e$) 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 $w = w_1 \dots w_k$. 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 $p_2(3|M'|)$ 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 $E = A$. ◀