TheoretiCS

Orbit-Finite-Dimensional Vector Spaces and Weighted Register Automata

Mikołaj Bojańczyk [a] ✉ iD

Joanna Fijalkow [a,b] ✉ iD

Bartek Klin [c] ✉ iD

Joshua Moerman [d] ✉ iD

a University of Warsaw, Poland

b University of Bordeaux, CNRS, LaBRI, Talence, France

c University of Oxford, UK

d Open Universiteit, Heerlen, The Netherlands

ABSTRACT. We develop a theory of vector spaces spanned by orbit-finite sets. Using this theory, we give a decision procedure for equivalence of weighted register automata, which are the common generalization of weighted automata and register automata for infinite alphabets. The algorithm runs in exponential time, and in polynomial time for a fixed number of registers. As a special case, we can decide, with the same complexity, language equivalence for unambiguous register automata, which improves previous results in three ways: (a) we allow for order comparisons on atoms, and not just equality; (b) the complexity is exponentially better; and (c) we allow automata with guessing.

1. Introduction

Weighted automata over a field were introduced in [20] by Schützenberger. Such an automaton is defined in the same way as a nondeterministic automaton, with a set Q of states and an input alphabet Σ, except that instead of having subsets for transitions, initial and final states, the

automaton has weight functions into the underlying field:

$$\underbrace{I \colon Q \to \mathbb{F}}_{\text{initial}} \quad \underbrace{\delta \colon Q \times \Sigma \times Q \to \mathbb{F}}_{\text{transition}} \quad \underbrace{F \colon Q \to \mathbb{F}}_{\text{final}}.$$

The weight of a run is obtained by multiplying the weights along all transitions, the initial weight of the first state, and the final weight of the last state. The automaton recognizes a weighted language, which is the function $L \colon \Sigma^* \to \mathbb{F}$ that maps a word to the sum of weights of all runs on that word.

Schützenberger proved that weighted automata can be minimized [20, Sec. B], which provides a polynomial time algorithm for equivalence. In contrast, it is undecidable whether some field element $a \in \mathbb{F}$ is achieved as the weight of some word, already in the special case of weighted automata that are probabilistic [17, Thm. 22]. Equivalence is undecidable for weighted automata over semirings that are not fields, e.g. for the min-plus semiring [14, Cor. 4.3].

One application of weighted automata (see [1] for other ones) is a polynomial time algorithm for language equivalence of unambiguous automata, i.e. nondeterministic finite automata that have at most one accepting run for every word. The algorithm is a simple reduction to equivalence of weighted automata: a nondeterministic automaton can be viewed as a weighted automaton over the field of rational numbers, such that the weight of a word is the number of accepting runs. For unambiguous automata, the number of accepting runs is either zero or one, and hence two unambiguous automata accept the same words if and only if the corresponding weighted automata are equivalent.

In this paper, we generalize weighted automata to infinite alphabets, motivated by the study of register automata, in particular the equivalence problem for unambiguous register automata [2, 9, 15, 16]. The kinds of infinite alphabets that we study are constructed using an infinite set $\mathbb{A}$ of atoms (also called data values) which can only be accessed in very limited ways; in the simplest case, they can only be compared for equality. Register automata are like finite automata, except that they additionally use finitely many registers to store atoms that occurred in the word. They were introduced by Kaminski and Francez [12, Def. 1], under the name of *finite memory automata*. This model has attracted much attention, and register automata are now one of the most widely studied infinite state systems. Their decidability landscape is rather complex: for example, emptiness is decidable for nondeterministic register automata [12, Thm. 1], but universality is not [18, Thm 5.1]. More robust results can be achieved for deterministic register automata, but at a considerable loss of expressive power.

One of the numerous (unfortunately non-equivalent) variants of register automata are unambiguous register automata, for which language equivalence was shown in [16, Thm. 1] to be decidable in 2ExpSpace, and in ExpSpace for a fixed number of registers. These upper bounds were improved to 2ExpTime and ExpTime, respectively, in [2, Thm. 2]. None of these proofs use weighted automata.

From the point of view of register automata, the main contributions of this paper are:

1. We introduce a weighted version of register automata, and we prove that their equivalence problem can be solved in ExpTime, and in polynomial time for a fixed number of registers.

2. We show that our weighted automata have a robust theory, in particular they can be described in several different ways and admit canonical syntactic automata.

3. As an immediate application of our equivalence algorithm for weighted automata, we show that the language equivalence problem for unambiguous register automata can be solved in the same complexity, improving exponentially on prior work [2, 16].

4. In [2, 16], the equivalence algorithms work only for automata which are non-guessing, in the sense that every pair (state, input letter) has finitely many outgoing transitions. Without the non-guessing assumption, decidability is only known in the case where one of the automata has just a single register [15, Thm. 10]. Our algorithm can be adapted so that it works for general unambiguous automata, without the non-guessing assumption, with the same complexity.

5. Apart from atoms with equality only, our algorithm for weighted automata, and its applications to unambiguous automata, also work for atoms equipped with a total order. Previous algorithms for ordered atoms assume that one of the automata has a single register [16, Thm. 2].

In our opinion, perhaps the most interesting aspect of this paper is not the applications described above, but the theory that is developed in order to obtain them: elements of linear algebra in the universe of sets with atoms [5], also known as nominal sets [19]. Our central objects of study are vector spaces which are spanned by orbit-finite set of vectors. A typical example is the vector space $\mathsf{Lin}(\mathbb{A}^k)$, which consists of finite linear combinations of k-tuples of atoms. Such a vector space has two kinds of structure: we can take linear combinations of vectors, and we can apply atom automorphisms to them. A natural concept – which appears in our equivalence algorithm for weighted automata – is the *equivariant subspaces*, i.e. subsets of the vector space closed under both linear combinations and atom automorphisms. Our principal technical result is that for every k, there is a finite (exponential in k) bound on the maximal length of chains

$$\underbrace{V_1 \subsetneq V_2 \subsetneq \cdots \subsetneq V_n}_{\text{equivariant subspaces}} \subseteq \mathsf{Lin}(\mathbb{A}^k).$$

This bound is enough to obtain an algorithm for checking equivalence on weighted automata. For other problems, such as automata minimization, we need to go beyond vectors spaces of the form $\mathsf{Lin}(\mathbb{A}^k)$ and consider the more general class of spaces which are spanned by orbit-finite sets. Since the axiom of choice fails for sets with atoms (see e.g. [19, Sec. 2.7]), these orbit-finitely spanned spaces may not admit equivariant bases. On the other hand, they are closed under

subspaces, Cartesian products, tensor products and dual spaces. We hope that these results can be used in future work, beyond applications to unambiguous automata.

The structure of the paper is as follows. In Section 2 we recall the basics of sets with atoms and orbit-finite sets, with particular focus on equality and ordered atoms. In Section 3 we introduce weighted orbit-finite automata and state our main result about checking their equivalence, Theorem 3.5. In Section 4 we leave automata aside, and we study vector spaces of the form Lin Q, for orbit-finite sets Q. We prove, in Theorem 4.8 and Corollary 4.9, that these spaces have finite length for equality and ordered atoms. These results are used in Section 5 to provide an algorithm for checking equivalence, thereby proving Theorem 3.5. In Section 6 we study the general notion of orbit-finitely spanned vector spaces and prove their numerous closure properties. This is used in Section 7 to introduce the class of orbit-finitely spanned automata, which have the same expressive power as weighted orbit-finite automata but admit minimization. As another application, in Section 8 we develop a new algorithm for equivalence checking for unambiguous register automata.

This is an extended version of the conference paper [7]. In addition to full proofs of several results, including a new and vastly simplified proof of Theorem 4.8, it settles several questions left open in [7]. In particular, we give an exponential lower bound on length (Section 4.3), an example of an atom structure where the finite length property fails (Section 4.4), and a proof that orbit-finitely spanned spaces over graph atoms are not closed under duals (Example 6.9 in Section 6).

2. Orbit-finite sets

Our paper is based on the approach to register automata that uses *sets with atoms*. The idea is to consider automata where all components are closed under atom automorphisms, and the states and input alphabets have finitely many elements up to atom automorphisms. This abstract view avoids cumbersome notation for representing states and transitions of register automata. We can also leverage existing results that treat atoms with more structure than equality only, e.g., with a total order.

Our notation is based on [5]. Fix a countably infinite relational structure $\mathbb{A}$, in the sense of model theory, i.e. an underlying set equipped with some relations. Elements of this fixed structure will be called *atoms*. In this paper we mostly focus on two such structures:

— the *equality atoms* $\mathbb{A} = (\mathbb{N}, =)$, and

— the *ordered atoms* $\mathbb{A} = (\mathbb{Q}, <)$.

The main results of this paper, notably the finite length property in Section 4, are proved for these atoms only, and we leave other atoms as future work.

An *atom automorphism* is any bijection of the underlying set of atoms which is consistent with all relations in $\mathbb{A}$. For the equality atoms this is any permutation of $\mathbb{N}$; for the ordered

atoms, any order-preserving bijection on $\mathbb{Q}$. In both settings, the set of all atoms is orbit-finite, which means that it has finitely many elements up to atom automorphisms. Other orbit-finite sets constructed with atoms include:

$$\underbrace{\mathbb{A}^k}_{k\text{-tuples of atoms}} \qquad \underbrace{\mathbb{A}^{(k)}}_{\substack{\text{non-repeating}\\ k\text{-tuples of atoms}}} \qquad \underbrace{\binom{\mathbb{A}}{k}}_{\text{sets of }k\text{ atoms}} . \tag{1}$$

To formally define "sets constructed with atoms" and the orbit-finite restriction, we use the cumulative hierarchy from set theory. (See [5, Sec. 3.1] or [19, Sec. 2.6] for a more detailed presentation.) The *cumulative hierarchy (over atoms $\mathbb{A}$)* is indexed by ordinal numbers, and defined as follows: on level 0 we find the atoms, and on the level indexed by an ordinal $\alpha > 0$ we find the atoms and all sets whose elements are from levels $< \alpha$. For example, on level 1 we find every subset of the atoms. Automorphisms $\pi : \mathbb{A} \to \mathbb{A}$ of the atoms act on sets in the cumulative hierarchy in the expected way.

In the cumulative hierarchy one can encode data structures such as tuples, words, relations, functions etc., using standard set-theoretic machinery. Automorphisms π then act on tuples as expected: $\pi((x_1, \ldots, x_n)) = (\pi(x_1), \ldots, \pi(x_n))$.

For a tuple of atoms $\bar{a} \in \mathbb{A}^*$, an *$\bar{a}$-automorphism* is an atom automorphism π such that $\pi(\bar{a}) = \bar{a}$. A tuple $\bar{a}$ is said to *support* an element x of the cumulative hierarchy if $\pi(x) = x$ for every $\bar{a}$-automorphism π. We say that an element x of the cumulative hierarchy is *finitely supported* if it is supported by some tuple of atoms in $\mathbb{A}^*$. An element is called *equivariant* if it is supported by the empty tuple. An equivariant set may contain non-equivariant elements; e.g., the set of all atoms is equivariant, but its elements are not. A *set with atoms* is defined to be a set in the cumulative hierarchy which is finitely supported and with each of its members (and their members, and so on) finitely supported.

Sets with atoms provide a relaxed notion of finiteness, called orbit-finiteness. We say that two sets with atoms x and y are in the same *$\bar{a}$-orbit* if there is some $\bar{a}$-automorphism π such that $y = \pi(x)$. Being in the same $\bar{a}$-orbit is clearly an equivalence relation; its equivalence classes are called *$\bar{a}$-orbits*. A set with atoms x is called *orbit-finite* if there is some atom tuple $\bar{a}$ such that x is a finite union of $\bar{a}$-orbits.

EXAMPLE 2.1. — Over equality atoms, the set $\mathbb{A}$ of atoms is orbit-finite, also for every $k \in \{1, 2, \ldots\}$ the sets $\mathbb{A}^{(k)}$ and $\binom{\mathbb{A}}{k}$ from (1) are orbit-finite, with a single equivariant orbit each. The set $\mathbb{A}^k$ is also orbit-finite, with the number of equivariant orbits equal to the k-th Bell number, i.e., the number of equivalence relations on the set $\{1, \ldots, k\}$.

— Over ordered atoms, the sets $\mathbb{A}$ and $\binom{\mathbb{A}}{k}$ have single equivariant orbit each. The set $\mathbb{A}^{(k)}$ is also orbit-finite, with $k!$ equivariant orbits; one of these orbits is the increasing k-tuples of atoms. Each of these orbits is in equivariant bijection with $\binom{\mathbb{A}}{k}$.

— Over ordered atoms, for every atom $a \in \mathbb{A}$ the set $\mathbb{A} - \{a\}$ has two a-orbits: the open intervals $(-\infty; a)$ and $(a; \infty)$.

— The set $\mathbb{A}^*$ is not orbit-finite, neither for equality nor for ordered atoms, as words of different length are necessarily in distinct orbits.

◆

An atom structure $\mathbb{A}$ is called *oligomorphic* if for every k, the set $\mathbb{A}^k$ is orbit-finite. For oligomorphic atoms, orbit-finite sets behave well, e.g. they are closed under finite products and finitely supported subsets [5, Lem. 3.24]. In this paper we only consider oligomorphic atoms; in particular, equality and ordered atoms are oligomorphic.

Orbit-finite sets can be represented in a finite way so that they can be used as inputs for algorithms. Two examples of such representations are set builder expressions [5, Sec. 4.1] or the G-set representation [6, Sec. 8]. The reader does not need to know these representations in detail; the important thing is that they support basic operations such as products, Boolean operations, or inclusion and membership checks.

3. Weighted orbit-finite automata

We now introduce the main model for this paper, an orbit-finite generalization of weighted automata.

DEFINITION 3.1 (Weighted orbit-finite automaton). Fix an oligomorphic atom structure $\mathbb{A}$ and a field $\mathbb{F}$. A *weighted orbit-finite automaton* consists of orbit-finite sets Q and Σ, called the *states* and the *alphabet*, and finitely supported functions

$$\underbrace{I : Q \to \mathbb{F}}_{\text{initial}} \quad \underbrace{\delta : Q \times \Sigma \times Q \to \mathbb{F}}_{\text{transitions}} \quad \underbrace{F : Q \to \mathbb{F}}_{\text{final}}.$$

Furthermore, we always require the *non-guessing condition:*

(*) there are finitely many states with nonzero initial weight, and also for every state q and input letter a, there are finitely many states p such that the transition (q, a, p) has nonzero weight.

Weights of runs and words is defined in the same way as for the classical model of weighted automata with finitely many states. The non-guessing condition ensures that each word has only finitely many runs with non-zero weight; otherwise there could be difficulties in summing up the weights of infinitely many runs. Intuitively, the condition means that an automaton cannot spontaneously invent a fresh atom.

EXAMPLE 3.2. Consider any oligomorphic atoms $\mathbb{A}$ and the field of rational numbers. We define a weighted orbit-finite automaton which maps a word $w \in \mathbb{A}^*$ to the number of distinct

atoms that appear in w. The states are

$$\underbrace{\{\bot\}}_{\substack{\text{initial weight 1} \\ \text{final weight 0}}} \quad + \quad \underbrace{\mathbb{A}}_{\substack{\text{initial weight 0} \\ \text{final weight 1}}}$$

for some $\bot \notin \mathbb{A}$. The transition weight is 1 for all triples

$$(\bot, a, \bot) \ (\bot, a, a) \ (a, b, a) \qquad \text{for } a \neq b \in \mathbb{A},$$

and 0 for all other triples. For every input word, all runs have weight 0, except for the following runs, which have weight 1: start in $\bot$, stay there until the last occurrence of some atom a, and then stay in state a until the end of the word. Since the number of runs with weight 1 is the number of distinct atoms a that appear in the word, the output of the automaton is the number of distinct atoms.

This automaton is equivariant, in the sense that its state space and all three weight functions are equivariant. ◆

EXAMPLE 3.3. The automaton in the previous example is a special case of a more general construction: counting accepting runs of a nondeterministic automaton. Define a *nondeterministic orbit-finite automaton* like a nondeterministic finite automaton, except that all components (the alphabet, states, transitions, initial and accepting sets) are required to be orbit-finite sets, see [5, Def. 5.7]. To such an automaton $\mathcal{A}$ one can associate a weighted orbit-finite automaton (over the field of rational numbers) as follows: the alphabet and states are the same, and:

— a state gets an initial weight 1 if it was initial in $\mathcal{A}$, otherwise its initial weight is 0;

— a state gets an final weight 1 if it was accepting in $\mathcal{A}$, otherwise its initial weight is 0;

— transition weight is 1 for all transitions present in $\mathcal{A}$ and 0 for the ones absent in $\mathcal{A}$.

This weighted automaton maps a word to the number of accepting runs of $\mathcal{A}$. The construction makes sense only if $\mathcal{A}$ is non-guessing in the sense that there are finitely many inputs states, and for every pair (state, letter) there are finitely many outgoing transitions. ◆

As mentioned in Section 2, orbit-finite sets can be represented in a finite way. Therefore, weighted orbit-finite automata can also be represented in a finite way (assuming that field elements can be represented in a finite way). This is because (a) a finitely supported relation on an orbit-finite set is also orbit-finite; and (b) a finitely supported function from an orbit-finite set to any set is itself an orbit-finite set.

EXAMPLE 3.4. A special case of weighted orbit-finite automata, where the finite representation is easier to see, is a weighted k-register automaton. This is a weighted orbit-finite automaton where the input alphabet is finitely many disjoint copies of the atoms, the states are finitely many copies of $(\mathbb{A} + \{\bot\})^k$, and all weight functions are equivariant. For equality and ordered atoms, the weight functions can be finitely represented using quantifier-free formulas, see [5,

p. 6]; as a result, weighted register automata appear as a straightforward weighted version of register automata as introduced in [12]. We will focus on this special case in Section 8.   ◆

We now state the main result of this paper, which is an algorithm for checking the equivalence of weighted orbit-finite automata, assuming that the atom structure is either the equality atoms or the ordered atoms. We do not know if the problem is decidable for other atom structures.

THEOREM 3.5. *Assume that the atoms are the equality atoms* $(\mathbb{N}, =)$ *or the ordered atoms* $(\mathbb{Q}, <)$. *The equivalence problem for equivariant[1] weighted orbit-finite automata can be solved in deterministic time*

$$2^{\mathrm{poly}(k)} \cdot n^{O(k)}$$

where

   n is the orbit count, i.e. the number of equivariant orbits in the disjoint union of the two state sets;

   k is the atom dimension of the state spaces of the automata, i.e. the smallest k such that every state in both automata is supported by at most k atoms.

In particular, the equivalence problem is in EXPTIME*, and polynomial time when the atom dimension k is fixed.*

A lower bound for the problem is PSPACE, which is the complexity of language equivalence of deterministic register automata. We do not know the exact complexity; it is worth pointing out that for weighted finite automata there is an equivalence algorithm in randomized polylogarithmic parallel time [13, Sec. 3.2] that uses the Isolating Lemma.

We now present a proof strategy for the theorem, which will be carried out in the next sections. The first observation is that the equivalence problem reduces to the zeroness problem, which asks whether a single weighted orbit-finite automaton outputs zero for every word. The reduction is as follows: given two equivariant weighted orbit-finite automata $\mathcal{A}_1$ and $\mathcal{A}_2$, we create a new weighted orbit-finite automaton $\mathcal{A}_1 - \mathcal{A}_2$, which is obtained by taking the disjoint union of $\mathcal{A}_1$ and $\mathcal{A}_2$, and flipping the sign of the final weights in $\mathcal{A}_2$. The new automaton maps all words to zero if and only if the original two automata were equivalent. Also, in the reduction, the atom dimension does not change, neither does the orbit count.

It remains to give an algorithm for zeroness of a single equivariant weighted orbit-finite automaton, with states Q. Our proof follows the same lines as Schützenberger's algorithm. Write Lin Q for the set of finite linear combinations (over the field $\mathbb{F}$) of states, seen as a vector space with basis Q. For an input word w, define its *configuration* to be the vector in Lin Q which

---

1   This theorem would also work for finitely supported automata, but the notation for the complexity and orbit counts would be more involved.

maps a state q to the sum of pre-weights of all runs over w that end in state q; the pre-weight of a run is defined by multiplying the initial weight of the first state and the weights of all transitions, without taking into account the final weights. Thanks to the non-guessing condition (*) in Definition 3.1, each configuration is indeed a finite linear combination, since there are finitely many runs with nonzero pre-weight. Consider the chain

$$V_0 \subseteq V_1 \subseteq V_2 \subseteq \cdots \subseteq \mathsf{Lin}\, Q$$

where V_i is the subspace of $\mathsf{Lin}\, Q$ that is spanned by the configurations of words of length at most i. (We do not intend to compute the subspaces in this chain, whatever that would mean; the chain is only used in the analysis of the algorithm.) Because the automaton is equivariant, one can easily see that each subspace V_n is also equivariant.

In the finite-dimensional case studied by Schützenberger, where Q was a finite set, we could conclude that the chain must stabilize in a number of steps that is bounded by the finite dimension of the vector space $\mathsf{Lin}\, Q$. In the orbit-finite case, the vector space has infinite dimension, and thus it is not clear why the chain should stabilize in finitely many steps. Our main technical contribution is a proof that the chain does indeed stabilize, and furthermore the time to stabilize is consistent with the bounds in the statement of Theorem 3.5. This stabilization property is the subject of the next section, and in Section 5 we shall use the property to conclude the proof of Theorem 3.5.

4.   Finite length property

A standard result in linear algebra says that a vector space has finite dimension if and only if it has finite length, where the length is defined to be the longest chain of its subspaces. Indeed, the length of a finite-dimensional space is equal to its dimension. Because of this easy correspondence the notion of length is seldom explicitly applied to vector spaces, and it becomes more important only in more general structures such as modules over a ring. However, the situation becomes more interesting for chains of *equivariant* subspaces. We define:

DEFINITION 4.1 (Length). The *length* of an equivariant vector space V, denoted $\mathsf{length}(V)$, is the maximal length n of a chain of proper inclusions on equivariant subspaces of V:

$$V_1 \subsetneq V_2 \subsetneq \cdots \subsetneq V_n \subsetneq V.$$

If a maximal length does not exist, we say that V has infinite length. An atom structure $\mathbb{A}$ has the *finite length property* if for every number k, the vector space $\mathsf{Lin}\, \mathbb{A}^k$ has finite length.

The finite length property easily implies oligomorphicity. As we shall see in Section 4.4, the converse implication does not hold. But first, let us consider the two basic atom structures.

4.1 Equality and ordered atoms

The purpose of this section is to prove that both the equality atoms $(\mathbb{N}, =)$ and the ordered atoms $(\mathbb{Q}, <)$ have the finite length property (Corollary 4.9 and Theorem 4.8, respectively). Furthermore, the length of $\mathsf{Lin}\,\mathbb{A}^k$ grows exponentially (and not worse) with k. The results of this section apply to an arbitrary field $\mathbb{F}$.

Definition 4.1 speaks of equivariant subspaces. Later, in Theorem 4.10, we will show that for the equality and ordered atoms, allowing a fixed non-empty support would make no difference.

EXAMPLE 4.2. Consider the equality atoms. The vector space $\mathsf{Lin}\,\mathbb{A}$ has length 2, because it has exactly 3 equivariant subspaces which form a chain of two proper inclusions:

$$\{0\} \quad \subsetneq \quad V \quad \subsetneq \quad \mathsf{Lin}\,\mathbb{A},$$

where V is the subspace spanned by the set

$$\{a - b : a \neq b \in \mathbb{A}\}.$$

Equivalently, V is the vector space of all vectors where all coefficients sum up to 0. Let us prove that there are no other equivariant subspaces. Suppose that W is an equivariant subspace which contains some nonzero vector

$$w = \lambda_1 a_1 + \cdots + \lambda_n a_n \quad \text{where } \lambda_i \in \mathbb{F} \setminus \{0\}. \tag{2}$$

Since W is equivariant, it also contains the vector obtained from w by replacing a_n with some fresh atom b_n. By taking the difference of these vectors and dividing by λ_n, we see that W contains $a_n - b_n$. By equivariance, W contains all vectors of the form $a - b$ for distinct atoms a, b, and thus $V \subseteq W$.

We now show that W is either V or the entire space $\mathsf{Lin}\,\mathbb{A}$. Indeed, suppose that W contains some w as in (2) that is not in V. If $n > 1$ then we can subtract from w the vector $\lambda_1 \cdot (a_1 - a_2) \in V$, which results in another vector that is in W but not in V, with a smaller n. By repeating this process, we see that W contains a vector of the form $\lambda_1 a_1$ with $\lambda_1 \neq 0$, and hence, by equivariance, it is the entire space. ◆

Definition 4.1 differs from the classical definition of length in that we only consider equivariant subspaces. Some classical properties of length easily transfer to our case. One may even say that our definition is a special case of the classical one: an equivariant vector space can be seen as a module over the (non-commutative) group ring $\mathbb{F}[G]$, where $\mathbb{F}$ is the underlying field and G is the automorphism group of $\mathbb{A}$. To keep the presentation elementary we do not pursue this correspondence, but we remark that all properties of length which are valid for modules over (non-commutative) rings, remain true for our definition. For example, the following lemma

has the same proof (see Appendix A) as for the classical notion of length of a module (see e.g. [11, Prop. 4.12]), and works for arbitrary oligomorphic atoms:

LEMMA 4.3. *For any equivariant spaces $V \subseteq W$, and equivariant sets P, Q with their disjoint union $P + Q$:*

(i) $\text{length}(W) = \text{length}(V) + \text{length}(W/V)$;[2]
(ii) $\text{length}(\text{Lin}(P + Q)) = \text{length}(\text{Lin}\,P) + \text{length}(\text{Lin}\,Q)$;
(iii) *if there is an equivariant surjective function from P to Q then* $\text{length}(\text{Lin}\,Q) \leq \text{length}(\text{Lin}\,P)$.

From this we infer that in Definition 4.1, we could have equivalently talked about arbitrary equivariant orbit-finite set, instead of just sets of the form $\mathbb{A}^k$:

COROLLARY 4.4. *For atoms with the finite length property,* $\text{Lin}\,Q$ *has finite length for every equivariant orbit-finite Q.*

PROOF. An orbit-finite set is a finite disjoint union of single-orbit sets, and for oligomorphic atoms every single-orbit set Q is an image of a surjective function from $\mathbb{A}^k$ [6], for k the atom dimension of Q. Hence, we can use the closure properties from Lemma 4.3. ∎

The following lemma is the key step in proving the finite length property for the equality and ordered atoms:

LEMMA 4.5. *Consider the ordered atoms $\mathbb{A} = (\mathbb{Q}, <)$. For every k, the length of $\text{Lin}\binom{\mathbb{A}}{k}$ is finite, and satisfies*

$$\text{length}\left(\text{Lin}\binom{\mathbb{A}}{k}\right) \quad \leq \quad 1 + k \cdot \text{length}\left(\text{Lin}\binom{\mathbb{A}}{k-1}\right).$$

PROOF. For any set α of $2k$ atoms:

$$\underbrace{a_1 < b_1 < \cdots < a_k < b_k}_{\alpha} \tag{3}$$

and for any $I \subseteq k = \{1, \dots, k\}$, define $\alpha \bowtie I \in \binom{\mathbb{A}}{k}$ by:

$$\alpha \bowtie I = \{a_i \mid i \notin I\} \cup \{b_i \mid i \in I\}.$$

In words, $\alpha \bowtie I$ picks either a_i or b_i from α according to I. Define the *cog* (on α), $v^\alpha \in \text{Lin}\binom{\mathbb{A}}{k}$, to be the vector:

$$v^\alpha = \sum_{I \subseteq k} (-1)^{|I|}(\alpha \bowtie I).$$

For example, for $k = 2$, the cog on 4 atoms $\alpha = \{a_1, b_1, a_2, b_2\}$ ordered as in (3) is the vector:

$$v^\alpha \quad = \quad \{a_1, a_2\} - \{a_1, b_2\} - \{b_1, a_2\} + \{b_1, b_2\} \quad \in \quad \text{Lin}\binom{\mathbb{A}}{2}.$$

---

2   Here W/V denotes the quotient space as usual.

Notice that, for a fixed k, all cogs form a single orbit in $\mathrm{Lin}\binom{\mathbb{A}}{k}$.

CLAIM 4.6. *Every nontrivial equivariant subspace $V \subseteq \mathrm{Lin}\binom{\mathbb{A}}{k}$ contains a cog.*

PROOF. Pick any nonzero $v \in V$ and pick $\alpha = \{a_1, \ldots, a_k\} \in \binom{\mathbb{A}}{k}$ so that $v(\alpha) \neq 0$. Choose fresh atoms $b_1, \ldots, b_k$ to form a set as in (3), so that no atom present in v lies strictly between a_i and b_i for any i. (This is possible because α and v are both finite.) We say that an atom c is *present* in v if there is a set $\beta \in \binom{\mathbb{A}}{k}$ such that $v(\beta) \neq 0$ and $c \in \beta$.

For every $i = 1, \ldots, k$, choose an atom automorphism π_i such that:

— $\pi_i(a_i) = b_i$,
— $\pi_i(a_j) = a_j$ and $\pi_i(b_j) = b_j$ for $j \neq i$, and
— $\pi_i(c) = c$ for all other atoms c present in v.

This is possible thanks to our choice of the b_i's. Then put $v_0 = v$ and define $v_1, \ldots, v_k \in V$ by induction:

$$v_i = v_{i-1} - \pi_i(v_{i-1})$$

(note that $\pi_i(v_{i-1}) \in V$ since V is equivariant). It is easy to prove by induction on i that:

— $\bigl(\pi_i(v_{i-1})\bigr)(\alpha) = 0$, hence
— $v_i(\alpha) = v(\alpha)$;

in particular v_k is nonzero. Also by straightforward induction, each v_i has the following properties for every $\beta \in \binom{\mathbb{A}}{k}$:

— $v_i(\beta)$ is nonzero only if for each $1 \leq j \leq i$, β contains either a_j or b_j but not both,
— if β' arises from such β by replacing a_j with b_j (or *vice versa*) for exactly one $j \leq i$, while keeping the other components unchanged, then $v_i(\beta) + v_i(\beta') = 0$.

For $i = k$ this implies that v_k, divided by the scalar $v(\alpha)$, is a cog. ∎

Claim 4.6 implies that $\mathrm{Lin}\binom{\mathbb{A}}{k}$ has a unique least nontrivial equivariant subspace: the one spanned by all cogs. We shall now give an explicit description of that subspace.

A vector $v \in \mathrm{Lin}\binom{\mathbb{A}}{k}$ is called *balanced* if for every set $S \in \binom{\mathbb{A}}{k-1}$, and for every S-orbit $I \subseteq \mathbb{A}$ such that $I \cap S = \emptyset$:

$$\sum_{a \in I} v(S \cup \{a\}) = 0. \tag{4}$$

The above sum is formally infinite, but only finitely many summands in it are non-zero because v is a finite vector. Note that if $S = \{a_1, \ldots, a_{k-1}\}$ where $a_1 < \cdots < a_{k-1}$, then S-orbits disjoint from S are exactly the k open intervals:

$$(-\infty, a_1), \ (a_1, a_2), \ldots, (a_{k-2}, a_{k-1}), \ (a_{k-1}, +\infty). \tag{5}$$

Balanced vectors form an equivariant subspace of $\mathrm{Lin}\binom{\mathbb{A}}{k}$: if v is balanced then $\pi(v)$ is balanced for every atom automorphism π, and the sum of two balanced vectors is balanced. We denote

the space of balanced vectors by B. An immediate corollary of Claim 4.6 is that every cog is balanced.

CLAIM 4.7. *B is the subspace spanned by all cogs.*

PROOF. It is enough to show that every balanced vector is a linear combination of cogs. So consider any nonzero vector $v \in B$. Let $\beta_1 \sqsubseteq \beta_2 \sqsubseteq \ldots \sqsubseteq \beta_M \in \binom{\mathbb{A}}{k}$ be the set of all sets that appear in v with nonzero coefficients, ordered lexicographically. More precisely, we define $\beta \sqsubseteq \beta'$ if

— $\min(\beta) < \min(\beta')$, or
— $\min(\beta) = \min(\beta') = a$ and $\beta \setminus \{a\} \sqsubseteq \beta' \setminus \{a\}$.

Consider

$$\beta_{\min} = \{a_1, \ldots, a_k\},$$

the least set in v in the lexicographic order. Assume $a_1 < a_2 < \cdots < a_k$.

Since v is balanced, there must be some $b_1 \neq a_1$ with $b_1 < a_2$ such that the set $\beta' = \{b_1, a_2, \ldots, a_k\}$ has a nonzero coefficient in v. (Consider $S = \{a_2, \ldots, a_k\}$ in (4).) This means that $\beta_{\min} \sqsubseteq \beta'$, and so $a_1 < b_1$.

Symmetrically, there must be some $b_k \neq a_k$ with $b_k > a_{k-1}$ such that the set $\beta' = \{a_1, \ldots, a_{k-1}, b_k\}$ has a nonzero coefficient in v. This means that $\beta_{\min} \sqsubseteq \beta'$, and so $a_k < b_k$.

More generally, for every $1 < i < n$ there must be some $b_i \neq a_i$ with $a_{i-1} < b_i < a_{i+1}$ such that the set $\beta_{\min}$ with a_i replaced by b_i has a nonzero coefficient in v. Because $\beta_{\min}$ is the least set in v lexicographically, it must hold that $a_i < b_i$.

We obtain the following order of $2k$ atoms:

$$\underbrace{a_1 < b_1 < \cdots < a_k < b_k}_{\alpha}$$

Note that every atom in this set is present in some set that appears in v with a nonzero coefficient.

Let v^α be the cog based on α. Note that for every β' with a nonzero coefficient in v^α we have $\beta_{\min} \sqsubseteq \beta'$.

Let $\lambda \neq 0$ be the coefficient of $\beta_{\min}$ in v. Consider the vector

$$w = v - \lambda \cdot v^\alpha.$$

This has the following crucial properties:

— every atom present in w is also present already in v, i.e., no new atoms are introduced in w,
— for every β' with a nonzero coefficient in w we have $\beta_{\min} \sqsubseteq \beta'$,
— the coefficient of $\beta_{\min}$ in w is zero, i.e., $\beta_{\min}$ does not appear in w.

By induction on the (finite) lexicographic order on $\binom{\mathbb{A}}{k}$ restricted to atoms present in v, the vector w, and therefore also v, is a linear combination of cogs. This concludes the proof of Claim 4.7. ∎

Claims 4.6 and 4.7 together imply that B is the smallest nontrivial equivariant subspace of Lin $\binom{\mathbb{A}}{k}$. To finish the proof of Lemma 4.5, for every $i \in \{1, \ldots, k\}$ consider the equivariant linear map:

$$g_i : \text{Lin } \binom{\mathbb{A}}{k} \to \text{Lin } \binom{\mathbb{A}}{k-1}$$

defined by:

$$g_i(v)(S) = \sum_{a \in I} v(S \cup \{a\})$$

for any $S \in \binom{\mathbb{A}}{k-1}$, where I is the i'th orbit on the list (5). Tupling these functions for all i we obtain an equivariant linear map:

$$g : \text{Lin } \binom{\mathbb{A}}{k} \to \left(\text{Lin } \binom{\mathbb{A}}{k-1}\right)^k.$$

By definition, the kernel of g is B, so there is a subspace embedding:

$$\left(\text{Lin } \binom{\mathbb{A}}{k}\right) / B \subseteq \left(\text{Lin } \binom{\mathbb{A}}{k-1}\right)^k.$$

Note that B does not have any nontrivial equivariant subspaces, so $\text{length}(B) = 1$. By Lemma 4.3(i) applied to $V = B$ and $W = \text{Lin } \binom{\mathbb{A}}{k}$ we thus obtain:

$$\text{length}\left(\text{Lin } \binom{\mathbb{A}}{k}\right) = 1 + \text{length}\left(\left(\text{Lin } \binom{\mathbb{A}}{k}\right) / B\right) \leq$$
$$\leq 1 + \text{length}\left(\left(\text{Lin } \binom{\mathbb{A}}{k-1}\right)^k\right) = 1 + k \cdot \text{length}\left(\text{Lin } \binom{\mathbb{A}}{k-1}\right)$$

and conclude the proof of Lemma 4.5. ∎

The following is now easy:

THEOREM 4.8. *The ordered atoms* $\mathbb{A} = (\mathbb{Q}, <)$ *have the finite length property. For an equivariant orbit-finite set* Q, *the space* $\text{Lin } Q$ *has length at most*

$$\text{(orbit count of } Q) \cdot (1 + \text{atom dimension of } Q)!$$

PROOF. If Q has only one orbit, then there is an equivariant surjective function from $\binom{\mathbb{A}}{k}$ to Q, where k is the atom dimension of Q (see [5, Lem. 3.20]). By induction on k, using Lemma 4.5, we obtain

$$\text{length}\left(\text{Lin } \binom{\mathbb{A}}{k}\right) \leq (1 + k)! \tag{6}$$

The theorem then follows by Lemma 4.3(iii). Finally, the case of multi-orbit Q follows from Lemma 4.3(ii). ∎

From this it is easy to deduce the finite length property for equality atoms.

COROLLARY 4.9. *The equality atoms* $\mathbb{A} = (\mathbb{N}, =)$ *have the finite length property. For an equivariant orbit-finite set Q, the space* $\mathsf{Lin}\, Q$ *has length at most*

$$\text{(orbit count of } Q) \cdot k! \cdot (1 + k)!$$

where k is the atom dimension of Q.

PROOF. Extend $\mathbb{A}$ with any total order isomorphic to that of the rational numbers. Any set V, equivariant over equality atoms, remains equivariant when considered as a set over ordered atoms along this order. In particular, if V is a vector space, any chain of subspaces

$$V_1 \subsetneq V_2 \subsetneq \cdots \subsetneq V$$

equivariant over equality atoms, remains a chain over ordered atoms. As a result, the length of V over equality atoms does not exceed the length of V over ordered atoms.

The set $\mathbb{A}^{(k)}$, a single-orbit set over equality atoms, when seen as an equivariant set over ordered atoms, has exactly $k!$ disjoint orbits. Each of these orbits is equivariantly (over ordered atoms) isomorphic to $\binom{\mathbb{A}}{k}$. Theorem 4.8 then implies that $\mathsf{Lin}\left(\mathbb{A}^{(k)}\right)$, over equality atoms, has length at most $k! \cdot (1 + k)!$.

Every single-orbit equivariant set Q is an image of an equivariant surjective function from $\mathbb{A}^{(k)}$, where k is the atom dimension of Q; the lemma follows by Lemma 4.3(iii). The case of multi-orbit Q follows from Lemma 4.3(ii).   ■

4.2   Length of finitely supported chains

So far, we have bounded the length of chains of equivariant subspaces. This will be enough to prove Theorem 3.5 for equivariant automata in Section 5. However, for generalising that theorem to finitely supported automata a finite bound on the length of chains of finitely supported spaces is needed. Here it is important that the chains are uniformly supported, i.e., all subspaces in a chain are required to have the same support. (Uniformly supported chains have been considered before in nominal domain theory [21].)

It is natural to conjecture that the finite length property (Definition 4.1) implies a finite bound on the length of such chains, at least for the atom structures that guarantee the existence of least supports (see [6, Section 9]). However, we are only able to prove it for the equality and ordered atoms.

THEOREM 4.10. *Consider the equality or ordered atoms. Let $\bar{a}$ be a tuple of atoms. For every $\bar{a}$-supported orbit-finite set Q, there is a finite upper bound on the length of chains of $\bar{a}$-supported subspaces of* $\mathsf{Lin}\, Q$.

PROOF. Define the $\bar{a}$-length of a vector space to be the maximal length of an $\bar{a}$-supported chain of subspaces. In the special case when $\bar{a}$ is the empty tuple, we get the notion of length from Definition 4.1. Every $\bar{a}$-supported orbit-finite set of atom dimension k can be obtained, using images under $\bar{a}$-supported functions and disjoint unions, from $\bar{a}$-orbits contained in $\mathbb{A}^k$. Since Lemma 4.3 holds, with the same proof, for $\bar{a}$-length, it is enough to show that the $\bar{a}$-length is finite for $\mathrm{Lin}\, Q$ when Q is a single $\bar{a}$-orbit contained in $\mathbb{A}^k$. We now split into two proofs, depending on whether we deal with the equality or ordered atoms.

— *Equality atoms.* Choose some bijection

$$f : \underbrace{(\mathbb{A} - \bar{a})}_{\substack{\text{atoms that do} \\ \text{not appear in } \bar{a}}} \to \mathbb{A},$$

which is possible since both sets are countably infinite. (Note that f cannot be finitely supported.) Let $\ell \in \{1, \ldots, k\}$ be the number of coordinates that are not from $\bar{a}$ in some (equivalently, every) tuple from the $\bar{a}$-orbit Q. We can lift f to an injective function (in fact, a bijection)

$$g : Q \to \mathbb{A}^{(\ell)},$$

which erases the coordinates that use atoms from $\bar{a}$ and applies f to the remaining coordinates. One can easily see that

$$\mathrm{Lin}\, g : \mathrm{Lin}\, Q \to \mathrm{Lin}\, \mathbb{A}^{(\ell)}$$

maps $\bar{a}$-supported subspaces of $\mathrm{Lin}\, Q$ to equivariant subspaces of $\mathrm{Lin}\, \mathbb{A}^{(\ell)}$, and preserves strict inclusions. Therefore the $\bar{a}$-length of $\mathrm{Lin}\, Q$ is at most the (equivariant) length of $\mathrm{Lin}\, \mathbb{A}^{(\ell)}$, and the latter length is finite by Corollary 4.9.

— *Ordered atoms.* Let the atoms in $\bar{a}$ be

$$a_1 < \cdots < a_n.$$

For every $i \in \{0, 1, \ldots, n\}$, consider the interval

$$\mathbb{A}_i = \{a : a_i < a < a_{i+1}\},$$

where a_0 is $-\infty$ and a_{n+1} is ∞. Choose some order-preserving bijections

$$f_i : \mathbb{A}_i \to \mathbb{A}.$$

As in the previous item, let ℓ be the coordinates from Q which avoid atoms from $\bar{a}$. By erasing the coordinates which use atoms from $\bar{a}$, and applying the appropriate functions f_i to the remaining coordinates, we get an injective function

$$g : Q \to \mathbb{A}^\ell.$$

Since the functions f_i all have the same co-domain, namely $\mathbb{A}$, the image of the function will contain tuples that are not necessarily strictly increasing. Nevertheless, the linear lifting

$$g : \operatorname{Lin} Q \to \operatorname{Lin} \mathbb{A}^\ell$$

maps $\bar{a}$-supported subspaces in $\operatorname{Lin} Q$ to equivariant subspaces in $\operatorname{Lin} \mathbb{A}^\ell$, and hence by Theorem 4.8 we obtain a finite bound on $\bar{a}$-supported chains in $\operatorname{Lin} Q$.

$\blacksquare$

4.3   A lower bound on length

Theorem 4.8 and Corollary 4.9 give upper bounds on the length of chains of equivariant vector spaces. We will now show a lower bound which is not quite matching, but exponential in the atom dimension.

We will focus on the equality atoms, where the case of interest is $\operatorname{Lin}(\mathbb{A}^{(k)})$ for a number k.

For $\alpha \in \mathbb{A}^{(k)}$ and $I \subseteq \{1, \ldots, k\}$, let $\alpha|_I$ denote the tuple α restricted to the coordinates from I, and define $B_I(\alpha) \subseteq \mathbb{A}^{(k)}$ to be the set of those tuples β for which $\beta|_I = \alpha|_I$.

Given I, define the space $V_I \subseteq \operatorname{Lin}(\mathbb{A}^{(k)})$ as the set of all vectors v such that, for every $\alpha \in \mathbb{A}^{(k)}$:

$$\sum_{\beta \in B_I(\alpha)} v(\beta) \quad = \quad 0.$$

For example, for $I = \emptyset$ the space V_I is the space of those vectors where all coefficients add up to 0. For $I = \{1, \ldots, n\}$ we have $B_I(\alpha) = \{\alpha\}$ and the space V_I is trivial.

If $I \subseteq J \subseteq \{1, \ldots, n\}$ then each set $B_I(\alpha)$ is a disjoint union of sets of the form $B_J(\beta)$ for some tuples β. As a result, $V_J \subseteq V_I$. Our main observation is a kind of converse to this:

LEMMA 4.11. *For every* $I, J_1, \ldots, J_m \subseteq \{1, \ldots, k\}$*, if* $I \nsubseteq J_i$ *for every* $i = 1, \ldots, m$ *then:*

$$V_{J_1} \cap \cdots \cap V_{J_m} \nsubseteq V_I.$$

PROOF. Assume sets I and $J_1, \ldots, J_m$ as above. Fix some pairwise distinct atoms c_i for $i \notin I$ and a_i, b_i for $i \in I$. For every subset $H \subseteq I$, define the tuple $\alpha_H = (d_1, \ldots, d_k) \in \mathbb{A}^{(k)}$ by:

$$d_i = \begin{cases} a_i & \text{if } i \in H \\ b_i & \text{if } i \in I \setminus H \\ c_i & \text{if } i \notin I. \end{cases}$$

Define the vector:

$$v = \sum_{H \subseteq I} (-1)^{|H|} \cdot \alpha_H.$$

Then $v \in V_{J_i}$ for each $i = 1, \ldots, m$, but $v \notin V_I$.

$\blacksquare$

COROLLARY 4.12. *The space* $\mathsf{Lin}(\mathbb{A}^{(k)})$ *admits a chain of equivariant subspaces of length* 2^k.

PROOF. Consider any enumeration of all subsets of $\{1, \ldots, k\}$:

$$J_1, J_2, \ldots, J_{2^k}$$

such that $J_i \nsubseteq J_j$ for all $i > j$. This is easy to achieve: order the subsets by their size (increasingly), and subsets of equal size can be enumerated in any order. Then define spaces:

$$W_i = \bigcap_{j \leq i} V_{J_j}.$$

Obviously $W_1 \supseteq W_2 \supseteq \cdots \supseteq W_{2^k}$, and by Lemma 4.11 all the inclusions are strict. ■

4.4   Failure of finite length: vector atoms

In this section we consider, as the structure of atoms, the countable infinite-dimensional vector space over $\mathbb{Z}_2$. We will show that these atoms do not have the finite length property. This is a rather confusing case in the context of this paper: in general we study vectors built of atoms, and now we wish atoms themselves to be vectors. We will therefore present the structure and its basic properties step by step.

Let $\mathbb{D} = \{a, b, c, d, \ldots\}$ be a countably infinite set of generators. We consider, as the set of atoms,

$$\mathbb{A} = \mathsf{Lin}(\mathbb{D}),$$

the free vector space, over the field $\mathbb{Z}_2$, generated by $\mathbb{D}$. Vectors in this space (i.e., our atoms) are formal combinations of generators; because we work over $\mathbb{Z}_2$, these formal combinations are simply finite sets of generators, with addition defined by symmetric difference. We will represent atoms simply by listing the generators in them, with the zero vector denoted by 0:

$$0, a, b, c, \ldots, ab, ac, \ldots, bce, \ldots$$

and their addition will be denoted by $\oplus$, so that e.g.:

$$abc \oplus bcd = ad.$$

To distinguish atoms from generators we will, in this section only, denote them by u, v, w etc.

The space $\mathbb{A}$ can be seen as a relational structure in a standard way, with one constant $0 \in \mathbb{A}$ and one ternary relation

$$R(u, v, w) \iff u \oplus v = w.$$

The automorphisms of this structure are exactly linear automorphisms of $\mathbb{A}$ seen as a vector space. It is a standard result that the structure $\mathbb{A}$ is oligomorphic.

For $v \in \mathbb{A}$ and $X \subseteq \mathbb{A}$, we write as usual: $v \oplus X = \{v \oplus w \mid w \in X\}$.

We will consider finite-dimensional subspaces of $\mathbb{A}$. For a number $n > 0$, an n-dimensional subspace of $\mathbb{A}$ has exactly 2^n elements. One way to construct such a subspace is to take the subspace spanned by some chosen n generators from $\mathbb{D}$. For example,

$$\{0, a, b, c, ab, ac, bc, abc\}$$

is a 3-dimensional subspace spanned by a, b and c. Other subspaces exist; for example,

$$\{0, ab, ac, bc\}$$

is a 2-dimensional subspace.

For subsets $X, Y \subseteq \mathbb{A}$, we denote their symmetric difference (as sets) by $X \div Y$.

LEMMA 4.13. *For all $v \in \mathbb{A}$ and $X, Y, Z \subseteq \mathbb{A}$:*

(i) $X \cap (Y \div Z) = (X \cap Y) \div (X \cap Z)$,

(ii) $v \oplus (X \cap Y) = (v \oplus X) \cap (v \oplus Y)$,

(iii) $v \oplus (X \div Y) = (v \oplus X) \div (v \oplus Y)$.

PROOF. Routine exercise.   ∎

We will be interested in symmetric differences of finite families of vector spaces of the same finite dimension.

LEMMA 4.14. *For numbers $n > m > 0$, let $V \subseteq \mathbb{A}$ be an n-dimensional vector space and let $V_1, V_2, \ldots, V_k$ be all its m-dimensional subspaces. Then*

$$V_1 \div V_2 \div \cdots \div V_k = V. \tag{7}$$

PROOF. It is a standard exercise to show that

$$k = \frac{(2^n - 1)(2^n - 2) \cdots (2^n - 2^{m-1})}{(2^m - 1)(2^m - 2) \cdots (2^m - 2^{m-1})}.$$

Both the numerator and the denominator of this fraction are divisible by the same power of 2; specifically it is $\frac{m(m-1)}{2}$. As a result, k is an odd number.

For any vector $v \in V$, let $\alpha(v)$ be the number of subspaces V_i that contain v. Note that, due to the symmetry under automorphisms of V, for all non-zero $v, w \in V$ there is $\alpha(v) = \alpha(w)$; denote this number by α.

Now calculate:

$$k \cdot 2^m = \sum_{i=1}^{k} |V_i| = \sum_{v \in V} \alpha(v) = k + \sum_{v \in V \setminus \{0\}} \alpha(v) = k + (2^n - 1) \cdot \alpha.$$

The number on the left is even and k is odd, so α must be odd.

As a result, *every* vector in V belongs to an odd number of subspaces V_i, which yields (7).   ∎

The following basic property of finite-dimensional spaces over $\mathbb{Z}_2$ will be important for our purposes.

LEMMA 4.15. *For every finite family $A_1, A_2, \ldots, A_k$ of n-dimensional vector subspaces of $\mathbb{A}$, the symmetric difference $A_1 \div A_2 \div \cdots \div A_k$ is either empty or has size at least 2^n.*

PROOF. See Appendix B.   ■

To see that Lemma 4.15 is not entirely obvious, let us have a look at some examples. First of all, the symmetric difference from the lemma may be empty even if all A_i are distinct. For an example, consider $n = 2$, $k = 4$ and vector subspaces:

$$V_1 = \{0, a, b, ab\} \qquad\qquad V_2 = \{0, b, c, bc\}$$
$$V_3 = \{0, a, c, ac\} \qquad\qquad V_4 = \{0, ab, bc, ac\}.$$

The symmetric difference may also be of size strictly between 2^n and 2^{n+1}, so in particular it may not be a vector subspace. For an example, consider $n = 3$, $k = 2$ and vector subspaces:

$$V_1 = \{0, a, b, c, ab, bc, ac, abc\}$$
$$V_2 = \{0, a, d, e, ad, de, ae, ade\};$$

here $|V_1 \div V_2| = 12$. Another example, for $n = 2$ and $k = 3$, is:

$$V_1 = \{0, a, b, ab\} \qquad V_2 = \{0, b, c, bc\} \qquad V_3 = \{0, c, d, cd\};$$

here $|V_1 \div V_2 \div V_3| = 6$. If we additionally consider

$$V_4 = \{0, ab, bc, ac\},$$

the symmetric difference becomes $\{a, d, ac, cd\}$, showing that the symmetric difference may not be a vector subspace even if its size is 2^n.

So far in this section we studied properties of the vector atoms, but have not considered vector spaces *over* those atoms yet. In fact we will need to consider only one such vector space and its equivariant subspaces: Lin $\mathbb{A}$, the free vector space generated by the set of atoms, over the field $\mathbb{Z}_2$. Vectors in Lin $\mathbb{A}$ are formal linear combinations of atoms, such as:

$$0 + a + b + ab \qquad \text{or} \qquad a + bc + ac + bde. \tag{8}$$

The addition symbol above is used to build formal combinations of atoms, and it should not be confused with $\oplus$, which is an operation to build atoms from other atoms. Thanks to the choice of the field, combinations as above can be seen simply as finite sets of atoms, with symmetric difference $\div$ as addition.

The main result of this section is that Lin $\mathbb{A}$ has infinite length, so the vector atoms $\mathbb{A}$ does not have the finite length property.

THEOREM 4.16. Lin $\mathbb{A}$ *has an infinite chain of equivariant subspaces.*

PROOF. For any number n, let $W_n \subseteq \text{Lin}\,\mathbb{A}$ be the subspace spanned by those vectors in $\text{Lin}\,\mathbb{A}$ that, considered as subsets of $\mathbb{A}$, are n-dimensional vector subspaces of $\mathbb{A}$. (For example, in (8) the first vector is a subspace of $\mathbb{A}$, the second one is not.) It is easy to see that each space W_n is equivariant.

For any $n > m$, the inclusion $W_n \subseteq W_m$ follows from Lemma 4.14. Moreover, this inclusion is strict by Lemma 4.15. Indeed, every nonzero vector in W_n has at least 2^n elements, and W_m contains vectors of smaller size.

As a result, there is an infinite chain of subspaces:

$$\cdots \subsetneq W_3 \subsetneq W_2 \subsetneq W_1 \subsetneq \text{Lin}\,\mathbb{A}.$$

■

Note that the above infinite chain of subspaces is decreasing. We conjecture that for all oligomorphic atom structures $\mathbb{A}$, all spaces $\text{Lin}\,\mathbb{A}^k$ are *noetherian*, i.e. they do not admit *increasing* infinite chains of equivariant subspaces.

5.   The equivalence algorithm

In this section, we use results from Section 4 to complete the proof of Theorem 3.5. We assume that the atoms $\mathbb{A}$ are the equality or ordered atoms.

LEMMA 5.1. *Let $\mathcal{A}_1$ and $\mathcal{A}_2$ be equivariant weighted orbit-finite automata, and let n be their orbit count and k the atom dimension as in Theorem 3.5. If the recognized weighted languages are different, then this difference is witnessed by some input word of length at most*

$$2^{\text{poly}(k)} \cdot n.$$

PROOF. Let Q be the state space of the difference automaton $\mathcal{A}_1 - \mathcal{A}_2$ as described in the reduction from equivalence to zeroness. For an input word $w \in \Sigma^*$, let $[w] \in \text{Lin}\,Q$ be its corresponding configuration in the difference automaton, and let $V_i \subseteq \text{Lin}\,Q$ be the subspace that is spanned by configurations of input words of length at most i. By Theorem 4.8/Corollary 4.9, we know that the chain $\{V_i\}_i$ must stabilize in a number of steps that is bounded as in the statement of the lemma. This means that for every input word, its configuration is a linear combination of configurations of short words; in particular the difference automaton can produce a nonzero output if and only if it can produce a nonzero output on a short input word. ■

At this point, we could solve the equivalence problem by guessing a short differentiating word, leading to a nondeterministic algorithm for non-equivalence, with running time as in the bound from Lemma 5.1. We can, however, improve this to get a deterministic algorithm, by

using a reduction to the equivalence problem for finite weighted automata. Short input words necessarily use few atoms, and therefore the equivalence problem boils down to checking equivalence for input words that have few atoms. The latter problem is solved in the following lemma.

LEMMA 5.2. *Consider the following problem.*

— **Input.** *Two equivariant weighted orbit-finite automata $\mathcal{A}_1$ and $\mathcal{A}_2$ and a number $\ell \in \{1, 2, \ldots\}$;*

— **Question.** *Are the two automata equivalent on all input words supported by at most ℓ atoms?*

This problem can be solved in time polynomial in $n \cdot \ell^k$, where the parameters n and k are defined as in Theorem 3.5.

PROOF. Choose a tuple $\bar{a}$ of ℓ atoms. Since the automata are equivariant, they are equivalent on inputs with at most ℓ atoms if and only if they are equivalent on inputs words supported by $\bar{a}$.

By equivariance and condition (*) from Definition 3.1, if a state q and an input letter σ are both supported by $\bar{a}$, then the same is true for every state p such that (q, σ, p) is a transition with nonzero weight. Therefore, when restricted to input words that are supported by $\bar{a}$, both automata only use states that are supported by $\bar{a}$. These observations motivate the following definition: for $i \in \{1, 2\}$, let $\mathcal{A}_{i,\bar{a}}$ the weighted automaton that is obtained from $\mathcal{A}_i$ by restricting the states and alphabet to elements supported by $\bar{a}$, and restricting the transitions and weight functions to the new alphabet and states. This automaton is finite, with size bounded by $n \cdot \ell^k$. The lemma follows by applying the polynomial time algorithm for equivalence of finite weighted automata. ∎

The above two lemmas complete the proof of Theorem 3.5. The algorithm for zeroness is simply the algorithm from Lemma 5.2, with ℓ being the bound from Lemma 5.1.

Notice how our algorithm does not rely on an explicit computation of the chain of subspaces $\{V_i\}_i$ from the proof of Lemma 5.1. Doing so would require a method for calculating (a representation of) V_{i+1} from (a representation of) V_i, and a method of checking whether two subspaces of $\mathrm{Lin}\, Q$ are equal, to detect stabilisation. Variants of the latter problem, under the guise of solving orbit-finite systems of linear equations, are studied in [10].

6.   Vector spaces with atoms

In this section, we study in more depth vector spaces that are spanned by orbit-finite sets. Apart from their independent interest, these results will be used in Section 7 to minimize weighted automata, and in Section 8 to decide equivalence for unambiguous automata.

Vector spaces of the form Lin Q, where Q is an orbit-finite set, could be called *orbit-finite-dimensional*, since they have an orbit-finite basis. These spaces are general enough to treat weighted orbit-finite automata. Some other spaces admit orbit-finite bases too; in [10, Thm. 3.3], this is shown for a variant of Lin Q where vectors are infinite (but finitely supported) linear combinations, rather than finite ones. Nevertheless, vector spaces with orbit-finite bases are not very robust. The reason is that extracting a basis uses choice, see [3, Thm. 1], and the principle of choice fails for sets with atoms.

EXAMPLE 6.1. Recall the vector space Lin $\mathbb{A}$ that was discussed in Example 4.2, and the subspace V that was spanned by

$$X = \{a - b : a \neq b \in \mathbb{A}\}.$$

The set X is not a basis, since the vectors $a - b$ and $b - a$ (or $a - b$, $b - c$ and $a - c$) are linearly dependent. However, X is a single equivariant orbit, so it does not have any nonempty proper subset that is equivariant. Therefore, no equivariant subset of X is a basis. In fact, V does not have any equivariant basis, even if we allow bases that are not contained in X. It does have, however, a finitely supported basis contained in X: for any fixed atom $a_0 \in \mathbb{A}$, the set

$$\{a_0 - b : b \in \mathbb{A} - \{a_0\}\}$$

is a basis.   ◆

In the previous example, there was no equivariant basis, but there was a finitely supported one. In the next example there is no finitely supported basis at all.

EXAMPLE 6.2. Consider the equality atoms and the space $\mathsf{Lin}(\mathbb{A}^{(2)})$ over the field of rationals. A vector in this space can be visualized as a weighted directed finite graph, where vertices are atoms and the weight of an edge (a, b) is the corresponding coefficient in the vector. Consider the subspace of $\mathsf{Lin}(\mathbb{A}^{(2)})$ spanned by

$$X = \{(a, b) - (b, a) : a \neq b \in \mathbb{A}\}.$$

This subspace consists of graphs where for every pair of vertices, the connecting edges in both directions have opposite weights. We claim that there is no finitely supported subset of X that is a basis of this space. Indeed, suppose that $Y \subseteq X$ is a finitely supported basis. It is easy to see that for every two distinct atoms a, b, the set Y must contain one of the vectors

$$(a, b) - (b, a) \qquad \text{or} \qquad (b, a) - (a, b).$$

If the atoms a, b are fresh (i.e. they do not belong to the least support of Y), by swapping these two atoms we map one of the vectors to the other, hence *both* vectors must belong to Y. However, the two vectors are linearly dependent and so Y is not a basis.

Using a similar argument one can show that the subspace spanned by X does not have any finitely supported basis, even if we allow bases that are not contained in X. ◆

The above example shows that vector spaces with an orbit-finite basis are not closed under finitely supported subspaces. The same issue appears with finitely supported quotients (images of surjective linear maps). Indeed, the spaces from Examples 6.1 and 6.2 are images of equivariant linear maps from $\mathsf{Lin}(\mathbb{A}^{(2)})$ to, respectively, $\mathsf{Lin}(\mathbb{A})$ and $\mathsf{Lin}(\mathbb{A}^{(2)})$. These issues will become a problem in the next section, where we minimize weighted orbit-finite automata, since minimization will require taking subspaces and quotients. To deal with these issues, we notice that all spaces that we have considered so far, even if they lack an orbit-finite basis, have orbit-finite spanning sets. They are therefore instances of the following general notion:

DEFINITION 6.3 (Orbit-finitely spanned vector space). Fix a field $\mathbb{F}$ and an atom structure. A *vector space with atoms* is a vector space such that the underlying set V is a set with atoms, and the operations

$$\underbrace{+ : V \times V \to V}_{\text{binary addition}} \qquad \underbrace{\cdot : \mathbb{F} \times V \to V}_{\text{scalar multiplication}}$$

are finitely supported. A vector space with atoms is called *orbit-finitely spanned* if there is an orbit-finite subset of the vector space which spans it.

As we have seen in Example 6.2, some orbit-finitely spanned vector spaces do not have any orbit-finite basis. They do, however, have finite length:

LEMMA 6.4. *For atoms with the finite length property, equivariant orbit-finitely spanned vector spaces have finite length.*

PROOF. Let V be spanned by an equivariant orbit-finite set Q, and let

$$f \colon \mathsf{Lin}\,Q \to V$$

be the unique linear map which extends the inclusion of Q in V. This is a surjective finitely supported linear map, so (by Lemma 4.3(iii)) the length of V is not greater than the length of $\mathsf{Lin}\,Q$, which is finite by the finite length property. ■

THEOREM 6.5 (Closure properties). *If the atoms have the finite length property, then equivariant orbit-finitely spanned vector spaces are closed under equivariant quotients, equivariant subspaces, direct sums and tensor products.*

PROOF. Quotients and direct sums are immediate. For the tensor product $U \otimes V$, suppose U and V are spanned by orbit-finite sets Q and P. Since $\mathsf{Lin}\,Q \otimes \mathsf{Lin}\,P$ is isomorphic to $\mathsf{Lin}(Q \times P)$, it follows that $U \otimes V$ is spanned by $Q \times P$, which is orbit-finite.

We are left with the subspaces. Suppose that V is equivariant and orbit-finitely spanned, and let U be an equivariant subspace of V. Construct an increasing chain of equivariant subspaces of U

$$U_0 \subseteq U_1 \subseteq U_2 \subseteq \cdots \subseteq U$$

as follows:

— U_0 is the trivial space,
— if $U_n \neq U$ then pick any vector $u_{n+1} \in U \setminus U_n$, and let U_{n+1} be the subspace of U spanned by all vectors in U_n and all vectors in the orbit of u_{n+1}.

Each U_n is an equivariant subspace of the orbit-finitely-spanned V, so by Lemma 6.4 the chain cannot grow forever, hence $U = U_n$ for some n. Of course, U_n is spanned by all vectors in the orbits of $u_1, \ldots, u_n$. ∎

Thanks to Theorem 4.10, in the equality and ordered atoms, Lemma 6.4 and Theorem 6.5 hold also for orbit-finitely spanned vector spaces that are not necessarily equivariant.

In the remainder of this section, we discuss another closure property of orbit-finitely spanned vector spaces: if V is an equivariant orbit-finitely spanned vector space, then the same is true for its *finitely supported dual*, which is the vector space of finitely supported linear maps from V to $\mathbb{F}$. We will prove closure under finitely supported duals for the equality and ordered atoms.

Before considering finitely supported duals, we discuss a slightly simpler vector space. For Q an orbit-finite set, define

$$Q \xrightarrow{\text{fs}} \mathbb{F}$$

to be the set of finitely supported functions from Q to the field $\mathbb{F}$. This set can be viewed as a vector space, with coordinate-wise addition and scalar multiplication. We have

$$\text{Lin}\, Q \ \subseteq \ Q \xrightarrow{\text{fs}} \mathbb{F},$$

and the inclusion is strict when Q is infinite, as witnessed by the function that returns 1 on all arguments.

EXAMPLE 6.6. For the equality atoms, the space $\mathbb{A} \xrightarrow{\text{fs}} \mathbb{F}$ is spanned by the functions

$$\{f_{\mathbb{A}}\} \ \cup \ \{f_a : a \in \mathbb{A}\} \tag{9}$$

where $f_{\mathbb{A}}$ maps all atoms to 1, while f_a maps a to 1 and the remaining atoms to 0. Indeed, every function $f : \mathbb{A} \to \mathbb{F}$ supported by an atom tuple $\bar{a}$ can be presented as:

$$f = \underbrace{f(c)}_{\substack{c \text{ is some fresh} \\ \text{atom not in } \bar{a}}} \cdot f_{\mathbb{A}} + \sum_{a \in \bar{a}} (f(a) - f(c)) \cdot f_a.$$

The functions from (9) are linearly independent, hence they form an orbit-finite basis of $\mathbb{A} \xrightarrow{\text{fs}} \mathbb{F}$.

A similar result holds for the ordered atoms, except that there the spanning set is

$$\{f_{\mathbb{A}}\} \cup \{f_a : a \in \mathbb{A}\} \cup \{f_{>a} : a \in \mathbb{A}\}$$

where $f_{\mathbb{A}}$ and f_a are as before and $f_{>a}$ maps b to 1 precisely if $b > a$.   ◆

The above example shows that $\mathbb{A} \xrightarrow{\text{fs}} \mathbb{F}$ is orbit-finitely spanned. This is true for every orbit-finite set, not just $\mathbb{A}$:

THEOREM 6.7. *Assume the equality or the ordered atoms. If Q is an orbit-finite set, then $Q \xrightarrow{\text{fs}} \mathbb{F}$ is orbit-finitely spanned.*

PROOF. To simplify notation, we assume that Q is equivariant. By Theorem 6.5 this implies the general case anyway, because every orbit-finite set Q is contained in some equivariant orbit-finite set $\bar{Q}$, hence $Q \xrightarrow{\text{fs}} \mathbb{F}$ is a subspace of $\bar{Q} \xrightarrow{\text{fs}} \mathbb{F}$. Furthermore, for the same reasons, it is safe to assume that Q is a single-orbit set.

We need to exhibit an orbit-finite set Φ of finitely supported functions, so that every finitely supported function $f : Q \to \mathbb{F}$ is a linear combination of functions from Φ. It is enough to show this for the case where f is the characteristic function of a finitely supported set $R \subseteq Q$, since such characteristic functions are easily seen to span the space of all finitely supported functions. So consider such a set R, supported by some tuple $\bar{a}$ of atoms. R is then a disjoint union of $\bar{a}$-orbits, and it is enough to consider the case where R is a single $\bar{a}$-orbit.

From here, the arguments for the equality and the ordered atoms begin to differ.

For the case of **ordered atoms**, we know that the single-orbit Q is in equivariant bijection with $\binom{\mathbb{A}}{k}$ for some number k, so it is enough to deal with $Q = \binom{\mathbb{A}}{k}$. Let R be the $\bar{a}$-orbit of some

$$r = \{a_1, \dots, a_k\} \in \binom{\mathbb{A}}{k}.$$

Each a_i belongs to some $\bar{a}$-orbit of atoms. This orbit is an interval, so it is equal to the orbit determined by at most two atoms from $\bar{a}$. (One atom is enough if a_i belongs to $\bar{a}$, or if it is larger than (or smaller than) every atom in $\bar{a}$.) Picking these atoms for each $i = \{1, \dots, k\}$, we obtain a tuple $\bar{b} \subseteq \bar{a}$ of at most $2k$ atoms, such that R is equal to the $\bar{b}$-orbit of r. As a consequence, R is supported by $\bar{b}$.

As a result, we may take Φ to be the set of characteristic functions of subsets of Q supported by at most $2k$ atoms. Since k is fixed for a given Q, this is an orbit-finite set.

The argument for the **equality atoms** is only a little more complicated. Here for Φ we take the set of characteristic functions of subsets of Q supported by at most k atoms.

We know that for some number k there is an equivariant surjection from $\mathbb{A}^{(k)}$ to Q, so it is enough to deal with $Q = \mathbb{A}^{(k)}$. Let R be the $\bar{a}$-orbit of some

$$r = (a_1, \ldots, a_k) \in \mathbb{A}^{(k)}.$$

We proceed by induction on the number of atoms in r that are *not* in $\bar{a}$. If this number is zero, then R is the singleton of r and the characteristic function of R is in Φ, so there is nothing to do. If there are some atoms in r that are not present in $\bar{a}$, denote by $I \subseteq \{1, \ldots, k\}$ the set of coordinates where these atoms occur in r. Then an $s \in \mathbb{A}^{(k)}$ belongs to R if and only if:

(i) s is equal to r on all coordinates apart from those in I,

(ii)   atoms in s on coordinates from I do not belong to $\bar{a}$.

Let $\hat{R} \subseteq Q$ be the set of tuples s that satisfy condition (i) above; obviously $R \subseteq \hat{R}$. Moreover, $\hat{R}$ has a support of size $k - |I|$, so $\chi_{\hat{R}}$, the characteristic function of $\hat{R}$, belongs to Φ.

Let $\{r_1, \ldots, r_n\}$ be the set of all k-tuples in $\mathbb{A}^{(k)}$ that can be obtained from r by replacing some (not necessarily all, but at least one) atoms on coordinates from I by some atoms from $\bar{a}$. Let R_i be the $\bar{a}$-orbit of r_i, for each i. The sets R_i are pairwise disjoint, and their union is equal to the difference $\hat{R} \setminus R$, therefore there is a linear equation:

$$\chi_R = \chi_{\hat{R}} - \sum_{i=1}^{n} \chi_{R_i}.$$

By the inductive assumption, each χ_{R_i} is a linear combination of functions from Φ, which completes the proof. ∎

COROLLARY 6.8. *Assume the equality or ordered atoms. If V is an orbit-finitely spanned vector space, then the same is true for its finitely supported dual.*

PROOF. Recall that the finitely supported dual of V is the vector space of linear maps from V to $\mathbb{F}$. If V is spanned by an orbit-finite set Q, then its finitely supported dual embeds into the space $Q \xrightarrow{\text{fs}} \mathbb{F}$ by precomposing with the inclusion of Q in V; the latter space is orbit-finitely spanned by Theorem 6.7. The corollary follows since, by Theorem 6.5, orbit-finitely spanned vector spaces are closed under subspaces. ∎

In contrast to the finite-dimensional case, orbit-finitely spanned vector spaces are not isomorphic to their finitely supported duals. Indeed, Example 6.6 shows that the finitely supported dual of Lin $\mathbb{A}$ is equivariantly isomorphic to the space Lin$(1 + \mathbb{A})$. There is no linear isomorphism between them which is equivariant, or even finitely supported.

The following example shows that the restriction to equality or ordered atoms was important in Theorem 6.7. This example hints on the difficulties one may encounter when generalizing our theory to, say, arbitrary oligomorphic atom structures.

EXAMPLE 6.9. Consider the graph atoms [5, Section 7.3.1], i.e. the case when the atom structure $\mathbb{A}$ is the Rado graph. This is an undirected graph with the key property that for all finite sets $U \subseteq S \subseteq \mathbb{A}$ there is an atom $a \in \mathbb{A}$ that has an edge to every atom in U and does not have an edge to any atom in $S \setminus U$.

We will show that the vector space $\mathbb{A} \xrightarrow{\text{fs}} \mathbb{F}$, for $\mathbb{F} = \mathbb{Z}_2$, does not have an orbit-finite spanning set.

Assume to the contrary that such an orbit-finite spanning set $X \subseteq \mathbb{A} \xrightarrow{\text{fs}} \mathbb{F}$ exists. Let k be a number such that every function in X is supported by at most k atoms; such a number exists since X is orbit-finite. Fix any finite set T of more than k atoms. Define $\chi_T : \mathbb{A} \to \mathbb{F}$ to be the characteristic function of the common neighborhood of T, i.e. the function which maps an atom to 1 if it has an edge to all atoms in T, and otherwise maps the atom to 0. This function is supported by T, so it belongs to $\mathbb{A} \xrightarrow{\text{fs}} \mathbb{F}$.

Since the whole space is spanned by X, there must be a linear equation:

$$\chi_T = \sum_{i=1}^{n} f_i \tag{10}$$

with $f_i \in X$. Each f_i is supported by some set $S_i \subseteq \mathbb{A}$ with $|S_i| \leq k$. Denote

$$S = T \cup \bigcup_{i=1}^{n} S_i.$$

For every subset $U \subseteq T$, pick an atom a_U such that:

— a_U is a neighbour of every atom in U, and

— a_U is *not* a neighbour of any atom in $S \setminus U$.

Such an atom exists by the universal property of the Rado graph.

Note that $\chi_T(a_U) = 1$ if and only if $U = T$.

Now consider a function f_i from (10), supported by S_i as above. Pick some atom $b \in T \setminus S_i$; this is possible since $|T| > |S_i|$. Since f_i is supported by S_i, the value of f_i on any atom a depends only on how a is connected by edges to atoms in S_i. This means that, for any $U \subseteq T$,

$$f_i(a_{U \setminus \{b\}}) = f_i(a_U) = f_i(a_{U \cup \{b\}})$$

(for every U exactly one of these equations is trivial). As a result, the number of U's for which $f_i(a_U) = 1$ is even. Summing up (10) over all $U \subseteq T$ we therefore obtain:

$$1 = \sum_{U \subseteq T} \chi_T(a_U) \overset{(10)}{=} \sum_{U \subseteq T} \sum_{i=1}^{n} f_i(a_U) = \sum_{i=1}^{n} \sum_{U \subseteq T} f_i(a_U) = \sum_{i=1}^{n} 0 = 0,$$

a contradiction.   ◆

7.   Minimization

Schützenberger's original paper on weighted automata contained a minimization procedure. We now describe a version of that procedure in the orbit-finite setting using the theory of orbit-finitely spanned vector spaces from Section 6.

Consider a weighted orbit-finite automaton $\mathcal{A}$ with states Q. We assume that the automaton is reachable, in the following sense: every vector in $\mathrm{Lin}\,Q$ is a linear combination of configurations corresponding to input words. For a vector $v \in \mathrm{Lin}\,Q$, define the weighted language of v to be the weighted language recognized by the automaton obtained from $\mathcal{A}$ by setting the initial map to v, i.e. the initial weight of a state is its coefficient in the vector v. Define the *syntactic congruence* $\sim$ to be the equivalence relation on $\mathrm{Lin}\,Q$ which identifies two vectors if the corresponding weighted languages are equal. It is not hard to see that $\sim$ is closed under both linear combinations and applying atom automorphisms; as a result, one can speak of an equivariant quotient vector space $(\mathrm{Lin}\,Q)/_\sim$. Equivalently, this is the quotient of the vector space $\mathrm{Lin}\,Q$ under the subspace which consists of vectors v whose corresponding weighted language is 0 everywhere. This quotient space is orbit-finitely spanned by the equivalence classes of states in Q.

In the finite dimensional case studied by Schützenberger, a minimal automaton is obtained by choosing some basis for this vector space, and using it as the state space of the minimal automaton. This idea, however, will not work in the orbit-finite setting, due to the difficulties with finding a basis that were described in Examples 6.1 and 6.2.

EXAMPLE 7.1. For the equality atoms and the field of rational numbers, consider the weighted language $L \colon \mathbb{A}^* \to \mathbb{F}$:

$$
L(w) = \begin{cases} 1 & \text{if } w = aba \text{ for some } a \neq b \in \mathbb{A}, \\ -1 & \text{if } w = abb \text{ for some } a \neq b \in \mathbb{A}, \\ 0 & \text{otherwise.} \end{cases}
$$

This is recognized by a weighted orbit-finite automaton where the set of states Q is

$$
\underbrace{\{\bot\}}_{\substack{\text{initial weight 1} \\ \text{final weight 0}}} + \underbrace{\mathbb{A} + \mathbb{A}^{(2)}}_{\substack{\text{initial weight 0} \\ \text{final weight 0}}} + \underbrace{\{\top\}}_{\substack{\text{initial weight 0} \\ \text{final weight 1}}}
$$

and the transitions with nonzero weight are

$$
\underbrace{(\bot, a, a),\ (a, b, (a, b)),\ ((a, b), a, \top),}_{\text{weight 1}}\ \underbrace{((a, b), b, \top)}_{\text{weight -1}}
$$

for every $a \neq b \in \mathbb{A}$. For the syntactic congruence $\sim$, it is not hard to see that $(a, b) \sim -(b, a)$ for every $a \neq b$. The quotient space $\mathrm{Lin}\, Q/_\sim$ is:

$$\mathrm{Lin}(\,\{\bot, \top\} + \mathbb{A}\,) \ \oplus \ X,$$

where X is the space from Example 6.2. This space has no finitely supported basis.   ◆

This example shows that in the orbit-finite setting the minimization procedure can leave the realm of weighted orbit-finite automata as defined in Definition 3.1. To overcome this issue, we use an alternative model for weighted automata, which we call *orbit-finitely spanned automata*. These are deterministic automata where the state spaces are orbit-finitely spanned vector spaces and all weight functions are linear maps.

DEFINITION 7.2 (Orbit-finitely spanned automaton). An orbit-finitely spanned automaton consists of:

1. an orbit-finite input alphabet Σ;
2. an orbit-finitely spanned vector space V;
3. a finitely supported transition function

$$\delta : V \times \Sigma \ \to \ V$$

such that $v \mapsto \delta(v, a)$ is a linear map for every $a \in \Sigma$;
4. an initial vector $v_0 \in V$;
5. a finitely supported linear map $F : V \to \mathbb{F}$.

An orbit-finitely spanned automaton recognizes a weighted language as expected: given an input word, it computes an element of $\mathbb{F}$ by starting in the initial vector, then applying the transition functions corresponding to the input letters, and finally applying the final function F.

THEOREM 7.3. *Orbit-finitely spanned automata and weighted orbit-finite automata recognize the same weighted languages.*

PROOF. The proof resembles an analogous construction in the finite-dimensional setting, with one important difference. If one converts a weighted orbit-finite automaton with states Q to an orbit-finitely spanned automaton in the natural way, then the resulting vector space $\mathrm{Lin}\, Q$ has an orbit-finite basis. Not every orbit-finitely spanned automaton arises this way though, because we do not require the vector space to have an orbit-finite basis. A key step in the proof is that every orbit-finitely spanned automaton is equivalent to one whose state space has a basis. A detailed argument is provided in Appendix C.   ■

One advantage of orbit-finitely spanned automata is that they can be minimized. Define a *homomorphism* between orbit-finitely spanned automata $\mathcal{A}$ and $\mathcal{B}$ to be a finitely supported

linear map from the state space of $\mathcal{A}$ to the state space of $\mathcal{B}$, which is consistent with the structure of the automata in the natural way, see [5, Sec. 6.2]. If $\mathcal{A}$ is an orbit-finitely spanned automaton that is reachable (i.e. its vector space is spanned by vectors that can be reached via input words), and $\sim$ is its syntactic congruence (defined in the same way as for weighted orbit-finite automata), then there is a well defined quotient automaton $\mathcal{A}_{/\sim}$. This construction, dating back to Schützenberger's original paper [20] and described in detail e.g. in [4, Thm. 8.4], does not depend on $\mathcal{A}$ being orbit-finitely spanned; in fact it applies to any weighted language over any alphabet. However, if $\mathcal{A}$ is orbit-finitely spanned then so is $\mathcal{A}_{/\sim}$. The quotient automaton admits a (surjective) homomorphism from every reachable orbit-finitely spanned automaton that recognizes the same weighted language as $\mathcal{A}$. This property uniquely defines $\mathcal{A}_{/\sim}$ up to isomorphism, and so it can be called *the* minimal automaton.

We finish this section with a third perspective on weighted languages, this time phrased in terms of monoids. Define an *orbit-finitely spanned monoid* to be a monoid $(M, \cdot, 1)$ where the underlying set M is an orbit-finitely spanned vector space, and the monoid operation is bi-linear (i.e. linear in each of the two coordinates). We say that a weighted language $L \colon \Sigma^* \to \mathbb{F}$ is recognized by such a monoid if the diagram

$$
\begin{array}{ccc}
 & M & \\
{}^{h}\nearrow & & \searrow^{F} \\
\Sigma^* \xrightarrow{\quad L \quad} & & \mathbb{F}
\end{array}
$$

commutes for some finitely supported monoid homomorphism h and some finitely supported linear map F.

The following result is somewhat unexpected, because in the non-weighted setting, orbit-finite automata and orbit-finite monoids do not recognize the same languages [5, Exercise 91].

THEOREM 7.4. *Orbit-finitely spanned monoids recognize the same weighted languages as weighted orbit-finite automata and orbit-finitely spanned automata.*

PROOF. From an orbit-finitely spanned monoid we can easily construct an orbit-finitely spanned automaton, with the same underlying vector space. For the converse direction, starting with a weighted orbit-finite automaton with states Q, we build a monoid out of finitely supported functions:

$$
f \colon (Q \times Q) \xrightarrow{\text{fs}} \mathbb{F}
$$

such that for every $p \in Q$ there are finitely many states $q \in Q$ such that $f(p, q) \neq 0$. By Theorem 6.7, this vector space is orbit-finitely spanned. The monoid operation is defined by:

$$
(f \cdot g)(p, q) = \sum_{r \in Q} f(p, r) \cdot g(r, q),
$$

with the sum being finite by the assumption on f. This operation is finitely supported and bilinear. The recognizing homomorphism is built using the same construction as when converting a nondeterministic automaton into a monoid.   ∎

An advantage of the monoid approach is the symmetry between reading the input word left-to-right and right-to-left. In particular, the languages recognized by orbit-finitely spanned monoids are easily seen to be closed under reversals; this is harder to see for the remaining models.

8.   Application to unambiguous automata

A classical application of weighted automata is a polynomial-time algorithm for language equivalence of unambiguous finite automata, i.e., nondeterministic automata with at most one accepting run for every input word. Two unambiguous finite automata are equivalent (i.e. they recognize the same language) if and only if they have the same number of accepting runs for every input word (since the number of accepting runs is zero or one). For every nondeterministic finite automaton, one can easily construct in polynomial time a weighted finite automaton which maps every input word to the number of accepting runs of the nondeterministic finite automaton; and therefore two unambiguous finite automata are equivalent if and only if the corresponding weighted finite automata are equivalent.

In this section, we show how this result can be lifted from finite to orbit-finite automata. Consider first the case of unambiguous orbit-finite automata which are non-guessing, in the sense that they have finitely many initial states, and for every state q and input letter σ, there are finitely many transitions of the form (q, a, p). As explained in Example 3.3, for such automata we can easily count runs using a weighted orbit-finite automaton, and thus we can solve the language equivalence problem in the same time as in Theorem 3.5. However, our techniques apply also to unambiguous orbit-finite automata without the non-guessing restriction.

THEOREM 8.1. *Assume that the atoms are either $(\mathbb{N}, =)$ or $(\mathbb{Q}, <)$. The equivalence problem for equivariant[3] unambiguous register automata[4], which are allowed to use guessing, is in* ExpTime, *and in polynomial time when the number of registers is fixed.*

This improves on previous work [16, 15] in that: (a) we allow unrestricted guessing; (b) we allow ordered atoms and not just equality atoms; and (c) we improve the previous upper bounds of 2ExpSpace for an unbounded number of registers and ExpSpace for a fixed number of registers.

---

3   As for Theorem 3.5, this theorem would also work for finitely supported automata, but the notation would become more involved.

4   We formulate the theorem using unambiguous register automata, see [16, Sect. 2] and not for general unambiguous orbit-finite automata, so that it can be more easily compared to existing results in the literature. However, the entire proof works for the general case.

The rest of this section is devoted to proving the above theorem. The main observation is that an orbit-finitely spanned automaton can count accepting runs for a nondeterministic orbit-finite automaton, as stated in the following lemma.

LEMMA 8.2. *Consider an equivariant nondeterministic orbit-finite automaton $\mathcal{A}$, which has finitely many accepting runs for every input word. There is an equivariant orbit-finitely spanned automaton $\mathcal{B}$, over the field of rational numbers, which outputs for every word the number of accepting runs of $\mathcal{A}$. Furthermore, the length of the vector space used by the automaton $\mathcal{B}$ is at most*

$$2^{poly(k)} \cdot n^{O(k)}. \tag{11}$$

where k and n are as in Theorem 3.5.

PROOF. Let Σ be the input alphabet, and let Q be the state space of $\mathcal{A}$. Without loss of generality we assume that every state can reach some accepting state; the remaining states can be eliminated from the automaton without affecting the recognized language or the numbers of accepting runs [5, Cor. 9.12].

For an input word $w \in \Sigma^*$, define its *configuration*

$$[w] \in Q \xrightarrow{\text{fs}} \mathbb{F}$$

to be the function which maps each state q to the number of runs on w that begin in an initial state and end in the state q. The configuration produces only finite numbers, because there cannot be a state q that can be reached via infinitely many runs over the same input word w; otherwise we could append some word to w and get infinitely many accepting runs. Define

$$V = \mathsf{Span}\{[w] : w \in \Sigma^*\} \subseteq Q \xrightarrow{\text{fs}} \mathbb{F}$$

to be the subspace of $Q \xrightarrow{\text{fs}} \mathbb{F}$ that is spanned by configurations. Although the definition of V uses a spanning set that is not necessarily orbit-finite (because Σ^* is not orbit-finite), the space V is orbit-finitely spanned, as an equivariant subspace of an orbit-finitely spanned vector space, see Theorems 6.5 and 6.7.

We use V as the state space of a orbit-finitely spanned automaton $\mathcal{B}$. Let us first prove the length bound (11). It is enough to prove the bound for the length of $Q \xrightarrow{\text{fs}} \mathbb{F}$, since V is an equivariant subspace of it. To this end, note that the set Q can be decomposed as a disjoint union of n single-orbit sets with dimension at most k. Since

$$\mathsf{length}((Q_1+Q_2) \xrightarrow{\text{fs}} \mathbb{F}) = \mathsf{length}(Q_1 \xrightarrow{\text{fs}} \mathbb{F}) + \mathsf{length}(Q_2 \xrightarrow{\text{fs}} \mathbb{F})$$

it is enough to show that every equivariant single-orbit set P of atom dimension at most k satisfies

$$\text{length}(P \xrightarrow{\text{fs}} \mathbb{F}) \leq 2^{poly(k)}.$$

Since $P \xrightarrow{\text{fs}} \mathbb{F}$ embeds into $\mathbb{A}^{(k)} \xrightarrow{\text{fs}} \mathbb{F}$, it is enough to show that the length of the latter space is at most exponential in k. This follows from the proof of Theorem 6.7.

We now describe the remaining structure of the orbit-finitely spanned automaton. The initial state is the configuration $[\varepsilon]$, which maps initial states to 1 and non-initial states to 0. Let us now define the transition functions. For an input letter $\sigma \in \Sigma$, define a function $\delta_\sigma : V \to V$ as follows

$$\sum_{i \in I} \alpha_i [w_i] \quad \mapsto \quad \sum_{i \in I} \alpha_i [w_i \sigma].$$

We need to justify that this is well-defined. A potential problem is that the same element of V might have several decompositions as weighted sums of configurations, and the output of δ_σ should not depend on the choice of decomposition. Consider an element of V with two decompositions:

$$\sum_{w \in W} \alpha_w [w] = \sum_{w \in W} \beta_w [w],$$

for some finite set $W \subseteq \Sigma^*$ of input words, and some coefficients $\alpha_w, \beta_w \in \mathbb{F}$. We need to show that δ_σ produces the same output for both decomposition, i.e.

$$\sum_{w \in W} \alpha_w [w\sigma] = \sum_{w \in W} \beta_w [w\sigma]. \tag{12}$$

Both sides in (12) are functions from Q to $\mathbb{F}$, so to prove the equality we need to show that both sides give the same output for every state $q \in Q$. For a fixed q, let P be the set of states $p \in Q$ such that the automaton has a transition (p, σ, q), and furthermore p appears in the configuration of some $w \in W$ with nonzero coefficient. An important observation is that P is a finite set: because the automaton $\mathcal{A}$ has finitely many accepting runs for every input word, the set P contains finitely many states for every word in W. For every word $w \in W$ we have

$$[w\sigma](q) = \sum_{p \in P} [w](p).$$

Therefore, to prove (12), we need to show

$$\sum_{\substack{w \in W \\ p \in P}} \alpha_w [w](p) = \sum_{\substack{w \in W \\ p \in P}} \beta_w [w](p).$$

This equality is indeed true, because our assumption implies a stronger equality, namely that for every $p \in P$ we have

$$\sum_{w \in W} \alpha_w[w](p) = \sum_{w \in W} \beta_w[w](p).$$

The function δ_σ is clearly a linear map, and thus we can set the transition function of the automaton to be $\delta(v, \sigma) = \delta_\sigma(v)$. The automaton is defined so that after reading an input word w, its state is $[w]$. The final map simply takes a configuration to the sum of all coefficients that are accepting states. This concludes the proof of Lemma 8.2. ∎

PROOF OF THEOREM 8.1. Consider two unambiguous register automata $\mathcal{A}_1$ and $\mathcal{A}_2$ for which we want to decide equivalence. Apply Lemma 8.2 to each one of them, yielding orbit-finitely spanned automata $\mathcal{B}_1$ and $\mathcal{B}_2$. Using a product construction, we get another orbit-finitely spanned automaton $\mathcal{B}$ that outputs zero for words where the automata $\mathcal{A}_1$ and $\mathcal{A}_2$ agree, and nonzero for other words. The length of the vector space in $\mathcal{B}$ is at most twice the length of the vector spaces in $\mathcal{B}_1$ and $\mathcal{B}_2$, and hence it is at most

$$2^{poly(k)} \cdot n^{O(k)} \tag{13}$$

where k is the maximal number of registers used by the automata, and n is the sum of the numbers of control states. As in the proof of Theorem 3.5, we conclude that the automata $\mathcal{A}_1$ and $\mathcal{A}_2$ are equivalent if and only if they are equivalent using input words and runs that use a number of atoms as bounded by (13), and the latter equivalence can be tested using Schützenberger's polynomial time algorithm for equivalence on weighted finite automata. ∎

References

[1] **Borja Balle and Mehryar Mohri.** Learning weighted automata. *International Conference on Algebraic Informatics*, pages 1–21. Springer, 2015. DOI (2)

[2] **Corentin Barloy and Lorenzo Clemente.** Bidimensional linear recursive sequences and universality of unambiguous register automata. *38th International Symposium on Theoretical Aspects of Computer Science (STACS)*, volume 187 of *LIPIcs*, 8:1–8:15. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2021. DOI (2, 3)

[3] **Andreas Blass.** Existence of bases implies the axiom of choice. *Contemporary Mathematics*, 31:31–33, 1984. (23)

[4] **Mikołaj Bojańczyk.** An automata toolbox, 2024. Draft from February 6, 2024. URL (31)

[5] **Mikołaj Bojańczyk.** Slightly infinite sets, 2019. Draft from September 11, 2019. URL (3–7, 14, 28, 31, 33, 40)

[6] **Mikołaj Bojańczyk, Bartek Klin, and Sławomir Lasota.** Automata theory in nominal sets. *Log. Methods Comput. Sci.* 10(3), 2014. DOI (6, 11, 15)

[7] **Mikołaj Bojańczyk, Bartek Klin, and Joshua Moerman.** Orbit-finite-dimensional vector spaces and weighted register automata. *36th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS)*, 2021. DOI (1, 4)

[8] **Mikołaj Bojańczyk and Rafal Stefanski.** Single use register automata for data words. *CoRR*, abs/1907.10504, 2019. DOI (40, 41)

[9] **Thomas Colcombet**. Unambiguity in automata theory. *Descriptional Complexity of Formal Systems (DCFS)*, volume 9118 of *Lecture Notes in Computer Science*, pages 3–18. Springer, 2015. DOI (2)

[10] **Arka Ghosh**, **Piotr Hofman**, and **Slawomir Lasota**. Solvability of orbit-finite systems of linear equations. *37th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS)*, 2022. DOI (22, 23)

[11] **K. R. Goodearl and R. B. Warfield Jr**. An Introduction to Noncommutative Noetherian Rings. London Mathematical Society Student Texts. Cambridge University Press, 2nd edition, 2004. DOI (11)

[12] **Michael Kaminski and Nissim Francez**. Finite-Memory Automata. *Theor. Comput. Sci.* 134(2):329–363, 1994. DOI (2, 8)

[13] **Stefan Kiefer**, **Andrzej S. Murawski**, **Joël Ouaknine**, **Björn Wachter**, and **James Worrell**. On the complexity of equivalence and minimisation for q-weighted automata. *Log. Methods Comput. Sci.* 9(1), 2013. DOI (8)

[14] **Daniel Krob**. The equality problem for rational series with multiplicities in the tropical semiring is undecidable. *19th International Colloquium on Automata, Languages, and Programming (ICALP)*, pages 101–112, Berlin, Heidelberg. Springer Berlin Heidelberg, 1992. DOI (2)

[15] **Antoine Mottet and Karin Quaas**. On the containment problem for unambiguous single-register automata with guessing. *CoRR*, abs/1905.12445, 2019. DOI (2, 3, 32)

[16] **Antoine Mottet and Karin Quaas**. The containment problem for unambiguous register automata and unambiguous timed automata. *Theory of Computing Systems*, 65:706–735, 2021. DOI (2, 3, 32)

[17] **Masakazu Nasu and Namio Honda**. Mappings induced by PGSM-mappings and some recursively unsolvable problems of finite probabilistic automata. *Information and Control*, 15(3):250–273, 1969. DOI (2)

[18] **Frank Neven**, **Thomas Schwentick**, and **Victor Vianu**. Finite state machines for strings over infinite alphabets. *ACM Trans. Comput. Log.* 5(3):403–435, 2004. DOI (2)

[19] **Andrew M. Pitts**. Nominal Sets: Names and Symmetry in Computer Science, volume 57 of *Cambridge Tracts in Theoretical Computer Science*. Cambridge University Press, 2013. DOI (3, 5)

[20] **Marcel Paul Schützenberger**. On the definition of a family of automata. *Information and control*, 4(2–3):245–270, 1961. DOI (1, 2, 31)

[21] **David Turner and Glynn Winskel**. Nominal domain theory for concurrency. *18th EACSL Annual Conference on Computer Science Logic (CSL)*, volume 5771 of *Lecture Notes in Computer Science*, pages 546–560, 2009. DOI (15)

A.   Proof of Lemma 4.3

For item (i), let $g : W \to W/V$ be the (equivariant) quotient map. For the "$\geq$" direction, for any chains:

$$V_1 \subsetneq \cdots \subsetneq V_n \subsetneq V \qquad\qquad (n \text{ proper inclusions})$$

$$U_1 \subsetneq \cdots \subsetneq U_m \subsetneq W/V \qquad\qquad (m \text{ proper inclusions})$$

the chain

$$V_1 \subsetneq \cdots \subsetneq V_n \subsetneq V \subseteq g^{-1}(U_1) \subsetneq \cdots \subsetneq g^{-1}(U_m) \subsetneq W$$

has $n + m$ proper inclusions. For the "$\leq$" direction, for any chain of proper inclusions:

$$W_1 \subsetneq \cdots \subsetneq W_n \subsetneq W \qquad\qquad (14)$$

consider chains:

$$W_1 \cap V \subseteq \cdots \subseteq W_n \cap V \subseteq V \qquad\qquad (15)$$

$$\overrightarrow{g}(W_0) \subseteq \cdots \subseteq \overrightarrow{g}(W_n) \subseteq W/V. \qquad\qquad (16)$$

All these inclusions are not necessarily proper. However, if the i-th inclusion in the first chain is an equality:

$$W_i \cap V = W_{i+1} \cap V$$

then there is some $v \in W_{i+1} \setminus W_i$ such that $v \notin V$. Since V is the kernel of g, the i-th inclusion in the second chain is then proper:

$$\overrightarrow{g}(W_i) \subsetneq \overrightarrow{g}(W_{i+1}),$$

so the length of (14) does not exceed the sum of lengths of (15) and (16).

For item (ii), apply item (i) to $W = \mathrm{Lin}(P + Q)$ and $V = \mathrm{Lin}\,P$, noting that then $W/V = \mathrm{Lin}\,Q$.

For item (iii), assume an equivariant surjection $q : P \to Q$, let $\bar{q} : \mathrm{Lin}\,P \to \mathrm{Lin}\,Q$ be the unique linear extension of q, and apply item (i) to $W = P$ and $V = \ker(\bar{q})$, noting that then $W/V = \mathrm{Lin}\,Q$.

B.   Proof of Lemma 4.15

We shall proceed by induction on n. In the process we will need to consider a slightly more general notion: an *affine subspace* of $\mathbb{A}$ is a set of the form $A = v \oplus V$, for an atom $v \in \mathbb{A}$ and a vector subspace $V \subseteq \mathbb{A}$. This is called n-dimensional if V is.

For any $a \in \mathbb{D}$, let $\mathbb{A}_0^a$ be the set of all those atoms that do not contain a. This, of course, is a vector subspace of $\mathbb{A}$. Similarly, $\mathbb{A}_1^a$ is the set of all those atoms that do contain a. This is not a

vector subspace, but it is an affine subspace of $\mathbb{A}$; indeed $\mathbb{A}_1^a = v \oplus \mathbb{A}_0^a$ for any $v \in \mathbb{A}_1^a$. Obviously $\mathbb{A} = \mathbb{A}_0^a \uplus \mathbb{A}_1^a$.

LEMMA B.1. *For any n-dimensional affine space A, a generator $a \in \mathbb{D}$ and $i = \{0, 1\}$, the intersection $\mathbb{A}_i^a \cap A$:*

— *is empty, or*

— *is equal to A, or*

— *is an $(n - 1)$-dimensional affine space.*

PROOF. First, consider the special case where A is a vector space, $A = V$. Then, for $i = 0$, $\mathbb{A}_0^a$ is a vector space, hence $\mathbb{A}_0^a \cap V$ is a vector space as well, and either it is equal to V or it is $(n - 1)$-dimensional. For $i = 1$, assume $\mathbb{A}_1^a \cap V$ is nonempty and pick some $v \in \mathbb{A}_1^a \cap V$. Then $\mathbb{A}_1^a = v \oplus \mathbb{A}_0^a$ and, since V is a vector space, $V = v \oplus V$. Then:

$$\mathbb{A}_1^a \cap V = (v \oplus \mathbb{A}_0^a) \cap (v \oplus V) = v \oplus (\mathbb{A}_0^a \cap V).$$

where the second equality is by Lemma 4.13(ii). The intersection on the right is a vector space; it cannot be equal to V (because v is not in it), so it must be $(n - 1)$-dimensional.

Now consider the general case of an arbitrary affine space A. Let $A = v \oplus V$ for some atom v and an n-dimensional vector space V. For $i = 0$:

$$\mathbb{A}_0^a \cap (v \oplus V) = v \oplus (\mathbb{A}_{v(a)}^a \cap V),$$

where $v(a) \in \{0, 1\}$ is the coefficient of a in v. Indeed, for the left-to-right inclusion, take any $w \in V$ such that $(v \oplus w)(a) = 0$ or, equivalently, $v(a) = w(a)$. Then $w \in \mathbb{A}_{v(a)}^a \cap V$. For the right-to-left inclusion, do the same backwards.

Then use the special case for the vector space V and $i = v(a)$.

The argument for $i = 1$ is essentially the same. ∎

PROOF OF LEMMA 4.15. We prove a more general statement, where $A_1, A_2, \ldots, A_k$ are affine n-dimensional subspaces, and not necessarily vector subspaces. We proceed by induction on the dimension n. For the base case, $n = 1$, observe that for every finite family $A_1, A_2, \ldots, A_k$ of finite-dimensional affine subspaces of $\mathbb{A}$, the symmetric difference $A_1 \div A_2 \div \cdots \div A_k$ has even size. Indeed, each A_i is of even size, and sets of even size are closed under taking symmetric differences.

For the induction step, consider a family of n-dimensional affine spaces $A_1, \ldots, A_k$, and denote

$$X = A_1 \div \cdots \div A_k.$$

Assume that X is nonempty. Then, as we have seen above, is must contain at least two different atoms; pick some generator a on which some atoms in X differ. Then both intersections

$$\mathbb{A}_0^a \cap X \qquad \text{and} \qquad \mathbb{A}_1^a \cap X$$

are nonempty.

By Lemma 4.13(i):

$$\mathbb{A}_0^a \cap X = (\mathbb{A}_0^a \cap A_1) \div \cdots \div (\mathbb{A}_0^a \cap A_k),$$
$$\mathbb{A}_1^a \cap X = (\mathbb{A}_1^a \cap A_1) \div \cdots \div (\mathbb{A}_1^a \cap A_k).$$

By Lemma B.1, both expressions to the right are symmetric differences of $(n-1)$-dimensional affine spaces. Some care is needed when (as is allowed by Lemma B.1), $\mathbb{A}_i^a \cap A_j = A_j$. In this case, present A_j as

$$A_j = v \oplus V,$$

use Lemma 4.14 to represent V as a symmetric difference of all its $(n-1)$-dimensional subspaces, and apply Lemma 4.13(iii).

By the inductive assumption $\mathbb{A}_0^a \cap X$ is either empty or has size at least 2^{n-1}, and similarly for $\mathbb{A}_1^a \cap X$. But we have picked a so that both these intersections are nonempty, so X has size at least 2^n as required. ∎

C.   Equivalence of linear and weighted automata

We begin with the easier inclusion.

LEMMA C.1. *For every weighted orbit-finite automaton $\mathcal{W}$ there is an orbit-finitely spanned automaton $\mathcal{L}$ which recognizes the same weighted language.*

PROOF. If Q are the states of $\mathcal{W}$, then the vector space of $\mathcal{L}$ is $\operatorname{Lin} Q$, and the transition function

$$\delta : \operatorname{Lin} Q \times \mathbb{A} \to \operatorname{Lin} Q$$

is defined so that for every input letter $a \in \Sigma$ and states $q, p \in Q$ (which are basis vectors of the vector space $\operatorname{Lin} Q$), the coefficient for state p in the vector $\delta(q, a)$ is the weight of the transition (q, a, p). In other words, linear map $\delta(\_, a)$ is defined by a $Q \times Q$ matrix where the coefficient in a cell (q, p) is the weight of the transition (q, a, p). Note how condition (*) from Definition 3.1 ensures that the $\delta(q, a)$ is a well-defined vector, in the sense that it has finitely many nonzero coordinates. The initial vector of $\mathcal{L}$ consists of the initial weights in $\mathcal{W}$, and the final function is the linear extension of the original final weight function. ∎

EXAMPLE C.2. Recall the weighted orbit-finite automaton from Example 7.1. We have already described the state space of the corresponding orbit-finitely spanned automaton:

$$V = \mathrm{Lin}(\{\perp, \top\} + \mathbb{A}) \oplus X,$$

and we will now define the rest of it. We define the transition function on the generators of the state space:

$$\delta(\perp, a) = a \quad \delta(\top, a) = 0 \quad \delta(a, b) = (a, b) - (b, a)$$

$$\delta((a, b) - (b, a), c) = \begin{cases} \top & \text{if } c = a, \\ -\top & \text{if } c = b, \\ 0 & \text{otherwise.} \end{cases}$$

Formally, one has to show that δ is well-defined on X, i.e. it satisfies

$$\delta((a, b) - (b, a), c) = -\delta((b, a) - (a, b), c).$$

Similarly for the final function, we define it for generators:

$$F(\perp) = 0 \quad F(\top) = 1 \quad F(a) = 0 \quad F((a, b) - (b, a)) = 0.$$

The initial vector is simply $v_0 = \perp$. This automaton accepts the same language L and is in fact minimal. ◆

Call an orbit-finitely spanned automaton *basic* if its state vector space has a basis. For a basic automaton an equivalent weighted orbit-finite automaton can be easily produced, by using the basis as the states. Therefore, to complete the proof of Theorem 7.3, we prove the following:

LEMMA C.3. *For every orbit-finitely spanned automaton, there is a basic one that recognizes the same weighted language.*

PROOF. Consider an orbit-finitely spanned automaton $\mathcal{A}$ where the state space V is spanned by an orbit-finite set Q. Define a *polynomial orbit-finite set*, see [8, Definition 1], to be any set which is a finite disjoint union of sets of the form $\mathbb{A}^k$. As for every orbit-finite set, there exists a polynomial orbit-finite set P with a surjective finitely supported function from P to Q. Extend this function to surjective a linear map

$$h : \mathrm{Lin}\, P \to V.$$

We will define a orbit-finitely spanned automaton $\mathcal{B}$ with the state space $\mathrm{Lin}\, P$ so that h becomes a homomorphism of orbit-finitely spanned automata, that is: a finitely supported linear map between the underlying vector spaces, which is consistent with the initial states, transition functions and final functions in the expected way, see [5, Sec. 6.2]. If two orbit-finitely spanned

automata are connected by a homomorphism, then they recognize the same weighted language. Therefore, to prove the lemma it remains to define the initial state, transition function and final function in $\mathcal{B}$ so that h is a homomorphism.

For the initial state in $\mathcal{B}$ we choose some vector that is mapped by h to the initial state of $\mathcal{A}$, and for the final function we use the composition of h and the final function of $\mathcal{A}$. To define the transition function, we construct a finitely supported function γ which makes the following diagram commute:

$$
\begin{array}{ccc}
P \times \Sigma & \xrightarrow{\ \ \gamma\ \ } & \mathsf{Lin}\,P \\
{\scriptstyle (h,id)}\downarrow & & \downarrow{\scriptstyle h} \\
V \times \Sigma & \xrightarrow[\text{transition function of } \mathcal{A}]{} & V
\end{array}
$$

To this end, consider the composition of the following relations: the function (h, id), the transition function of $\mathcal{A}$, and the inverse of h. This is a finitely supported binary relation

$$R \subseteq (P \times \Sigma) \times \mathsf{Lin}\,P$$

such that every element of $P \times \Sigma$ is related with at least one element of $\mathsf{Lin}\,P$ (thanks to the surjectivity of h). By the Uniformization Lemma from [8, Lemma 20], there exists a finitely supported function γ which is contained in R, thus proving the lemma. It is worth pointing out that the Uniformization Lemma changes supports: even if R is equivariant, it could be the case that γ needs non-empty support.

Using linearity, γ extends to a finitely supported function

$$\bar{\gamma} : \mathsf{Lin}\,P \times \Sigma \to \mathsf{Lin}\,P$$

which is a linear map for every fixed input letter; the resulting function can then be used as the transition function for $\mathcal{B}$. The commuting diagram above ensures that h is a homomorphism of automata. ∎