

Modelling Infinite Structures with Atoms

Mikołaj Bojańczyk*

University of Warsaw

Abstract. The topic is a variant of first-order logic defined in the framework of sets with atoms, which allows formulas to use orbit-finite boolean operations. The main contribution is a notion of model for this logic, which admits the compactness theorem.

1 Introduction

This paper studies a variant of first-order logic in sets with atoms (also known as nominal sets, also known as Fraenkel-Mostowski sets, also known as permutation models). The logic was introduced in [3], its intention is to describe properties of objects such as data words, data trees, or data graphs. The focus of [3] was the study of finite models, i.e. the logic made most sense for finite data words, finite data trees, and finite data graphs. When studying infinite objects, such as infinite data words, the approach from [3] runs against difficulties, e.g. the compactness theorem does not hold. The goal of this paper is a more thorough model-theoretic study of that logic, with a focus on the problems with infinite structures. This paper proposes a different semantics of first-order logic than the one in [3], which admits compactness, and has a sound and complete proof system.

Apart from the technical results on the logic, this paper discusses how to model and how not to model objects that talk about infinitely many atoms, while respecting the finite support constraint in the definition of sets with atoms. The running example in this paper is infinite data words: how can they be modelled using sets with atoms, and how logics can be used to express their properties.

2 Sets with Atoms

Sets with atoms were introduced in set theory by Fraenkel in 1922 and rediscovered for the computer science community, by Gabbay and Pitts [5]. Sets with atoms are now widely studied in the semantics community, under the name of nominal sets, see the book [9]. This paper is part of a research programme which studies a notion of finiteness which only makes sense in sets with atoms, called “orbit-finiteness”. The research programme is to see what happens to discrete mathematics when sets are replaced by sets with atoms, and finiteness is replaced by orbit-finiteness. Two examples of this research programme include the study of finite automata in [2], and the study of programming languages [4].

* Author supported by ERC Starting Grant “Sosna”.

What are sets with atoms? Sets with atoms are an extended notion of a set – such sets are allowed to contain “atoms”. The existence of atoms is postulated as an axiom. The key role in the theory is played by automorphisms of atoms. For instance, if a, b, c, d are atoms, and the atoms have no structure except for equality, then the sets

$$\{a, \{a, b, c\}, \{a, c\}\} \quad \{b, \{b, c, d\}, \{b, d\}\}$$

are equal up to automorphisms of atoms. If the atoms are real numbers, equipped with the successor relation $x = y + 1$ and linear order $x < y$, then the sets

$$\{-1, 0, 0.3\} \quad \{5.2, 6.2, 6.32\}$$

are equal up to automorphism of the atoms, but the sets

$$\{0, 2\} \quad \{5.3, 8.3\}$$

are not.

The formal definition is parametrized by a notion of atoms, which is given as a relational structure, and induces a notion of automorphism. (One can also consider atoms with function symbols, but we do not do this here.) A *set with atoms* is any set that can contain atoms or other sets with atoms, in a well-founded way¹. The key notion is the notion of a *legal* set of atoms, defined below. Suppose that X is a set with atoms. If π is an automorphism of atoms, then π can be applied to X , by renaming all atoms that appear in X , and appear in elements of X , and so on. We say that a set S of atoms is a *support* of the set X if X is invariant under every automorphism of atoms which is the identity on S . (For instance, the set of all atoms is supported by the empty set, because every automorphism maps the set to itself.) A set with atoms is called *legal* if it has some finite support, each of its elements has some finite support, and so on recursively.

Orbit-finiteness. Sets with atoms are a good abstraction for some infinite systems because they have a more relaxed notion of finiteness, which requires finitely many elements, but only up to automorphisms of atoms. More precisely, we say that a set with atoms X is *orbit-finite* if it is included in a union of finitely many single-orbit sets. A single orbit set is obtained from an element x (which may itself be a set) by applying all possible atom automorphisms.

Consider for example sets with atoms where the atoms have no structure, and therefore automorphisms are arbitrary permutations. The set of atoms itself is orbit-finite, actually has only one orbit, because every atom can be mapped to every other atom by a permutation. Likewise, the set of pairs of atoms has two elements up to permutation, namely (a, a) and (a, b) for $a \neq b$. The set of triples of atoms has five orbits: all coordinates are equal, all coordinates are distinct, and three orbits where exactly two coordinates are equal.

¹ Formally speaking, sets with atoms are defined by induction on their rank, which is an ordinal number. Sets of a given rank can contain atoms and sets of lower rank.

Homogeneous structure. A structure is called *homogeneous* if every finite partial automorphism (i.e. an isomorphism between two finite substructures) can be extended to a full automorphism. Examples of homogeneous structures include: the natural numbers with only equality (we call this the *equality atoms*), and the ordered rational numbers (we call this the *total order atoms*). We assume in this paper that the atom structure is a countable homogeneous relational structure over a finite relational vocabulary. This guarantees that the notion of orbit-finiteness is relatively well behaved, e.g. orbit-finite sets are closed under finite Cartesian products and finitely-supported subsets.

3 First-Order Logic and Flat Models

The topic of this paper is a study of first-order logic in sets with atoms. In this section, we take a first attempt at defining first-order logic and the models that it will be evaluated in. Later in the paper, we will present more sophisticated versions of both the logic and the models.

We are interested mainly in legal sets with atoms, and therefore we adopt the convention that, unless otherwise stated, all sets with atoms are implicitly assumed to be legal, such as in the following definition. The definition of vocabulary is the same as usual, only sets are required to be sets with atoms. A vocabulary is any set with atoms Σ , whose elements are called *predicates*, together with a finitely supported arity function, which maps predicates of Σ to their arities, which are natural numbers. We will write Σ_n for the set of n -ary predicates, i.e. the inverse image of n under the arity function. In this paper, we only study countable vocabularies.

Running example. [A vocabulary for describing data words] In a running example, we will talk about infinite data words. The idea is to model an ω -word (i.e. positions are the natural numbers) where each position is labelled by an atom. The vocabulary for this will contain one binary predicate $x < y$ (which will be used to order the positions), and one unary predicate $a(x)$ for every atom $a \in \mathbb{A}$ (which will be used to determine the labels of the positions). Observe that this vocabulary is infinite, but orbit-finite. The predicate $x < y$ is a singleton orbit, while all the predicates $a(x)$ are in the same orbit. \square

3.1 Flat Models and Stand First-Order Logic

Flat models. A *flat model* is obtained by taking the standard definition of a model and interpreting it in sets with atoms. As we shall later see, this notion of model does not behave as well as the standard notion of model.

A *flat model* \mathfrak{A} over a vocabulary Σ consists of

- A set with atoms A called the *universe* of the model.
- For every predicate $R \in \Sigma$, an interpretation $R^{\mathfrak{A}} \subseteq A^{\text{arity}(R)}$.

A flat model is a special case of a set with atoms (it is a pair of a set and a set of interpretations functions). In this special case, legality means: a) the

universe should be finitely supported; b) for every R , the set $R^{\mathfrak{A}}$ should be finitely supported; and c) the function $R \mapsto R^{\mathfrak{A}}$ must also be finitely supported, as a function from the vocabulary to sets of tuples of elements in the universe.

However, legal flat models turn out to be so weak that we will consider possibly illegal flat models in some constructions. That is why, for flat models, we make an exception to the convention that every object is implicitly legal, and explicitly write “legal flat model” or “possibly illegal flat model”.

Running example. [Legal flat model describe only data words with finitely many atoms] Consider an infinite word over the alphabet of the atoms

$$w = a_1 a_2 \cdots \in \mathbb{A}^\omega.$$

Such a word is a function from natural numbers to atoms. In this case, legality means that only finitely many different atoms can appear in the word.

Based on the word w , we can define a flat model, denoted by \underline{w} , as follows. The universe is the natural numbers, with $<$ is interpreted as the usual order on natural numbers. For every atom $a \in \mathbb{A}$, the interpretation $a^{\underline{w}}$ is the set of positions in w that have label a . It is not difficult to see that \underline{w} is a flat model, which is legal as long as w was legal.

What would happen if we tried to define \underline{w} for a word w with infinitely many different atoms? For instance, what about a word that is an enumeration of all atoms (i.e. every atom appears on exactly one position)? For such a word, every individual predicate interpretation $a^{\underline{w}}$ is a finitely supported set, namely it contains one natural number (the number of the atom a in the enumeration). What is wrong is that the function

$$a \in \mathbb{A} \quad \mapsto \quad a^{\underline{w}} \subseteq \mathbb{N}$$

is not finitely supported. □

First-order logic. To express properties of flat models, we can use standard first-order logic. Formulas can be built out of the predicates in the vocabulary, using standard logical operations \vee , \wedge , \neg , \exists and \forall . Semantics in flat models are defined in the standard way, by induction on the structure of formulas. This definition ignores the additional structure of flat models and vocabularies, namely the possibility of acting on both by using atom automorphisms.

Running example. [First-order definable properties of data words] Suppose that a is an atom, and consider the formula

$$\forall x \exists y \quad x < y \wedge a(x).$$

When evaluated in a structure of the form \underline{w} , as defined in Example 3.1, the formula says that the atom a appears infinitely often. □

A single formula of first-order logic can only use only finitely many predicates from the vocabulary. Therefore, as far as using a single formula of the logic is concerned, it makes little sense to have a flat model specify the interpretations of

all predicates in an infinite vocabulary. The infinite vocabulary *does* make sense when we consider an infinite set of formulas, as we will do in the Section 3.2, or when we consider a variant of first order-logic that admits orbit-finite boolean operations, as we will do starting with Section 4.

3.2 Compactness Fails in Legal Flat Models

We have already identified one problem with legal flat models: they can only model data words with finitely many distinct atoms, which means they are insufficient to model interesting data words. Here we identify another, related, problem: the compactness theorem fails.

Running example. [Compactness Fails for Legal Flat Models.] Let $\varphi_{<}$ be a formula which says that $<$ is a linear order. Let Γ be the following set of formulas:

$$\{\varphi_{<}\} \cup \{\exists x a(x)\}_{a \in \mathbb{A}} \cup \{\forall x a(x) \Rightarrow \neg b(x)\}_{a \neq b \in \mathbb{A}}.$$

When seen as an infinite conjunction of formulas, the set Γ says that every atom appears exactly once as a label of a position. Observe that Γ is infinite, but orbit-finite (three orbits). We will show that every finite subset of Γ is satisfiable in a legal flat model, but Γ itself is not, and therefore compactness fails for legal flat models.

Every finite subset $\Delta \subseteq \Gamma$ is satisfiable in a finite linear order where the positions are labelled by the atoms that appear in Δ .

To prove that the whole set Γ has no legal flat model, imagine for the sake of contradiction that it does have a legal flat model \mathfrak{A} . Suppose that S is a finite support of the model. Based on the model, we define a relation $<_{\mathfrak{A}}$ on atoms as follows: $a <_{\mathfrak{A}} b$ if there are elements x, y in the universe of the model \mathfrak{A} such that

$$x \in a^{\mathfrak{A}} \quad y \in b^{\mathfrak{A}} \quad x <^{\mathfrak{A}} y.$$

The function $\mathfrak{A} \mapsto <_{\mathfrak{A}}$ is a \emptyset -supported function from legal flat models to relations on the atoms, and therefore if \mathfrak{A} is finitely supported then also $<_{\mathfrak{A}}$ is finitely supported. By the properties axiomatised in Γ , the relation $<_{\mathfrak{A}}$ must be a linear order on atoms. But a finitely supported linear order on the (equality) atoms does not exist. Indeed, if S were the support of such an order, then for atoms $a, b \notin S$, the automorphism π which swaps a with b fixes all atoms in S , and witnesses

$$a <_{\mathfrak{A}} b \quad \text{iff} \quad b = \pi(a) <_{\mathfrak{A}} a = \pi(b).$$

This proves that Γ has no legal flat model.

Observe that the reasoning above is specific to the equality atoms (because we assumed that there is some automorphism that fixes S and swaps atoms a and b). Indeed, the set Γ does have a legal flat model when the total order atoms are used. The universe is the atoms, ordered by their built-in order, and the interpretation says that each atom is labelled by itself,

$$a^{\mathfrak{A}} = \{a\},$$

which is an \emptyset -supported interpretation. The model we have just described is not a data word because the positions are densely ordered. If we extended the formula $\varphi_{<}$ to also say that every position has a successor, then we would have an example for failure of compactness in the total order atoms, because the total order atoms cannot be ordered in a finitely supported way so that every atom has a successor. \square

The failure of compactness is actually related to the weak modelling power of legal flat models. Indeed, as illustrated in the example, the compactness theorem would imply that some legal flat model represents a data word where all positions have different data values.

The statement of the compactness theorem, for which we presented a counterexample above, would be: “if every finite subset of Γ has a legal flat model, then also Γ has a legal flat model.” One could imagine a weaker form the theorem, with a stronger assumption: “if every orbit-finite subset of Γ has a legal flat model, then also Γ has a legal flat model.” The counterexample above no longer works, because the set Γ , which has no legal flat models, is already orbit-finite itself. The weaker statement is also false.

Running example. [Orbit-finite compactness also fails for legal flat models.] Let φ_{succ} be a formula which says that every element has a $<$ -successor:

$$\forall x \exists y (x < y \wedge \forall z (z \leq x \vee z \geq y)).$$

For $n \in \mathbb{N}$, define $succ_n(x, y)$ to be the formula which says that y is the n -fold successor of x . For every $n \in \mathbb{N}$, define Γ_n to be the set of formulas, which says that positions at distance n have different labels:

$$\{\forall x \forall y \text{ succ}_n(x, y) \Rightarrow \neg(a(x) \wedge a(y))\}_{a \in \Delta}.$$

Define Γ to be $\varphi_{<}$, φ_{succ} , together with the union of all Γ_n . We claim that every orbit-finite subset of Γ has a legal flat model, but Γ itself has no legal flat model.

For an orbit-finite subset $\Delta \subseteq \Gamma$, let n be the biggest number such that

$$\Gamma_n \cap \Delta \neq \emptyset.$$

We define the model for Δ to be the ultimately periodic data word, which uses $n + 1$ atoms a_1, \dots, a_{n+1} in a periodic arrangement.

To prove that Γ itself has no legal flat model, use a reasoning similar to the one in the previous example. Choose some element x_1 of the universe, and let a_1 be its (unique) label. By following successors, we get infinitely many distinct atoms a_2, a_3, \dots together with a linear order. \square

4 First-Order Logic with Orbit-Finite Boolean Operations

This section, presents another variant of first-order logic, which was introduced in [3]. The motivation behind this logic is two-fold. First, the logic is more expressive. Second, it is a natural variant of first-order logic consistent with the paradigm of replacing “finite” by “orbit-finite”. The central question of this paper is finding the right semantics for this richer logic.

Definition of the logic. Since the vocabulary is a set with atoms, it makes sense to apply atom automorphisms to predicates in the vocabulary, and therefore also to formulas built out of these predicates. Therefore it makes sense to say that a set of formulas is orbit-finite. Suppose that we extend first-order logic, as defined previously, by allowing the boolean connectives (disjunction and conjunction) to take range over a (legal) orbit-finite set of formulas. Call the resulting logic *first-order logic with orbit-finite boolean operations*.

Observe that first-order logic with orbit-finite boolean operations can be seen as a fragment of $\mathcal{L}_{\omega_1\omega}$. Recall that $\mathcal{L}_{\omega_1\omega}$ is a variant of first-order logic which allows boolean operations ranging over countable sets of formulas. Since the atoms are assumed to be countable, orbit-finite sets are also countable, and therefore our logic is a special case of $\mathcal{L}_{\omega_1\omega}$. One of the differences is that $\mathcal{L}_{\omega_1\omega}$ has uncountably many formulas, while there are only countably many formulas of first-order logic with orbit-finite boolean operations (under the assumption that the vocabulary is countable).

Example 1. Suppose that the vocabulary has one unary predicate $a(x)$ for every atom a , as in the running example. We write a formula with one free variable x , which holds in positions that satisfy at most one unary predicate of the form $a(x)$.

For an atom a , consider the set of formulas

$$\Gamma_a \stackrel{\text{def}}{=} \{-b(x) : b \neq a\}.$$

This set is legal (supported by a), and also orbit-finite (it is included in the single-orbit set $\{a(x) : a \in \mathbb{A}\}$). Therefore, we are allowed to form a new formula by taking the disjunction over Γ_a . We adopt a notational convention where this formula is denoted by

$$\varphi_a \stackrel{\text{def}}{=} \bigwedge_{b \in \mathbb{A} - \{a\}} \neg b(x).$$

Generally speaking, the notational convention says that if I is an orbit-finite set and $i \mapsto \varphi_i$ is a finitely supported function from I to already defined formulas, then $\bigwedge_{i \in I} \varphi_i$ denotes the formula $\bigwedge \{\varphi_i : i \in I\}$.

To the formula φ_a , we can also apply an atom automorphism, by applying the automorphism to every conjunct in parallel. It is not difficult to see that if an atom automorphism maps the atom a to the atom a' , then it maps the formula φ_a to the formula $\varphi_{a'}$. It follows that the set

$$\Gamma \stackrel{\text{def}}{=} \{\varphi_a : a \in \mathbb{A}\}$$

is itself orbit-finite, namely it is one orbit under the action of automorphisms. Therefore, we are allowed to form a new formula by taking the disjunction over this set. This formula, according to our notational convention, is denoted by

$$\bigvee_{a \in \mathbb{A}} \bigwedge_{b \in \mathbb{A} - \{a\}} \neg b(x).$$

It says that position x satisfies at most one predicate of the form $a(x)$.

Semantics. As mentioned above, the central question of this paper is about the “right” semantics of first-order logic with orbit-finite boolean operations.

One idea for the semantics is to use flat models. A formula of first-order logic with orbit-finite boolean operations can be evaluated in a flat model, be it legal or not. The definition is standard, and can be seen as a special case of logics with infinitary boolean operations: $\bigvee_{i \in I} \varphi_i$ is true in a flat model (under an appropriate evaluation of the free variables) if and only if φ_i is true for some i ; likewise for \bigwedge .

Running example. Recall the set of formulas Γ from Example 3.2, which said that the universe is totally ordered, every atom a labels some position, and every position is labelled by at most one atom. In the example, we showed that there is no legal flat model which satisfies all formulas from Γ . Since Γ is orbit-finite, its conjunction $\bigwedge \Gamma$ is a well-formed formula of first-order logic with orbit-finite boolean operations. This formula is not satisfied in any legal flat model. On the other hand, the formula can be satisfied in illegal flat models, e.g. a model of the form \underline{w} where w is an enumeration of all atoms. \square

So far, we have two candidates for the semantics: satisfiability in legal flat models, and satisfiability in possibly illegal flat models. The first definition fails the compactness theorem, while the second definition uses notions that are not allowed in sets with atoms (as we shall later see, the second definition does admit the compactness theorem). The rest of this paper presents arguments in favour of the second definition. Section 5, presents a proof system which is sound and complete for the second definition. Section 6 defines stratified models, which are legal objects that satisfy the same formulas as possibly illegal flat models.

Differences with [3]. The syntax of first-order logic with orbit-finite boolean operations and its semantics are the same as in [3], except that [3] made the following restrictions, which are lifted in this paper:

- In [3], all elements in the universe of a relational structure are required to have empty support, which means that if x is an element of the universe and π is an atom automorphism, then $\pi(x) = x$.
- In [3], there is a finite partition $\Sigma = \Sigma_1 \cup \dots \cup \Sigma_n$ of the vocabulary vocabulary. For every block Σ_i all predicates in Σ_i have the same arity, call it n_i , and a given n_i -tuple can satisfy at most one predicate from Σ_i .

5 A Proof System

Consider sequent calculus, with the only difference being that the rules for Boolean operations are adapted to cover the orbit-finite arity of the operations. The rule for introducing \vee becomes

$$\frac{\Gamma \vdash \varphi_i, \Delta \quad \text{for some } i \in I}{\Gamma \vdash \bigvee_{i \in I} \varphi_i, \Delta} \text{ (introduce } \vee \text{),}$$

while the rule for eliminating \vee requires an orbit-finite set of premises

$$\frac{\Gamma, \varphi_i \vdash \psi, \Delta \quad \text{for all } i \in I}{\Gamma, \bigvee_{i \in I} \varphi_i \vdash \psi, \Delta} \quad (\text{eliminate } \vee).$$

A proof is defined to be a well-founded tree², where nodes are labelled by sequents. Every node and its children must be consistent with one of the proof rules in the standard way, with one subtle difference to be described later. The tree together with the labelling by sequents must be finitely supported, i.e. a legal set with atoms. Note that the requirement on the proof tree being finitely supported implies that if a node in the proof tree has an orbit-finite set of children proving different sequents, then the dependence of the sub-proof on the sequent must be finitely supported.

The subtle difference in the definition of proofs concerns the proof rules that require their premise to be true for “some $i \in I$ ”, namely the proof rules “introduce \vee ” and “eliminate \wedge ”. As opposed to the standard sequent calculus, a proof in our system is allowed to have some redundancy, by proving the premise for more than one $i \in I$. In other words, a more exact wording for the “introduce \vee ” rule would be

$$\frac{\Gamma \vdash \varphi_i \quad \text{for all } i \text{ in a nonempty subset of } I}{\Gamma \vdash \bigvee_{i \in I} \varphi_i} \quad (\text{introduce } \vee).$$

The need for redundancy is illustrated in the following example.

Example 2. Consider the total order atoms. This example is essentially in propositional logic, since all predicates have arity zero. The vocabulary has one zero-ary predicate $P_{a,b}$ for every pair of distinct atoms $a < b$; the vocabulary has one orbit. We prove the following sequent

$$\bigwedge_{a \in \mathbb{A}} \bigwedge_{b \in \mathbb{A}, b > a} P_{a,b} \quad \vdash \quad \bigwedge_{a \in \mathbb{A}} \bigvee_{b \in \mathbb{A}, b > a} P_{a,b}. \quad (1)$$

We use the rule for introducing \wedge , which means that the root of the proof tree for (1) has one child per atom $a \in \mathbb{A}$, containing a proof of the sequent

$$\bigwedge_{a \in \mathbb{A}} \bigwedge_{b \in \mathbb{A}, b > a} P_{a,b} \quad \vdash \quad \bigvee_{b \in \mathbb{A}, b > a} P_{a,b}. \quad (2)$$

Since the proof for (1) must be legal, the proof (2) must depend on a in a finitely supported way. (More precisely, the function that maps a to the proof of (2) must be a finitely supported function from atoms to proofs). This is why we can not choose, for every a , some unique $b \neq a$, because this can not be done in a finitely

² Without orbit-finite boolean operations, one can also use proofs which are sequences of sequents, such that each sequent follows from earlier sequents via proof rules. Such proofs will not work with orbit-finite operations, since in some cases they would require imposing a linear order on the atoms.

supported way. (This is why we use total order atoms. In the equality atoms, there is a finitely supported, but not emptily supported, function that maps every atom to some other atom. Namely, choose two atoms b_1, b_2 , and then that maps b_1 to b_2 and all other atoms to b_1 .) Therefore, the proof of (2) will have one child per $b > a$, containing a proof of

$$\bigwedge_{a \in \mathbb{A}} \bigwedge_{b > a} P_{a,b} \quad \vdash \quad P_{a,b}. \quad (3)$$

In particular, the proof for (2) has infinitely many child sub-proofs (although one orbit of them), even though just one would be enough.

The definition of a proof requires the tree to be well-founded; but there might be paths of unbounded length since the tree is not finitely branching (only orbit-finitely branching). One can show that if a sequent has a proof, then it has a proof with an orbit-finite set of nodes, and depth bounded by some natural number.

Completeness and compactness. The following theorem shows that the proof system is sound and complete with respect to semantics in (possibly illegal) flat models.

Theorem 1 (Completeness). *Let Γ be a finitely supported set of sentences of first-order logic with orbit-finite boolean operations, and φ a single sentence. There is a proof of $\Gamma \vdash \varphi$ if and only if φ is satisfied in every (possibly illegal) flat model which satisfies all sentences from Γ .*

The theorem can be proved using the Henkin constant saturation method, to be found in logic textbooks, e.g. [1]. Below we show that the logic also has compactness. Note the difference with $\mathcal{L}_{\omega_1\omega}$, which has completeness (under the assumption that Γ is countable), but does not have compactness (even for countable Γ), see [8].

Corollary 1 (Compactness). *Let Γ be a finitely supported set of sentences of first-order logic with orbit-finite boolean operations. If every orbit-finite subset of Γ is satisfied by a (possibly illegal) flat model, then also Γ is satisfied by a (possibly illegal) flat model.*

Proof (sketch). Toward a contradiction, assume that Γ is not satisfied by any (possibly illegal) flat model. Then, by the right-to-left implication in Theorem 1, there is a proof of $\Gamma \vdash \perp$. As remarked before, this proof can be chosen so that it is orbit-finite, and therefore only uses an orbit-finite number of premises from Γ . By the left-to-right implication in Theorem 1, these premises cannot be satisfied in any (possibly illegal) flat model, a contradiction with the assumption.

Example 3. In Corollary 1, the assumption is that all orbit-finite subsets of Γ are satisfied by (possibly illegal) flat models. The statement is no longer true when this assumption is weakened to finite subsets. Consider a set Γ of formulas which has one formula P_a for every atom, and also the formula $\bigvee_{a \in A} \neg P_a$. Every finite subset of Γ is satisfied by a flat model, even a legal one. The whole set is not satisfied in any (possibly illegal) flat model.


6 Stratified Models

As we have seen in Theorem 1, the “natural” proof system coincides with the semantics in possibly illegal flat models³. In this section, we present an alternative semantics, which only uses legal sets with atoms, and which is equivalent to the semantics in possibly illegal flat models, and therefore is equivalent to the proof system⁴.

The general idea is to present a possibly illegal flat model by giving a set of finitely supported local versions of the model. A local version does not contain full information about the model, but it does have sufficient information to evaluate formulas with a given support. When one needs to evaluate formulas with a bigger support, then a more precise local view can be taken. The local views need to be consistent with each other; our notion of consistency is modelled using a variant of directed sets, as presented below.

6.1 Ultimately Directed Sets

Suppose that V is a partially ordered legal set with atoms, which means that both the underlying set and the partial order are legal sets with atoms. If S is a set of atoms, then we say that $v \in V$ is S -greater than $w \in V$ if $\pi(v) \geq w$ holds for some automorphism π that fixes S . We say that V is S -directed if for every two elements of $v_1, v_2 \in V$, some element of $v \in V$ is S -greater than both of them. This definition is illustrated below:



The bigger S becomes, the more difficult it is to be S -directed, because of the requirements on π_1 and π_2 . The standard notion of directed set is recovered in the limit, when S is the set of all atoms.

Define $V|_{\geq v}$ to be V restricted to elements that are greater or equal to v . A set V is called *ultimately directed* if for every $v \in V$ and every finite set of atoms S , there is some $w \geq v$ such that $V|_{\geq w}$ is S -directed.

Example 4. Consider the equality atoms. The set \mathbb{A}^* with the prefix ordering is not \emptyset -directed, because the words aab and abb cannot be extended to words in the same orbit. If we restrict \mathbb{A}^* to words where each letter appears exactly once, call this set $\mathbb{A}^{(*)}$, then the set becomes \emptyset -directed; it even has a stronger property:

³ The dual approach would be to find a proof system that coincides with the semantics in legal flat models. It seems that this is the approach taken in [6], although for a different (and intuitionistic) logic, so the exact connection is not clear to the author.

⁴ Another idea would be to somehow relax, but not completely lift, the finite support condition in the definition of sets with atoms so that some interesting illegal flat models become legal. Such an approach is pursued in [7].

every two elements are comparable up to atom automorphisms. Furthermore, if $a_1 \cdots a_n$ is a finite word in $\mathbb{A}^{(*)}$, then $\mathbb{A}^{(*)}|_{\geq a_1 \cdots a_n}$ is $\{a_1, \dots, a_n\}$ -directed. It follows that $\mathbb{A}^{(*)}$ is ultimately directed.

Example 5. This example generalises the previous one to arbitrary atom structures, e.g. the total order atoms. Choose some non-repeating enumeration of the atoms $a_1 a_2 \cdots$. We claim that

$$V = \{\pi(a_1 \cdots a_n) : n \in \mathbb{N}, \pi \text{ is an atom automorphism.}\},$$

equipped with the prefix order, is ultimately directed. (Although the enumeration of atoms is not a legal object, the set V that it yields is.) It is not difficult to see that for every $b_1 \cdots b_n \in V$, the set $V|_{\geq b_1 \cdots b_n}$ is S -directed if and only if $S \subseteq \{b_1, \dots, b_n\}$. This implies that V itself is ultimately directed.

Actually, we did not even use the property that $a_1 a_2 \cdots$ was an enumeration of all atoms, it could even have been an enumeration of some infinite subset of the atoms. However, when the enumeration does use all the atoms, it has the following property: for every $v \in V$ and every finite set of atoms S , there is some node $w \geq v$ such w supports all the atoms in S (i.e. an automorphism of the atoms that fixes w must also fix all atoms in S). This additional property will be useful in some constructions.

A disadvantage of the previous example was that it required some enumeration of the atoms, which is an illegal object. This can be avoided by using a hack called a fake atom structure. A *fake atom structure* is defined to be an isomorphic copy of the atom structure, but where all elements of the universe have empty support. In particular, even infinite subsets of fake atom structure are legal.

Example 6. Consider the equality atoms. Recall the standard set-theoretic encoding of natural numbers, where for instance 2 is encoded by $\{\emptyset, \{\emptyset\}\}$. Each natural number is built without using atoms, and therefore the relational structure $(\mathbb{N}, =)$ is an example of a fake atom structure for the equality atoms: it is constructed inside legal sets with atoms, does not use atoms, but is isomorphic to the atom structure (a countable set of abstract atoms with equality only). Note that any isomorphism of this fake atom structure with the real atom structure would itself be an illegal object, since it would in particular induce an order on the atoms.

Example 7. Choose some fake atom structure, and let V be the set of isomorphisms between finite substructures of the fake atoms, and finite substructures of the atoms, ordered by extension. One can think of a partial isomorphism $f \in V$ as a finite tuple of atoms, indexed by finitely many elements from the universe of the fake atom structure.

Claim. For every $f \in V$ with range S , $V|_{\geq f}$ is S -directed.

The claim implies that V is ultimately directed. This is because for every finite set of atoms S , every element of V can be extended so that its range includes S .

6.2 Stratified Models

We now propose our definition of models which will turn out to be equivalent to (possibly illegal) flat models. Let S be a finite set of atoms. An S -stratified model over a vocabulary Σ consists of

- **The frame.** An partially ordered set with atoms V which is called the *frame*. The frame must be countable, S -directed and ultimately directed, and have a least element called the *root*. Elements of the frame are called nodes.
- **The local models.** A labelling which maps each node v of the frame to a legal flat model $\mathfrak{A}(v)$, called the *local model in v* , whose vocabulary is a subset of Σ . This labelling must satisfy the following conditions:
 1. The labelling function $v \mapsto \mathfrak{A}(v)$ is *weakly S -supported*, which means that it can be extended to some S -supported function.
 2. If $w \geq v$, then the local model in w extends the local model in v in the following sense. The universe of $\mathfrak{A}(v)$ is a subset of the universe of $\mathfrak{A}(w)$, and the interpretation of $\mathfrak{A}(v)$ is obtained from the interpretation of $\mathfrak{A}(w)$ by restricting to the vocabulary and universe of $\mathfrak{A}(v)$.
 3. For every predicate R in Σ there is a node $w \geq v$ such that R is in the vocabulary of the local model in w .

If v is a node in the frame of an S -stratified model \mathfrak{A} , then the *submodel* in node v , denoted by $\mathfrak{A}|_{\geq v}$, is defined by restricting the frame to nodes greater or equal to v , and keeping the same labelling by local models. Submodels of an S -stratified model are also S -stratified.

Semantics in stratified models. We now define the semantics of first-order logic with orbit-finite boolean operations in stratified models. Suppose that \mathfrak{A} is an S -stratified model and φ is an S -supported formula. If φ has free variables, then we also need a valuation ν , which maps the free variables of the formula to S -supported elements in the universe of the local model in the root of \mathfrak{A} . Given these ingredients, we can define the truth value, which is denoted by a triple turnstile

$$\mathfrak{A}, \nu \vDash \varphi.$$

We underline that the above is only defined when both the formula φ and the valuation ν are supported by the parameter S such that \mathfrak{A} is an S -stratified model. The definition is by induction on the structure of φ .

- For a predicate formula, define

$$\mathfrak{A}, \nu \vDash R(x_1, \dots, x_n) \tag{4}$$

to be true if there is some node w in the frame such that the local model in w has R in its vocabulary and $R(x_1, \dots, x_n)$ holds in that local model, under the valuation ν . Using the assumptions that the frame is S -directed, and that the labelling by local models is weakly S -supported, one can show that this definition does not depend on the choice of v , i.e. if it $R(x_1, \dots, x_n)$ holds in some local model whose vocabulary contains R , then it holds in all local models whose vocabulary contains R .

- Negation is defined in the standard way.
- For an orbit-finite disjunction, define

$$\mathfrak{A}, \nu \models \bigvee_{i \in I} \varphi_i$$

to be true if there is some finite set of atoms $T \supseteq S$, a submodel \mathfrak{B} of \mathfrak{A} that is T -stratified, and some $i \in I$ such that φ_i is T -supported and

$$\mathfrak{B}, \nu \models \varphi_i.$$

The definition for conjunction is obtained by using De Morgan's law.

- For an existential quantification, we define

$$\mathfrak{A}, \nu \models \exists x. \varphi(x)$$

to be true if there is some finite set of atoms $T \supseteq S$, a submodel \mathfrak{B} of \mathfrak{A} that is T -stratified, and some T -supported element a in the universe of the local model in the root of \mathfrak{B} such that

$$\mathfrak{B}, \nu[x \rightarrow a] \models \varphi(x)$$

The definition for \forall is obtained by using De Morgan's law.

Running example. [A stratified model describing a data word] We will describe a \emptyset -stratified model, call it \mathfrak{A} , that describes a data word (for equality atoms) where all letters are different. The frame is non-repeating sequences of atoms, which was shown in Example 4 to be ultimately directed. In a node $a_0 \cdots a_n$ of the frame, the local model looks as follows. The universe is always the natural numbers, regardless of the node (the universe is not built using atoms, i.e. all of its elements have empty support). The vocabulary contains the order predicate $<$, regardless of the node, which is interpreted in the standard way. Furthermore, the vocabulary contains the predicates $a_0(x), \dots, a_n(x)$, which are interpreted so that predicate $a_i(x)$ holds exactly in position i , and nowhere else. In particular, each local model has only finitely many positions satisfying the unary predicates.

Recall the property “ $<$ is a linear order and every atom appears in exactly one position”, which was studied previously in the running example, and which is expressed by the sentence

$$\varphi_{<} \wedge \bigwedge_{a \in \mathbb{A}} \exists x a(x) \wedge \bigwedge_{a \neq b \in \mathbb{A}} \forall x a(x) \Rightarrow \neg b(x).$$

This property was false in all legal flat models, but it is true in the stratified model \mathfrak{A} we have just described. To illustrate the semantics, we show that

$$\mathfrak{A} \models \bigwedge_{a \in \mathbb{A}} \exists x a(x).$$

(There are no free variables, so there is no valuation.) Recall that \models is not defined for any choice of model and formula: the formula needs to be supported by a set S such that the model is S -stratified. In this case, the formula has empty support and the model is \emptyset -stratified, so \models is defined. By unravelling the semantics of \bigwedge , we need to show that for every atom $a \in \mathbb{A}$, and every node $a_1 \cdots a_n$ in the model, if the submodel $\mathfrak{M}|_{\geq a_0 \cdots a_n}$ is S -stratified for some $S \ni a$, then that submodel satisfies $\exists x a(x)$. It is not difficult to see that the submodel in node $a_0 \cdots a_n$ is S -stratified for some $S \ni a$ if and only if $a \in \{a_0 \cdots a_n\}$. Such submodels satisfy $\exists x a(x)$. Observe that even when a is fixed, then the position x that satisfies $a(x)$ will depend on the choice of the node $a_0 \cdots a_n$. For instance, if the node is chosen to be abc then x will be the position zero, and if the node is chosen to be cba then x will be position two. \square

Running example. [An “isomorphic” model] Choose some bijection f of the natural numbers. Consider a \emptyset -stratified model that is defined the same way as in the previous example, with the difference that in node $a_0 \cdots a_n$, the predicate $a_i(x)$ holds exactly in position $f(i)$. One can show that this model satisfies the same sentences as the previous one. \square

Stratified models are equivalent to (possibly illegal) flat models. The following theorem shows stratified models are equivalent to (possibly illegal) flat models that (possibly illegal) flat models, at least when restricted to uniformly supported sets of sentences. A set is called *uniformly supported* if there is some support S that supports all elements of the set. Every finite set of sentences is uniformly supported.

Theorem 2. *Let Γ be a uniformly finitely supported set of sentences of first-order logic with orbit-finite boolean operations. Then Γ has a (possibly illegal) flat model if and only if it has a stratified model.*

The following proposition shows the left-to-right implication in Theorem 2.

Proposition 1. *Let S be a finite support.*

1. *For every (possibly illegal) flat model there is an S -stratified model which satisfies the same S -supported sentences.*
2. *For every S -stratified model there is a (possibly illegal) flat model which satisfies the same S -supported sentences.*

In the proof of the first item in the proposition, we construct a stratified model satisfying the following properties: all local models have the same universe; all elements of this universe have empty support; the frame is obtained from Example 7. Therefore, Theorem 2 would still be true under a definition of stratified model that would require these properties.

Acknowledgement. I would like to thank Nathanaël Fijalkow, Bartosz Klin, Sawomir Lasota, and especially Szymon Toruńczyk for inspiring discussions.

References

1. Adamowicz, Z., Zbierski, P.: *Logic of Mathematics: A Modern Course of Classical Logic*. Pure and Applied Mathematics: A Wiley Series of Texts, Monographs and Tracts. Wiley (2011)
2. Bojanczyk, M., Klin, B., Lasota, S.: Automata with group actions. In: LICS, pp. 355–364 (2011)
3. Bojańczyk, M., Place, T.: Toward model theory with data values. In: Czumaj, A., Mehlhorn, K., Pitts, A., Wattenhofer, R. (eds.) ICALP 2012, Part II. LNCS, vol. 7392, pp. 116–127. Springer, Heidelberg (2012)
4. Bojanczyk, M., Torunczyk, S.: Imperative programming in sets with atoms. In: D’Souza, D., Kavitha, T., Radhakrishnan, J. (eds.) FSTTCS. LIPIcs, vol. 18, pp. 4–15. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik (2012)
5. Gabbay, M.J., Pitts, A.M.: A new approach to abstract syntax with variable binding. *Formal Asp. Comput.* 13(3-5), 341–363 (2002)
6. Gabbay, M.J.: Fresh logic: proof-theory and semantics for fm and nominal techniques. *J. Applied Logic* 5(2), 356–387 (2007)
7. Gabbay, M.J.: Finite and infinite support in nominal algebra and logic: nominal completeness theorems for free. *J. Symb. Log.* 77(3), 828–852 (2012)
8. Lopez-Escobar, E.G.K.: An interpolation theorem for denumerably long formulas. *Fundam. Math.* 57, 253–272 (1965)
9. Pitts, A.M.: *Nominal Sets: Names and Symmetry in Computer Science*. Cambridge Tracts in Theoretical Computer Science, vol. 57. Cambridge University Press (2013)