

Trakhtenbrot's Theorem in Coq

A Constructive Approach to Finite Model Theory

Dominik Kirst¹ , Dominique Larchey-Wendling² 

¹Saarland University, Saarland Informatics Campus, Saarbrücken, Germany

²Université de Lorraine, CNRS, LORIA, Vandœuvre-lès-Nancy, France

kirst@ps.uni-saarland.de dominique.larchey-wendling@loria.fr

Abstract. We study finite first-order satisfiability (FSAT) in the constructive setting of dependent type theory. Employing synthetic accounts of enumerability and decidability, we give a full classification of FSAT depending on the first-order signature of non-logical symbols. On the one hand, our development focuses on Trakhtenbrot's theorem, stating that FSAT is undecidable as soon as the signature contains an at least binary relation symbol. Our proof proceeds by a many-one reduction chain starting from the Post correspondence problem. On the other hand, we establish the decidability of FSAT for monadic first-order logic, i.e. where the signature only contains at most unary function and relation symbols, as well as the enumerability of FSAT for arbitrary enumerable signatures. All our results are mechanised in the framework of a growing Coq library of synthetic undecidability proofs.

1 Introduction

In the wake of the seminal discoveries concerning the undecidability of first-order logic by Turing and Church in the 1930s, a broad line of work has been pursued to characterise the border between decidable and undecidable fragments of the original decision problem. These fragments can be grouped either by syntactic restrictions controlling the allowed function and relation symbols or the quantifier prefix, or by semantic restrictions on the admitted models (see [1] for a comprehensive description).

Concerning signature restrictions, already predating the undecidability results, Löwenheim had shown in 1915 that monadic first-order logic, admitting only signatures with at most unary symbols, is decidable [12]. Therefore, the successive negative results usually presuppose non-trivial signatures containing an at least binary symbol.

Turning to semantic restrictions, Trakhtenbrot proved in 1950 that only admitting finite models, the satisfiability problem over non-trivial signatures is still undecidable [16]. Moreover, the situation is somewhat dual to the unrestricted case, since finite satisfiability (FSAT) is still enumerable while, conversely, unrestricted validity is enumerable. As a consequence, finite validity cannot be characterised by a complete finitary deduction system and, resting on finite model theory, various natural problems in database theory are undecidable.

Conventionally, Trakhtenbrot's theorem is proved by (many-one) reduction from the halting problem for Turing machines (see e.g. [1,11]). An encoding of a given Turing machine M can be given as a formula φ_M such that the models of φ_M correspond to the runs of M . Specifically, the finite models of φ_M correspond to terminating runs of M and so a decision procedure for finite satisfiability of φ_M is enough to decide whether M terminates or not.

Although this proof strategy is in principle explainable on paper, already the formal definition of Turing machines, not to mention their encoding in first-order logic, is not ideal for mechanisation in a proof assistant. So for our Coq mechanisation of Trakhtenbrot's theorem, we follow a different strategy by starting from the Post correspondence problem (PCP), a simple matching problem on strings. Similar to the conventional proof, we proceed by encoding every instance R of PCP as a formula φ_R such that R admits a solution iff φ_R has a finite model. Employing the framework of synthetic undecidability [6,9], the computability of φ_R from R is guaranteed since all functions definable in constructive type theory are computable without reference to a concrete model of computation.

Both the conventional proof and our elaboration actually produce formulas in a custom signature well-suited for the encoding of the seed decision problems. The sharper version of Trakhtenbrot's theorem, stating that a signature with at least one binary relation (or one binary function and one unary relation) is enough to turn FSAT undecidable, is in fact left as an exercise in e.g. Libkin's book [11]. However, at least in a constructive setting, this generalisation is non-trivial and led us to mechanising a chain of signature transformations eliminating and compressing function and relation symbols step by step.

Complementing the undecidability result, we further formalise that FSAT is enumerable for enumerable signatures and decidable for monadic signatures. Again, both of these standard results come with their subtleties when explored in a constructive approach of finite model theory.

In summary, the main contributions of this paper are threefold:

- we provide an axiom-free Coq mechanisation comprising a full classification of finite satisfiability with regards to the signatures allowed;¹
- we present a streamlined proof strategy for Trakhtenbrot's theorem well-suited for mechanisation and simple to explain informally, basing on PCP;
- we give a constructive account of signature transformations and the treatment of interpreted equality typically neglected in a classical development.

The rest of the paper is structured as follows. We first describe the type-theoretical framework for undecidability proofs and the representation of first-order logic in Section 2. We then outline our variant of Trakhtenbrot's theorem for a custom signature in Section 3. This is followed by a development of enough constructive finite model theory (Section 4) to conclude some decidability results (Section 5) as well as the final classification (Section 6). We end with a brief discussion of the Coq development and future work in Section 7.

¹ Downloadable from <http://www.ps.uni-saarland.de/extras/fol-trakh/> and systematically hyperlinked with the definitions and theorems in this PDF.

2 First-Order Satisfiability in Constructive Type Theory

In order to make this paper accessible to readers unfamiliar with constructive type theory, we outline the required features of Coq's underlying type theory, the synthetic treatment of computability available in constructive mathematics, some properties of finite types, as well as our representation of first-order logic.

2.1 Basics of Constructive Type Theory

We work in the framework of a constructive type theory such as the one implemented in Coq, providing a predicative hierarchy of type universes \mathbb{T} above a single impredicative universe \mathbb{P} of propositions. On type level, we have the unit type $\mathbb{1}$ with a single element $* : \mathbb{1}$, the void type \emptyset , function spaces $X \rightarrow Y$, products $X \times Y$, sums $X + Y$, dependent products $\forall x : X. F x$, and dependent sums $\{x : X \mid F x\}$. On propositional level in \mathbb{P} , these types are denoted using the usual logical notation ($\top, \perp, \rightarrow, \wedge, \vee, \forall, \exists$).

We employ the basic inductive types of Booleans ($\mathbb{B} ::= \text{tt} \mid \text{ff}$), of Peano natural numbers ($n : \mathbb{N} ::= 0 \mid \mathbb{S} n$), the option type ($\mathbb{O} X ::= \lceil x \rceil \mid \emptyset$), and lists ($l : \mathbb{L} X ::= [] \mid x :: l$). We write $|l|$ for the *length* of a list, $l + m$ for the *concatenation* of l and m , $x \in l$ for (list) *membership*, and simply $f [x_1; \dots; x_n] := [f x_1; \dots; f x_n]$ for the *map* function. We denote by X^n the type of (polymorphic) vectors of length $n : \mathbb{N}$ and by \mathbb{F}_n the finite types understood as indices $\{0, \dots, n - 1\}$. The definitions/notations for lists are shared with vectors $\mathbf{v} : X^n$. Moreover, when $i : \mathbb{F}_n$ and $x : X$, we denote by \mathbf{v}_i the i -th component of \mathbf{v} , and by $\mathbf{v}[x/i]$ the vector \mathbf{v} with i -th component updated to value x .

2.2 Synthetic (Un-)decidability

We review the main ingredients of our synthetic approach to decidability and undecidability [5,6,8,10,14], based on the computability of all functions definable in constructive type theory.² We first introduce standard notions of computability theory without referring to a formal model of computation, e.g. Turing machines.

Definition 1. A *problem* or predicate $p : X \rightarrow \mathbb{P}$ is

- *decidable* if there is $f : X \rightarrow \mathbb{B}$ with $\forall x. p x \leftrightarrow f x = \text{tt}$.
- *enumerable* if there is $f : \mathbb{N} \rightarrow \mathbb{O} X$ with $\forall x. p x \leftrightarrow \exists n. f n = \lceil x \rceil$.

These notions generalise to predicates of higher arity. Moreover, a type X is

- *enumerable* if there is $f : \mathbb{N} \rightarrow \mathbb{O} X$ with $\forall x. \exists n. f n = \lceil x \rceil$.
- *discrete* if equality on X (i.e. $\lambda x y : X. x = y$) is decidable.
- a *data type* if it is both enumerable and discrete.

² A result shown and applied for many variants of constructive type theory and which Coq designers are committed to maintain as Coq evolves.

Using the expressiveness of dependent types, we equivalently tend to establish the decidability of a predicate $p : X \rightarrow \mathbb{P}$ by giving a function $\forall x : X. p x + \neg p x$. Next, decidability and enumerability transport along many-one-reductions:

Definition 2. A problem $p : X \rightarrow \mathbb{P}$ (*many-one*) *reduces* to $q : Y \rightarrow \mathbb{P}$, written $p \preceq q$, if there is a function $f : X \rightarrow Y$ such that $p x \leftrightarrow q(f x)$ for all $x : X$.³

Fact 3. Assume $p : X \rightarrow \mathbb{P}$, $q : Y \rightarrow \mathbb{P}$ and $p \preceq q$: (1) if q is decidable, then so is p ; (2) if X and Y are data types and q is enumerable, then so is p .

Item (1) implies that we can justify the undecidability of a target problem by reduction from a seed problem known to be undecidable, such as the halting problem for Turing machines. This is in fact the closest rendering of undecidability available in a synthetic setting, since the underlying type theory is consistent with the assumption that every problem is decidable.⁴ Nevertheless, we believe that in the intended effective model for synthetic computability, a typical seed problem is indeed undecidable and so are the problems reached by verified reductions⁵. More specifically, since the usual seed problems are not co-enumerable, (2) implies that the reached problems are not co-enumerable either.

Given its simple inductive characterisation involving only basic types of lists and Booleans, the (binary) Post correspondence problem (BPCP) is a well-suited seed problem for compact encoding into first-order logic.

Definition 4. Given a list $R : \mathbb{L}(\mathbb{LB} \times \mathbb{LB})$ of pairs s/t of Boolean strings⁶, we define *derivability* of a pair s/t from R (denoted by $R \triangleright s/t$) and *solvability* (denoted by $\text{BPCP } R$) by the following rules:

$$\frac{s/t \in R}{R \triangleright s/t} \quad \frac{s/t \in R \quad R \triangleright u/v}{R \triangleright (s \# u)/(t \# v)} \quad \frac{R \triangleright s/s}{\text{BPCP } R}$$

Fact 5. Given a list $R : \mathbb{L}(\mathbb{LB} \times \mathbb{LB})$, the derivability predicate $\lambda s t. R \triangleright s/t$ is decidable. However, the halting problem for Turing machines reduces to BPCP.

Proof. We give of proof of the decidability of $R \triangleright s/t$ by induction on $|s| + |t|$. We also provide a trivial proof of the equivalence of two definitions of BPCP. See [5,8] for details on the reduction from the halting problem to BPCP. \square

2.3 Finiteness

Definition 6. A type X is *finite* if there is a list l_X with $x \in l_X$ for all $x : X$ and a predicate $p : X \rightarrow \mathbb{P}$ is *finite* if there is a list l_p with $\forall x. p x \leftrightarrow x \in l_p$.

³ Or equivalently, the dependent characterisation $\forall x : X. \{y : Y \mid p x \leftrightarrow q y\}$.

⁴ i.e. the axiom $\forall p : \mathbb{P}. p + \neg p$, “decidable” being understood as in Definition 1.

⁵ This synthetic treatment of undecidability is discussed in more detail in [6]

⁶ Notice that the list R is viewed as a (finite) set of pairs $s/t \in R$ (hence ignoring the order or duplicates), while s and t , which are also lists, are viewed as strings (hence repetitions and ordering matters for s and t).

Note that in constructive settings there are various alternative characterisations of finiteness⁷ (bijection with \mathbb{F}_n for some n ; negated infinitude for some definition of infiniteness; etc.) and we opted for the above since it is easy to work with while transparently capturing the expected meaning. One can distinguish *strong* finiteness in \mathbb{T} (i.e. $\{l_X : \mathbb{L} X \mid \forall x. x \in l_X\}$) from *weak* finiteness in \mathbb{P} (i.e. $\exists l_X : \mathbb{L} X. \forall x. x \in l_X$), the list l_X being required computable in the strong case. Strong finiteness implies weak finiteness but the converse holds only in restricted contexts, e.g. inside proofs of statements in \mathbb{P} .

We present three important tools for manipulating finite types: the *finite pigeon hole principle* (PHP) here established without assuming discreteness, the *well-foundedness* of strict orders over finite types, and *quotients* over strongly decidable equivalences that map onto \mathbb{F}_n . The proofs are given in Appendix A.

For the finite PHP, the typical classical proof requires the discreteness of X to design transpositions/permuations. Here we avoid discreteness completely, the existence of a duplicate being established without actually computing one.

Theorem 7. (Finite PHP) *Let $R : X \rightarrow Y \rightarrow \mathbb{P}$ be a binary relation and $l : \mathbb{L} X$ and $m : \mathbb{L} Y$ be two lists where m is shorter than l ($|m| < |l|$). If R is total from l to m ($\forall x. x \in l \rightarrow \exists y. y \in m \wedge R x y$) then the values at two distinct positions in l are related to the same y in m , i.e. there exist $x_1, x_2 \in l$ and $y \in m$ such that l has shape $l = \dots + x_1 :: \dots + x_2 :: \dots$ and $R x_1 y$ and $R x_2 y$.*

Using the PHP, one can constructively show that, for a strict order over a finite type X , any descending chain has length bounded by the size⁸ of X .

Fact 8. *Every strict order on a finite type is well-founded.*

Coq's type theory does not provide quotients in general but one can build computable quotients in certain conditions, here for a decidable equivalence relation of which representatives of equivalence classes are computable into a list.

Theorem 9. (Finite decidable quotient) *Let $\sim : X \rightarrow X \rightarrow \mathbb{P}$ be a decidable equivalence with $\{l_r : \mathbb{L} X \mid \forall x \exists y. y \in l_r \wedge x \sim y\}$, i.e. finitely many equivalence classes.⁹ Then one can compute the quotient X/\sim onto \mathbb{F}_n for some n , i.e. $n : \mathbb{N}$, $c : X \rightarrow \mathbb{F}_n$ and $r : \mathbb{F}_n \rightarrow X$ s.t. $\forall p. c(r p) = p$ and $\forall x y. x \sim y \leftrightarrow cx = cy$.*

Using Theorem 9 with identity over X as equivalence, we get bijections between finite, discrete types and the type family $(\mathbb{F}_n)_{n:\mathbb{N}}$.¹⁰

Corollary 10. *If X is a finite and discrete type then one can compute $n : \mathbb{N}$ and a bijection from X to \mathbb{F}_n .*

⁷ and these alternative characterisations are not necessarily constructively equivalent.

⁸ i.e. the length of the enumerating list of X .

⁹ hence l_r denotes a list of *representatives* of equivalence classes.

¹⁰ For a given X , the value n (usually called cardinal) is unique by the PHP.

2.4 Representing First-Order Logic

We briefly outline our representation of the syntax and semantics of first-order logic in constructive type theory (c.f. [7]). Concerning the *syntax*, we describe terms and formulas as dependent inductive types over a signature $\Sigma = (\mathcal{F}_\Sigma; \mathcal{P}_\Sigma)$ of function symbols $f : \mathcal{F}_\Sigma$ and relation symbols $P : \mathcal{P}_\Sigma$ with arities $|f|$ and $|P|$, using binary connectives $\square \in \{\dot{\rightarrow}, \dot{\wedge}, \dot{\vee}\}$ and quantifiers $\dot{\nabla} \in \{\dot{\forall}, \dot{\exists}\}$

$$\begin{array}{ll} t : \text{Term}_\Sigma ::= x \mid f \mathbf{t} & (x : \mathbb{N}, f : \mathcal{F}_\Sigma, \mathbf{t} : \text{Term}_\Sigma^{|f|}) \\ \varphi, \psi : \text{Form}_\Sigma ::= \perp \mid P \mathbf{t} \mid \varphi \dot{\square} \psi \mid \dot{\nabla} \varphi & (P : \mathcal{P}_\Sigma, \mathbf{t} : \text{Term}_\Sigma^{|P|}) \end{array}$$

where the vectors \mathbf{t} are of the expected lengths $|f|$ and $|P|$, respectively. Negation is defined as the abbreviation $\dot{\neg} \varphi := \varphi \dot{\rightarrow} \perp$.

In the chosen de Bruijn representation [4], a bound variable is encoded as the number of quantifiers shadowing its binder, e.g. $\forall x. \exists y. P x u \rightarrow P y v$ may be represented by $\dot{\forall} \dot{\exists} P 1 4 \dot{\rightarrow} P 0 5$. The variables $2 = 4 - 2$ and $3 = 5 - 2$ in this example are the *free* variables, and variables that do not occur freely are called *fresh*, e.g. 0 and 1 are fresh. For the sake of legibility, we write concrete formulas with named binders and defer de Bruijn representations to the Coq development. For a formula φ over a signature Σ , we define the list $\text{FV}(\varphi) : \mathbb{L}\mathbb{N}$ of free variables, the list $\mathcal{F}_\varphi : \mathbb{L}\mathcal{F}_\Sigma$ of function symbols and the list $\mathcal{P}_\varphi : \mathbb{L}\mathcal{P}_\Sigma$ of relation symbols that actually occur in φ , all by recursion on φ .

Turning to *semantics*, we employ the standard (Tarski-style) model-theoretic semantics, evaluating terms in a given domain and embedding the logical connectives into the (constructive) meta-logic (cf. [17]):

Definition 11. A *model* \mathcal{M} over a domain $D : \mathbb{T}$ is described by a pair of functions $\forall f. D^{|f|} \rightarrow D$ and $\forall P. D^{|P|} \rightarrow \mathbb{P}$ written $f^\mathcal{M}$ and $P^\mathcal{M}$. Given a *variable assignment* $\rho : \mathbb{N} \rightarrow D$, we recursively extend it to a *term evaluation* $\hat{\rho} : \text{Term} \rightarrow D$ with $\hat{\rho} x := \rho x$ and $\hat{\rho}(f \mathbf{v}) := f^\mathcal{M}(\hat{\rho} \mathbf{v})$, and to the *satisfaction* relation $\mathcal{M} \models_\rho \varphi$ by

$$\begin{array}{ll} \mathcal{M} \models_\rho \perp := \perp & \mathcal{M} \models_\rho \varphi \dot{\square} \psi := \mathcal{M} \models_\rho \varphi \square \mathcal{M} \models_\rho \psi \\ \mathcal{M} \models_\rho P \mathbf{t} := P^\mathcal{M}(\hat{\rho} \mathbf{t}) & \mathcal{M} \models_\rho \dot{\nabla} \varphi := \nabla a : D. \mathcal{M} \models_{a \cdot \rho} \varphi \end{array}$$

where each logical connective $\dot{\square}/\dot{\nabla}$ is mapped to its meta-level counterpart \square/∇ and where we denote by $a \cdot \rho$ the de Bruijn extension of ρ by a , defined by $(a \cdot \rho) 0 := a$ and $(a \cdot \rho)(1 + x) := \rho x$.¹¹

A Σ -*model* is thus a dependent triple (D, \mathcal{M}, ρ) composed of a domain D , a model \mathcal{M} for Σ over D and an assignment $\rho : \mathbb{N} \rightarrow D$. It is *finite* if D is finite, and *decidable* if $P^\mathcal{M} : D^{|P|} \rightarrow \mathbb{P}$ is decidable for all $P : \mathcal{P}_\Sigma$.

Fact 12. Satisfaction $\lambda \varphi. \mathcal{M} \models_\rho \varphi$ is decidable for finite, decidable Σ -models.

Proof. By induction on φ ; finite quantification preserves decidability. \square

¹¹ The notation $a \cdot \rho$ illustrates that a is pushed ahead of the sequence ρ_0, ρ_1, \dots

In this paper, we are mostly concerned with *finite satisfiability* of formulas. However, since some of the compound reductions hold for more general or more specific notions, we introduce the following variants:

Definition 13. (Satisfiability) For a formula φ over a signature Σ , we write

- $\text{SAT}(\Sigma) \varphi$ if there is a Σ -model (D, \mathcal{M}, ρ) such that $\mathcal{M} \models_\rho \varphi$;
- $\text{FSAT}(\Sigma) \varphi$ if additionally D is finite and \mathcal{M} is decidable;
- $\text{FSATEQ}(\Sigma; \equiv) \varphi$ if the signature contains a distinguished binary relation symbol \equiv interpreted as equality, i.e. $x \equiv^{\mathcal{M}} y \leftrightarrow x = y$ for all $x, y : D$.

Notice that in a classical treatment of finite model theory, models are supposed to be given *in extension*, i.e. understood as tables providing computational access to functions and relations values. To enable this view in our constructive setting, we restrict to decidable relations in the definition of FSAT , and from now on, *finite satisfiability is always meant to encompass a decidable model*. One could further require the domain D to be discrete to conform more closely with the classical view; discreteness is in fact enforced by FSATEQ . However, we refrain from this requirement and instead show in Section 4.1 that FSAT and FSAT over discrete models are constructively equivalent.

3 Trakhtenbrot's Theorem (Approximated)

In this section, we show that BPCP reduces to $\text{FSATEQ}(\Sigma_{\text{BPCP}}; \equiv)$ for the special purpose signature $\Sigma_{\text{BPCP}} := (\{\star^0, e^0, f_{\text{tt}}^1, f_{\text{ff}}^1\}; \{P^2, \prec^2, \equiv^2\})$. To this end, we fix an instance $R : \mathbb{L}(\mathbb{LB} \times \mathbb{LB})$ of BPCP (to be understood as a finite set of pairs of Boolean strings) and we construct a formula φ_R such that φ_R is finitely satisfiable if and only if R has a solution.

Informally, we axiomatise a family \mathcal{B}_n of models over the domain of Boolean strings of length bounded by n and let φ_R express that R has a solution in \mathcal{B}_n . The axioms express enough equations and inversions of the constructions included in the definition of BPCP such that a solution for R can be recovered.

Formally, the symbols in Σ_{BPCP} are used as follows: the functions f_b and the constant e represent $b :: (\cdot)$ and $[]$ for the encoding of strings s as terms \bar{s} :

$$[\] \# \# \tau := \tau \quad b :: s \# \# \tau := f_b(\bar{s} \# \# \tau) \quad \bar{s} := \bar{s} \# \# e$$

The constant \star represents an undefined value for strings too long to be encoded in the finite model \mathcal{B}_n . The relation P represents derivability $R \triangleright \cdot / \cdot$ while \prec and \equiv represent strict suffixes and equality, respectively.

Expected properties of the intended interpretation can be captured formally as first-order formulas. First, we ensure that P is proper (only subject to defined values) and that \prec is a strict order (irreflexive and transitive):

$$\begin{aligned} \varphi_P &:= \forall xy. P x y \dot{\rightarrow} x \neq \star \wedge y \neq \star && (P \text{ proper}) \\ \varphi_\prec &:= (\forall x. x \not\prec x) \wedge (\forall xyz. x \prec y \dot{\rightarrow} y \prec z \dot{\rightarrow} x \prec z) && (\prec \text{ strict order}) \end{aligned}$$

Next, the image of f_b is forced disjoint from e and injective as long as \star is not reached. We also ensure that the images of f_{tt} and f_{ff} intersect only at \star :

$$\varphi_f := \left(\begin{array}{l} f_{\text{tt}} \star \equiv \star \wedge f_{\text{ff}} \star \equiv \star \\ \dot{\forall} x. f_{\text{tt}} x \not\equiv e \\ \forall x. f_{\text{ff}} x \not\equiv e \end{array} \right) \wedge \left(\begin{array}{l} \dot{\forall} xy. f_{\text{tt}} x \not\equiv \star \rightarrow f_{\text{tt}} x \equiv f_{\text{tt}} y \rightarrow x \equiv y \\ \dot{\forall} xy. f_{\text{ff}} x \not\equiv \star \rightarrow f_{\text{ff}} x \equiv f_{\text{ff}} y \rightarrow x \equiv y \\ \forall xy. f_{\text{tt}} x \equiv f_{\text{ff}} y \rightarrow f_{\text{tt}} x \equiv \star \wedge f_{\text{ff}} y \equiv \star \end{array} \right)$$

Furthermore, we enforce that P simulates $R \triangleright \cdot / \cdot$, encoding its inversion principle

$$\varphi_{\triangleright} := \dot{\forall} xy. P xy \rightarrow \dot{\bigvee}_{s/t \in R} \dot{\vee} \left\{ \begin{array}{l} x \equiv \bar{s} \wedge y \equiv \bar{t} \\ \dot{\exists} uv. P u v \wedge x \equiv \bar{s} \# u \wedge y \equiv \bar{t} \# v \wedge u/v \prec x/y \end{array} \right.$$

where $u/v \prec x/y$ is short for $(u \prec x \wedge v \equiv y) \dot{\vee} (v \prec y \wedge u \equiv x) \dot{\vee} (u \prec x \wedge v \prec y)$. Finally, φ_R is the conjunction of all axioms plus the existence of a solution:

$$\varphi_R := \varphi_P \wedge \varphi_{\prec} \wedge \varphi_f \wedge \varphi_{\triangleright} \wedge \dot{\exists} x. P xx.$$

Theorem 14. $\text{BPCP} \preceq \text{FSATEQ}(\Sigma_{\text{BPCP}}; \equiv)$.

Proof. The reduction $\lambda R. \varphi_R$ is proved correct by Lemmas 15 and 16. \square

Lemma 15. $\text{BPCP } R \rightarrow \text{FSATEQ}(\Sigma_{\text{BPCP}}; \equiv) \varphi_R$.

Proof. Assume $R \triangleright s/s$ holds for a string s with $|s| = n$. We show that the model \mathcal{B}_n over Boolean strings bounded by n satisfies φ_R . To be more precise, we choose $D_n := \mathbb{O}\{s : \mathbb{LB} \mid |s| \leq n\}$ as domain, i.e. values in D_n are either an (overflow) value \emptyset or a (defined) dependent pair $\ulcorner(s, H_s)\urcorner$ where $H_s : |s| \leq n$. We interpret the function and relation symbols of the chosen signature by

$$\begin{aligned} e^{\mathcal{B}_n} &:= [] & f_b^{\mathcal{B}_n} \emptyset &:= \emptyset & P^{\mathcal{B}_n} st &:= R \triangleright s/t \\ \star^{\mathcal{B}_n} &:= \emptyset & f_b^{\mathcal{B}_n} s &:= \text{if } |s| < n \text{ then } b :: s \text{ else } \emptyset & s \prec^{\mathcal{B}_n} t &:= s \neq t \wedge \exists u. u \# s = t \end{aligned}$$

where we left out some explicit constructors and the excluded edge cases of the relations for better readability. As required, \mathcal{B}_n interprets \equiv by equality $=_{D_n}$.

Considering the desired properties of \mathcal{B}_n , first note that D_n can be shown finite by induction on n . This however crucially relies on the proof irrelevance of the $\lambda x. x \leq n$ predicate.¹² The atoms $s \prec^{\mathcal{B}_n} t$ and $s \equiv^{\mathcal{B}_n} t$ are decidable by straightforward computations on Boolean strings. Decidability of $P^{\mathcal{B}_n} st$ (i.e. $R \triangleright s/t$) follows from Fact 5. Finally, since φ_R is a closed formula, any variable assignment ρ can be chosen to establish that \mathcal{B}_n satisfies φ_R , for instance $\rho := \lambda x. \emptyset$. Then showing $\mathcal{B}_n \models_\rho \varphi_R$ consists of verifying simple properties of the chosen functions and relations, with mostly straightforward proofs. \square

Lemma 16. $\text{FSATEQ}(\Sigma_{\text{BPCP}}; \equiv) \varphi_R \rightarrow \text{BPCP } R$.

¹² i.e. that for every $x : \mathbb{N}$ and $H, H' : x \leq n$ we have $H = H'$. In general, it is not always possible to establish finiteness of $\{x \mid P x\}$ if P is not proof irrelevant.

Proof. Suppose that $\mathcal{M} \models_{\rho} \varphi_R$ holds for some finite Σ_{BPCP} -model (D, \mathcal{M}, ρ) interpreting \equiv as equality and providing operations $f_b^{\mathcal{M}}, e^{\mathcal{M}}, \star^{\mathcal{M}}, P^{\mathcal{M}}$ and $\prec^{\mathcal{M}}$. Again, the concrete assignment ρ is irrelevant and $\mathcal{M} \models_{\rho} \varphi_R$ ensures that the functions/relations behave as specified and that $P^{\mathcal{M}} x x$ holds for some $x : D$.

Instead of trying to show that \mathcal{M} is isomorphic to some \mathcal{B}_n , we directly reconstruct a solution for R , i.e. we find some s with $R \triangleright s / s$ from the assumption that $\mathcal{M} \models_{\rho} \varphi_R$ holds. To this end, we first observe that the relation $u/v \prec^{\mathcal{M}} x/y$ as defined above is a strict order and thus well-founded as an instance of Fact 8.

Now we can show that for all x/y with $P^{\mathcal{M}} x y$ there are strings s and t with $x = \bar{s}$, $y = \bar{t}$ and $R \triangleright s/t$, by induction on the pair x/y using the well-foundedness of $\prec^{\mathcal{M}}$. So let us assume $P^{\mathcal{M}} x y$. Since \mathcal{M} satisfies φ_{\triangleright} there are two cases:

- there is $s/t \in R$ such that $x = \bar{s}$ and $y = \bar{t}$. The claim follows by $R \triangleright s/t$;
- there are $u, v : D$ with $P^{\mathcal{M}} u v$ and $s/t \in R$ such that $x = \bar{s} \# u$, $y = \bar{t} \# v$, and $u/v \prec^{\mathcal{M}} x/y$. The latter makes the inductive hypothesis applicable for $P^{\mathcal{M}} u v$, hence yielding $R \triangleright s'/t'$ for some strings s' and t' corresponding to the encodings u and v . This is enough to conclude $x = \overline{s + s'}$, $y = \overline{t + t'}$ and $R \triangleright (s + s')/(t + t')$ as wished.

Applying this to the assumed match $P^{\mathcal{M}} x x$ yields a solution $R \triangleright s/s$. \square

4 Constructive Finite Model Theory

Combined with Fact 5, Theorem 14 entails the undecidability (and non-co-enumerability) of FSATEQ over a custom (both finite and discrete) signature Σ_{BPCP} . By a series of signature reductions, we generalise these results to any signature containing an at least binary relation symbol. In particular, we explain how to reduce FSAT(Σ) to FSAT($\emptyset; \{\in^2\}$) for any discrete signature Σ , hence including Σ_{BPCP} . We also provide a reduction from FSAT($\emptyset; \{\in^2\}$) to FSAT($\{f^n\}; \{P^1\}$) for $n \geq 2$, which entails the undecidability of FSAT for signatures with one unary relation and an at least binary function. But first, let us show that FSAT is unaltered when further assuming discreteness of the domain.

4.1 Removing Model Discreteness and Interpreted Equality

We consider the case of models over a discrete domain D . Of course, in the case of FSATEQ($\Sigma; \equiv$) the requirement that \equiv is interpreted as a decidable binary relation which is equivalent to $=_D$ imposes the discreteness of D . But in the case of FSAT(Σ) nothing imposes such a restriction on D . However, as we argue here, we can always quotient D using a suitable decidable congruence, making the quotient a discrete finite type while preserving first-order satisfaction.

Definition 17. We write FSAT'(Σ) φ if FSAT(Σ) φ on a discrete model.

Let us consider a fixed signature $\Sigma = (\mathcal{F}_{\Sigma}; \mathcal{P}_{\Sigma})$. In addition, let us fix a finite type D and a (decidable) model \mathcal{M} of Σ over D . We can conceive an equivalence

over D which is a congruence for all the interpretations of the symbols by \mathcal{M} , namely *first-order indistinguishability* $x \doteq_{\Sigma} y := \forall \varphi \rho. \mathcal{M} \models_{x \cdot \rho} \varphi \leftrightarrow \mathcal{M} \models_{y \cdot \rho} \varphi$, i.e. first-order semantics in \mathcal{M} is not impacted when switching x with y .

The facts that \doteq_{Σ} is both an equivalence and a congruence are easy to prove but, with this definition, there is little hope of establishing decidability of \doteq_{Σ} . The main reason for this is that the signature may contain symbols of infinitely many arities. So we fix two lists $l_{\mathcal{F}} : \mathbb{L}\mathcal{F}_{\Sigma}$ and $l_{\mathcal{P}} : \mathbb{L}\mathcal{P}_{\Sigma}$ of function and relation symbols respectively and restrict the congruence requirement to these lists.

Definition 18. (Bounded first-order indistinguishability) We say that x and y are *first-order indistinguishable up to $l_{\mathcal{F}}/l_{\mathcal{P}}$* , and we write $x \doteq y$, if for any $\rho : \mathbb{N} \rightarrow D$ and any first-order formula φ built from the symbols in $l_{\mathcal{F}}$ and $l_{\mathcal{P}}$ only, we have $\mathcal{M} \models_{x \cdot \rho} \varphi \leftrightarrow \mathcal{M} \models_{y \cdot \rho} \varphi$.

Theorem 19. *First-order indistinguishability \doteq up to $l_{\mathcal{F}}/l_{\mathcal{P}}$ is a strongly decidable equivalence and a congruence for all the symbols in $l_{\mathcal{F}}/l_{\mathcal{P}}$.*

Proof. The proof is quite involved; see Appendix B. We characterise \doteq as a bisimulation, i.e. the greatest Kleene fixpoint of an operator which preserves decidability and of which the iterations converge after finitely many steps. \square

We use the strongly decidable congruence \doteq to quotient models onto discrete ones (in fact \mathbb{F}_n for some n) while preserving first-order satisfaction.

Theorem 20. *For every first-order signature Σ and formula φ over Σ , we have $\text{FSAT}(\Sigma) \varphi$ iff $\text{FSAT}'(\Sigma) \varphi$, and as a consequence, both reductions $\text{FSAT}(\Sigma) \preceq \text{FSAT}'(\Sigma)$ and $\text{FSAT}'(\Sigma) \preceq \text{FSAT}(\Sigma)$ hold.*

Proof. $\text{FSAT}(\Sigma) \varphi$ entails $\text{FSAT}'(\Sigma) \varphi$ is the non-trivial implication. Hence we consider a finite Σ -model (D, \mathcal{M}, ρ) of φ and we build a new finite Σ -model of φ which is furthermore discrete. We collect the symbols occurring in φ as the lists $l_{\mathcal{F}} := \mathcal{F}_{\varphi}$ (for functions) and $l_{\mathcal{P}} := \mathcal{P}_{\varphi}$ (for relations). By Theorem 19, first-order indistinguishability $\doteq : D \rightarrow D \rightarrow \mathbb{P}$ up to $\mathcal{F}_{\varphi}/\mathcal{P}_{\varphi}$ is a strongly decidable equivalence over D and a congruence for the semantics of the symbols occurring in φ . Using Theorem 9, we build the quotient D/\doteq on a \mathbb{F}_n for some $n : \mathbb{N}$. We transport the model \mathcal{M} along this quotient and because \doteq is a congruence for the symbols in φ , its semantics is preserved along the quotient. Hence, φ has a finite model over the domain \mathbb{F}_n which is both finite and discrete. \square

Theorem 21. *If \equiv is a binary relation symbol in the signature Σ , one has a reduction $\text{FSATEQ}(\Sigma; \equiv) \preceq \text{FSAT}(\Sigma)$.*

Proof. Given a list $l_{\mathcal{F}}$ (resp. $l_{\mathcal{P}}$) of function (resp. relation) symbols, we construct a formula $\psi(l_{\mathcal{F}}, l_{\mathcal{P}}, \equiv)$ over the function symbols in $l_{\mathcal{F}}$ and relation symbols in $(\equiv :: l_{\mathcal{P}})$ expressing the requirement that \equiv is an equivalence and a congruence for the symbols in $l_{\mathcal{F}}/l_{\mathcal{P}}$. Then we show that $\lambda \varphi. \varphi \wedge \psi(\mathcal{F}_{\varphi}, \equiv :: \mathcal{P}_{\varphi}, \equiv)$ is a correct reduction, where \mathcal{F}_{φ} and \mathcal{P}_{φ} list the symbols occurring in φ . \square

4.2 From Discrete Signatures to Singleton Signatures

Let us start by converting a discrete signature to a finite and discrete signature.

Lemma 22. *For any formula φ over a discrete signature Σ , one can compute a signature $\Sigma_{n,m} = (\mathbb{F}_n; \mathbb{F}_m)$, arity preserving maps $\mathbb{F}_n \rightarrow \mathcal{F}_\Sigma$ and $\mathbb{F}_m \rightarrow \mathcal{P}_\Sigma$ and an equi-satisfiable formula ψ over $\Sigma_{n,m}$, i.e. $\text{FSAT}(\Sigma) \varphi \leftrightarrow \text{FSAT}(\Sigma_{n,m}) \psi$.*

Proof. We use the discreteness of Σ and bijectively map the lists of symbols \mathcal{F}_φ and \mathcal{P}_φ onto \mathbb{F}_n and \mathbb{F}_m respectively, using Corollary 10. We structurally map φ to ψ over $\Sigma_{n,m}$ along this bijection, which preserves finite satisfiability. \square

Notice that n and m in the signature $\Sigma_{n,m}$ depend on φ , hence the above statement cannot be presented as a reduction between (fixed) signatures.

We now erase all function symbols by encoding them with relation symbols. To this end, let $\Sigma = (\mathcal{F}_\Sigma; \mathcal{P}_\Sigma)$ be a signature, we set $\Sigma' := (\emptyset; \{\equiv^2\} + \mathcal{F}_\Sigma^{+1} + \mathcal{P}_\Sigma)$ where \equiv is a new interpreted relation symbol of arity two and in the conversion, function symbols have arity lifted by one, hence the \mathcal{F}_Σ^{+1} notation.

Lemma 23. *For any finite¹³ type of function symbols \mathcal{F}_Σ , one has a reduction $\text{FSAT}'(\mathcal{F}_\Sigma; \mathcal{P}_\Sigma) \preceq \text{FSATEQ}(\emptyset; \{\equiv^2\} + \mathcal{F}_\Sigma^{+1} + \mathcal{P}_\Sigma; \equiv^2)$.*

Proof. The idea is to recursively replace a term t over Σ by a formula which is “equivalent” to $x \equiv t$ (where x is a fresh variable not occurring in t) and then an atomic formula like e.g. $P[t_1; t_2]$ by $\exists x_1 x_2. x_1 \equiv t_1 \wedge x_2 \equiv t_2$. We complete the encoding with a formula stating that every function symbol $f : \mathcal{F}_\Sigma$ is encoded into a total functional relation $P_f : \mathcal{F}_\Sigma^{+1}$ of arity augmented by 1. \square

Next, assuming that the function symbols have already been erased, we explain how to merge the relation symbols in a signature $\Sigma = (\emptyset; \mathcal{P}_\Sigma)$ into a single relation symbol, provided that there is an upper bound for the arities in \mathcal{P}_Σ .

Lemma 24. *The reduction $\text{FSAT}(\emptyset; \mathcal{P}_\Sigma) \preceq \text{FSAT}(\emptyset; \{Q^{1+n}\})$ holds when \mathcal{P}_Σ is a finite and discrete type of relation symbols and $|P| \leq n$ holds for all $P : \mathcal{P}_\Sigma$.*

Proof. This comprises three independent reductions, see Fact 25 below. \square

In the following, we denote by \mathcal{F}_Σ^n (resp. \mathcal{P}_Σ^n) the same type of function (resp. relation) symbols but where the arity is uniformly converted to n .

Fact 25. *Let $\Sigma = (\mathcal{F}_\Sigma; \mathcal{P}_\Sigma)$ be a signature:*

1. $\text{FSAT}(\mathcal{F}_\Sigma; \mathcal{P}_\Sigma) \preceq \text{FSAT}(\mathcal{F}_\Sigma; \mathcal{P}_\Sigma^n)$ if $|P| \leq n$ holds for all $P : \mathcal{P}_\Sigma$;
2. $\text{FSAT}(\emptyset; \mathcal{P}_\Sigma^n) \preceq \text{FSAT}(\mathcal{P}_\Sigma^0; \{Q^{1+n}\})$ if \mathcal{P}_Σ is finite;
3. $\text{FSAT}(\mathcal{F}_\Sigma^0; \mathcal{P}_\Sigma) \preceq \text{FSAT}(\emptyset; \mathcal{P}_\Sigma)$ if \mathcal{F}_Σ is discrete.

Proof. For the first reduction, every atomic formula of the form $P \mathbf{v}$ with $|\mathbf{v}| = |P| \leq n$ is converted to $P \mathbf{v}'$ with $\mathbf{v}' := \mathbf{v} + [x_0; \dots; x_0]$ and $|\mathbf{v}'| = n$ for an arbitrary term variable x_0 . The rest of the structure of formulas is unchanged.

For the second reduction, we convert every atomic formula $P \mathbf{v}$ with $|\mathbf{v}| = n$ into $Q(P :: \mathbf{v})$ where P now represents a constant symbol (Q is fixed).

For the last reduction, we replace every constant symbol by a corresponding fresh variable chosen above all the free variables of the transformed formula. \square

¹³ In the Coq code, we prove the theorem for finite or discrete types of function symbols.

4.3 Compressing n -ary Relations to Binary Membership

Let $\Sigma_n = (\emptyset; \{P^n\})$ be a singleton signature where P is of arity n . We now show that P can be compressed to a binary relation modelling membership via a construction using hereditarily finite sets [13] (useful only when $n \geq 3$).

Theorem 26. $\text{FSAT}'(\emptyset; \{P^n\}) \preceq \text{FSAT}(\emptyset; \{\dot{\in}^2\})$.

Technically, this reduction is one of the most involved in this work, although in most presentations of Trakhtenbrot's theorem, this is left as an “easy exercice,” see e.g. [11]. Maybe it is perceived so because it relies one the encoding of tuples in set theory, which is somehow natural for mathematicians¹⁴ but properly building the finite set model in constructive type theory was not that easy.

Here we only give an overview of the main tools. We encode an arbitrary n -ary relation $R : X^n \rightarrow \mathbb{P}$ over a finite type X in the theory of *membership* over the signature $\Sigma_2 = (\emptyset; \{\dot{\in}^2\})$. Membership is much weaker than set theory because the only required set-theoretic axiom is *extensionality*. Two sets are extensionally equal if their members are the same, and extensionality states that two extensionally equal sets belong to the same sets:

$$\forall xy. (\dot{\forall} z. z \dot{\in} x \leftrightarrow z \dot{\in} y) \dot{\rightarrow} \dot{\forall} z. x \dot{\in} z \dot{\rightarrow} y \dot{\in} z \quad (1)$$

As a consequence, no first-order formula over Σ_2 can distinguish two extensionally equal sets. Notice that the language of membership theory (and set theory) does not contain any function symbol, hence, contrary to usual mathematical practices, there is no other way to handle a set than via its characterising formula which makes it a very cumbersome language to work with formally. However, this is how we have to proceed in the Coq development but here, we stick to meta-level “terms” in the prose for simplicity.

The ordered pair of two sets p and q is encoded as $(p, q) := \{\{p\}, \{p, q\}\}$ while the n -tuple (t_1, \dots, t_n) is encoded as $(t_1, (t_2, \dots, t_n))$ (recursively). The reduction function which maps formulas over Σ_n to formulas over Σ_2 proceeds as follows. We reserve two first-order variables d (for the domain D) and r (for the relation R). We describe the recursive part of the reduction $\Sigma_{n \rightsquigarrow 2}^r$

$$\begin{array}{ll} \Sigma_{n \rightsquigarrow 2}^r(P \mathbf{v}) := \text{“tuple } \mathbf{v} \dot{\in} r\text{”} & \Sigma_{n \rightsquigarrow 2}^r(\dot{\forall} z. \varphi) := \dot{\forall} z. z \dot{\in} d \dot{\rightarrow} \Sigma_{n \rightsquigarrow 2}^r(\varphi) \\ \Sigma_{n \rightsquigarrow 2}^r(\varphi \dot{\square} \psi) := \Sigma_{n \rightsquigarrow 2}^r(\varphi) \dot{\square} \Sigma_{n \rightsquigarrow 2}^r(\psi) & \Sigma_{n \rightsquigarrow 2}^r(\dot{\exists} z. \varphi) := \dot{\exists} z. z \dot{\in} d \wedge \Sigma_{n \rightsquigarrow 2}^r(\varphi) \end{array}$$

ignoring the de Bruijn syntax (which would imply adding d and r as parameters). Notice that d and r should not occur freely in φ . In addition, we require that:

$$\begin{array}{ll} \varphi_1 := \dot{\in} \text{ is extensional} & \text{see Equation (1);} \\ \varphi_2 := \dot{\exists} z. z \dot{\in} d & \text{i.e. } d \text{ is non-empty;} \\ \varphi_3 := x_1 \dot{\in} d \wedge \dots \wedge x_k \dot{\in} d & \text{where } [x_1; \dots; x_k] = \text{FV}(\varphi). \end{array}$$

This gives us the reduction function $\Sigma_{n \rightsquigarrow 2}^r(\varphi) := \varphi_1 \wedge \varphi_2 \wedge \varphi_3 \wedge \Sigma_{n \rightsquigarrow 2}^r(\varphi)$.

¹⁴ in our case we use Kuratowski's encoding.

The *completeness* of the reduction $\Sigma_{n \rightsquigarrow 2}$ is the easy part. Given a finite model of $\Sigma_{n \rightsquigarrow 2}(\varphi)$ over Σ_2 , we recover a model of φ over Σ_n by selecting as the new domain the members of d and the interpretation of $P \mathbf{v}$ is given by testing whether the encoding of \mathbf{v} as a n -tuple is a member of r .

The *soundness* of the reduction $\Sigma_{n \rightsquigarrow 2}$ is formally difficult but given Theorem 27 below, it becomes somewhat easier. The model built in its proof is based on a constructive treatment of hereditarily finite sets [13], see Appendix C.

Theorem 27. *Given a decidable n -ary relation $R : X^n \rightarrow \mathbb{P}$ over a finite, discrete and inhabited type X , one can compute a finite and discrete type Y equipped with a decidable relation $\in : Y \rightarrow Y \rightarrow \mathbb{P}$, two distinguished elements $d, r : Y$ and a pair of maps $i : X \rightarrow Y$ and $s : Y \rightarrow X$ s.t.*

1. \in is extensional;
2. extensionally equal elements of Y are equal;
3. all n -tuples of members of d exist in Y ;
4. $\forall x : X. i x \in d$;
5. $\forall y : Y. y \in d \rightarrow \exists x. y = i x$;
6. $\forall x : X. s(i x) = x$;
7. $R \mathbf{v}$ iff $i(\mathbf{v})$ is a n -tuple member of r , for any $\mathbf{v} : X^n$.

4.4 From Discrete Signatures to the Binary Signature

Combining all the previous results, we give a reduction from any discrete signature to the binary singleton signature.

Theorem 28. $\text{FSAT}(\Sigma) \preceq \text{FSAT}(\emptyset; \{P^2\})$ holds for any discrete signature Σ .

Proof. Let us first consider the case of $\Sigma_{n,m} = (\mathbb{F}_n; \mathbb{F}_m)$, a signature over the finite and discrete types \mathbb{F}_n and \mathbb{F}_m . Then we have a reduction $\text{FSAT}(\mathbb{F}_n; \mathbb{F}_m) \preceq \text{FSAT}(\emptyset; \{P^2\})$ by combining Theorems 20, 21 and 26 and Lemmas 23 and 24.

Let us denote by $f_{n,m}$ the reduction $\text{FSAT}(\mathbb{F}_n; \mathbb{F}_m) \preceq \text{FSAT}(\emptyset; \{P^2\})$. Let us now consider a fixed discrete signature Σ . For a formula φ over Σ , using Lemma 22, we compute a signature $\Sigma_{n,m}$ and ψ over $\Sigma_{n,m}$ s.t. $\text{FSAT}(\Sigma) \varphi_1 \leftrightarrow \text{FSAT}(\mathbb{F}_n; \mathbb{F}_m) \psi$. The map $\lambda \varphi. f_{n,m} \psi$ is the required reduction. \square

Lemma 29. $\text{FSAT}(\emptyset; \{P^2\}) \preceq \text{FSAT}(\{f^n\}; \{Q^1\})$ when $n \geq 2$.

Proof. We encode the binary relation $\lambda x y. P[x; y]$ with $\lambda x y. Q(f[x; y; \dots])$, using the first two parameters of f to encode pairing. But since we need to change the domain of the model, we also use a fresh variable d to encode the domain as $\lambda x. Q(f[d; x; \dots])$ and we restrict all quantifications to the domain similarly to the encoding $\Sigma_{n \rightsquigarrow 2}^r$ of Section 4.3. \square

We finish the reduction chains with the weakest possible signature constraints. The following reductions have straightforward proofs.

Fact 30. *One has reductions for the three statements below (for $n \geq 2$):*

1. $\text{FSAT}(\emptyset; \{P^2\}) \preceq \text{FSAT}(\emptyset; \{P^n\})$;
2. $\text{FSAT}(\emptyset; \{P^n\}) \preceq \text{FSAT}(\Sigma)$ if Σ contains an n -ary relation symbol;
3. $\text{FSAT}(\{f^n\}; \{Q^1\}) \preceq \text{FSAT}(\Sigma)$ if Σ contains an n -ary fun. and a unary rel.

5 Decidability Results

Complementing the previously studied negative results, we now examine the conditions allowing for decidable satisfiability problems.

Lemma 31. (FSAT over a fixed domain) *Given a discrete signature Σ and a discrete and finite type D , one can decide whether or not a formula over Σ has a (finite) model over domain D .*

Proof. By Fact 12, satisfaction in a given finite model is decidable. It is also invariant under extensional equivalence, so we only need to show that there are finitely many (decidable) models over D up to extensional equivalence.¹⁵ \square

Lemma 32. *A formula over a signature Σ has a finite and discrete model if and only if it has a (finite) model over \mathbb{F}_n for some $n : \mathbb{N}$.*

Proof. If φ has a model over a discrete and finite domain D , by Corollary 10, one can bijectively map D to \mathbb{F}_n and transport the model along this bijection. \square

Lemma 33. $\text{FSAT}(\emptyset; \mathcal{P}_\Sigma)$ is decidable if \mathcal{P}_Σ is discrete with uniform arity 1.

Proof. By Lemma 22, we can assume $\mathcal{P}_\Sigma = \mathbb{F}_n$ w.l.o.g. We show that if φ has a finite model then it must have a model over domain $\{\mathbf{v} : \mathbb{B}^n \rightarrow \mathbb{B} \mid b\mathbf{v} = \text{tt}\}$ for some Boolean subset $b : (\mathbb{B}^n \rightarrow \mathbb{B}) \rightarrow \mathbb{B}$. Up to extensional equivalence, there are only finitely many such subsets b and we conclude with Lemma 31. \square

Lemma 34. *For any finite type \mathcal{P}_Σ of relation symbols and signatures of uniform arity 1, we have a reduction $\text{FSAT}(\mathbb{F}_n; \mathcal{P}_\Sigma) \preceq \text{FSAT}(\emptyset; \mathbb{L}\mathbb{F}_n \times \mathcal{P}_\Sigma + \mathcal{P}_\Sigma)$.*

Proof. We implemented a proof somewhat inspired by that of Proposition 6.2.7 (Grädel) in [1, pp 251] but the invariant suggested in the iterative process described there did not work out formally and we had to proceed in a single conversion step instead, switching from single symbols to lists of symbols. \square

If functions or relations have arity 0, one can always lift them to arity 1 using a fresh variable (of arbitrary value), like in Fact 25, item (1).

Fact 35. *The reduction $\text{FSAT}(\mathcal{F}_\Sigma; \mathcal{P}_\Sigma) \preceq \text{FSAT}(\mathcal{F}_\Sigma^1; \mathcal{P}_\Sigma^1)$ holds when all arities in Σ are less than 1, where \mathcal{F}_Σ^1 and \mathcal{P}_Σ^1 denote arities uniformly updated to 1.*

6 Signature Classification

We conclude with the exact classification of FSAT regarding enumerability, decidability, and undecidability depending on the properties of the signature.

Theorem 36. *Given $\Sigma = (\mathcal{F}_\Sigma; \mathcal{P}_\Sigma)$ where both \mathcal{F}_Σ and \mathcal{P}_Σ are data types, the finite satisfiability problem for formulas over Σ is enumerable.*

¹⁵ Without discreteness of Σ , it is impossible to build the list of models over $D = \mathbb{B}$.

Proof. Using Theorem 20 and Lemmas 31 and 32, one constructs a predicate $Q : \mathbb{N} \rightarrow \text{Form}_\Sigma \rightarrow \mathbb{B}$ s.t. $\text{FSAT}(\Sigma)\varphi \leftrightarrow \exists n. Q n \varphi = \text{tt}$. Then, it is easy to build a computable enumeration $e : \mathbb{N} \rightarrow \bigcup \text{Form}_\Sigma$ of $\text{FSAT}(\Sigma) : \text{Form}_\Sigma \rightarrow \mathbb{P}$. \square

Theorem 37. (Full Monadic FOL) $\text{FSAT}(\Sigma)$ is decidable if Σ is discrete with arities less or equal than 1, or if all relation symbols arities are 0.

Proof. If all arities are ≤ 1 , then by Fact 35, we can assume Σ is of uniform arity 1. Then for a formula φ over Σ with uniform arity 1 and let us decide FSAT for φ . By Theorem 22, we can compute a signature $\Sigma_{n,m} = (\mathbb{F}_n; \mathbb{F}_m)$ and a formula ψ over $\Sigma_{n,m}$ equi-satisfiable with φ . Using the reduction of Lemma 34, we compute a formula γ , equi-satisfiable with ψ , over a discrete signature of uniform arity 1, void of function symbols. We decide the satisfiability of γ by Lemma 33.

If all relation symbols have arity 0, regardless of \mathcal{F}_Σ , no term can occur in formulas, hence neither can function symbols. Starting from φ over $\Sigma = (\mathcal{F}_\Sigma; \mathcal{P}_\Sigma^0)$ where only \mathcal{P}_Σ is assumed discrete, we compute an equi-satisfiable formula ψ over $\Sigma' = (\emptyset; \mathcal{P}_\Sigma^0)$ and we are back to the previous case. \square

Theorem 38. (Full Trakhtenbrot) If Σ contains either an at least binary relation symbol, or a unary relation symbol and an at least binary function symbol, then BPCP reduces to $\text{FSAT}(\Sigma)$, hence $\text{FSAT}(\Sigma)$ is undecidable.

Proof. By Theorems 14, 21 and 28, Lemma 29 and Fact 30. \square

7 Discussion

The main part of our Coq development directly concerned with the classification of finite satisfiability consists of 10k loc, in addition to 3k loc of (partly reused) utility libraries. Most of the code comprises the signature transformations with more than 4k loc for reducing discrete signatures to membership. Comparatively, the initial reduction from BPCP to $\text{FSATEQ}(\Sigma_{\text{BPCP}})$ takes less than 500 loc.

Our mechanisation of first-order logic in principle follows previous developments [6,7] but also differs in a few aspects. Notably, we had to separate function from relation signatures to be able to express distinct signatures that agree on one sort of symbols computationally. Moreover, we found it favourable to abstract over the logical connectives in form of \Box and ∇ to shorten syntactic proofs. Finally, we did not use the Autosubst 2 support for de Bruijn syntax [15] to avoid its current dependency on the functional extensionality axiom.

We plan to have our development included in the Coq library of synthetic undecidability proofs [9]. In this context, we also want to generalise some of the intermediate signature reductions so that they become reusable for other undecidability proofs concerning first-order logic over arbitrary models.

As further future directions, we want to explore and mechanise the direct consequences of Trakhtenbrot's theorem such as the undecidability of query containment and equivalence in data base theory, or the undecidability of separation logic [2,3]. Also possible, though rather ambitious, would be to mechanise the classification of first-order satisfiability with regards to the quantifier prefix as comprehensively developed in [1].

References

1. E. Börger, E. Grädel, and Y. Gurevich. *The Classical Decision Problem*. Perspectives in Mathematical Logic. Springer-Verlag Berlin Heidelberg, 1997.
2. R. Brochenin, S. Demri, and E. Lozes. On the almighty wand. *Information and Computation*, 211:106–137, 2012.
3. C. Calcagno, H. Yang, and P. W. O’Hearn. Computability and Complexity Results for a Spatial Assertion Language for Data Structures. In R. Hariharan, V. Vinay, and M. Mukund, editors, *FST TCS 2001: Foundations of Software Technology and Theoretical Computer Science*, pages 108–119, Berlin, Heidelberg, 2001. Springer Berlin Heidelberg.
4. N. G. de Bruijn. Lambda calculus notation with nameless dummies, a tool for automatic formula manipulation, with application to the Church-Rosser theorem. *Indagationes Mathematicae (Proceedings)*, 75(5):381–392, Jan. 1972.
5. Y. Forster, E. Heiter, and G. Smolka. Verification of PCP-related computational reductions in Coq. In *International Conference on Interactive Theorem Proving*, pages 253–269. Springer, 2018.
6. Y. Forster, D. Kirst, and G. Smolka. On synthetic undecidability in Coq, with an application to the Entscheidungsproblem. In *International Conference on Certified Programs and Proofs*, pages 38–51. ACM, 2019.
7. Y. Forster, D. Kirst, and D. Wehr. Completeness Theorems for First-Order Logic Analysed in Constructive Type Theory. In *Symposium on Logical Foundations Of Computer Science, 2020, Deerfield Beach, Florida, U.S.A.*, Jan 2020.
8. Y. Forster and D. Larchey-Wendling. Certified undecidability of intuitionistic linear logic via binary stack machines and Minsky machines. In *Proceedings of the 8th ACM SIGPLAN International Conference on Certified Programs and Proofs*, pages 104–117. ACM, 2019.
9. Y. Forster, D. Larchey-Wendling, A. Dudenhefner, E. Heiter, D. Kirst, F. Kunze, G. Smolka, S. Spies, D. Wehr, and M. Wuttke. A Coq Library of Undecidable Problems. In *CoqPL 2020*, New Orleans, LA, United States, 2020.
10. D. Larchey-Wendling and Y. Forster. Hilbert’s Tenth Problem in Coq. In *4th International Conference on Formal Structures for Computation and Deduction*, volume 131 of *LIPICS*, pages 27:1–27:20, Feb 2019.
11. L. Libkin. *Elements of Finite Model Theory*. Springer Publishing Company, Incorporated, 1st edition, 2010.
12. L. Löwenheim. Über Möglichkeiten im Relativkalkül. *Mathematische Annalen*, 76:447–470, 1915.
13. G. Smolka and K. Stark. Hereditarily Finite Sets in Constructive Type Theory. In *Interactive Theorem Proving - 7th International Conference, ITP 2016, Nancy, France, August 22-27, 2016*, volume 9807 of *LNCS*, pages 374–390. Springer, 2016.
14. S. Spies and Y. Forster. Undecidability of higher-order unification formalised in coq. In *International Conference on Certified Programs and Proofs, CPP 2020, New Orleans, USA*, Jan 2020.
15. K. Stark, S. Schäfer, and J. Kaiser. Autosubst 2: reasoning with multi-sorted de Bruijn terms and vector substitutions. In *International Conference on Certified Programs and Proofs*, pages 166–180. ACM, 2019.
16. B. A. Trakhtenbrot. The impossibility of an algorithm for the decidability problem on finite classes. *Dokl. Akad. Nok. SSSR*, 70(4):569–572, 1950.
17. W. Veldman and F. Waaldijk. Some Elementary Results in Intuitionistic Model Theory. *The Journal of Symbolic Logic*, 61(3):745–767, 1996.

A Tools for Finite Types

Theorem 39. (Finite PHP) Let $R : X \rightarrow Y \rightarrow \mathbb{P}$ be a binary relation and $l : \mathbb{L}X$ and $m : \mathbb{L}Y$ be two lists where m is shorter than l ($|m| < |l|$). If R is total from l to m ($\forall x. x \in l \rightarrow \exists y. y \in m \wedge Rxy$) then the values at two distinct positions in l are related to the same y in m , i.e. there exist $x_1, x_2 \in l$ and $y \in m$ such that l has shape $l = \dots \# x_1 :: \dots \# x_2 :: \dots$ and Rx_1y and Rx_2y .

Proof. We start with the case where R is the identity relation $=_X$ on X , hence we want to establish that l contains a duplicate. We first prove the following generalised statement: if $|m| \leq |l|$ and $l \sqsubseteq_f m$ (i.e. $\forall x. x \in l \rightarrow x \in m$) then either l contains a duplicate or l and m are permutable. We establish the generalised statement by structural induction on m .

In particular, when $|m| < |l|$ then l and m cannot be permutable (because permutations preserve length), hence l must contain a duplicate. Generalizing from $=_X$ to an arbitrary relation $R : X \rightarrow Y \rightarrow \mathbb{P}$ is then a simple exercise. \square

Fact 40. Every strict order on a finite type is well-founded.

Proof. For a constructive proof, one can for instance show that descending chains cannot contain a duplicate (otherwise this would give an impossible cycle in a strict order), hence by the PHP, the length of descending chains is bounded by the length of the enumerating list of the finite type. \square

Theorem 41. (Finite decidable quotient) Let $\sim : X \rightarrow X \rightarrow \mathbb{P}$ be a decidable equivalence with $\{l_r : \mathbb{L}X \mid \forall x \exists y. y \in l_r \wedge x \sim y\}$, i.e. finitely many equivalence classes.¹⁶ Then one can compute the quotient X/\sim onto \mathbb{F}_n for some n , i.e. $n : \mathbb{N}$, $c : X \rightarrow \mathbb{F}_n$ and $r : \mathbb{F}_n \rightarrow X$ s.t. $\forall p. c(rp) = p$ and $\forall x y. x \sim y \leftrightarrow cx = cy$.

Proof. From the list l_r of representatives of equivalence classes, remove duplicate representatives using the strong decidability of \sim . This gives a list l'_r which now contains exactly one representative for each equivalence class. Convert l'_r to a vector \mathbf{v} . The function r (representative) is defined by $r := \lambda p. \mathbf{v}_p$. The function c (for class) is simple search: cx is the first (and unique) p such that $\mathbf{v}_p \sim x$. \square

B Discrete Domains

We give an account of the proof of Theorem 19 (numbered 46 later in Appendix B) stating the first-order indistinguishability \doteq up to two given lists $l_{\mathcal{F}}$ and $l_{\mathcal{P}}$ of function and relation symbols respectively, and defined by

$$x \doteq y := \forall(\rho : \mathbb{N} \rightarrow D) \varphi. \mathcal{F}_\varphi \subseteq l_{\mathcal{F}} \rightarrow \mathcal{P}_\varphi \subseteq l_{\mathcal{P}} \rightarrow \mathcal{M} \models_{x \cdot \rho} \varphi \leftrightarrow \mathcal{M} \models_{y \cdot \rho} \varphi$$

is a strongly decidable equivalence and a congruence for all the symbols in $l_{\mathcal{F}}/l_{\mathcal{P}}$. To remain simple, we avoid displaying the dependency on $l_{\mathcal{F}}$, $l_{\mathcal{P}}$, D and \mathcal{M} in

¹⁶ hence l_r denotes a list of *representatives* of equivalence classes.

the notation \doteq as they remain fixed in this section anyway. Equivalence and congruence of \doteq are easy, but congruence is of course limited to l_F/l_P :

$$\begin{aligned} \forall(f : \mathcal{F}_\Sigma)(\mathbf{v} \mathbf{w} : D^{|f|}). f \in l_F \rightarrow (\forall i : \mathbb{F}_{|f|}. \mathbf{v}_i \doteq \mathbf{w}_i) \rightarrow f^\mathcal{M} \mathbf{v} \doteq f^\mathcal{M} \mathbf{w} \\ \forall(P : \mathcal{P}_\Sigma)(\mathbf{v} \mathbf{w} : D^{|P|}). P \in l_P \rightarrow (\forall i : \mathbb{F}_{|P|}. \mathbf{v}_i \doteq \mathbf{w}_i) \rightarrow P^\mathcal{M} \mathbf{v} \leftrightarrow P^\mathcal{M} \mathbf{w} \end{aligned}$$

However, the definition of \doteq hints at no clue for its decidability. We therefore switch to an alternate definition of \doteq as a bisimulation.¹⁷ Using Kleene's fixpoint theorem, we would get \doteq as $\bigcap_{n < \omega} F^n(\lambda u v. \top)$ for some ω -continuous operator F . Hopefully, finiteness would ensure that only finitely many (as opposed to ω) iterations of the operator F are needed for the fixpoint to be reached, hence preserving finitary properties such as decidability.

So let us define the operators $F_F, F_P : (D \rightarrow D \rightarrow \mathbb{P}) \rightarrow (D \rightarrow D \rightarrow \mathbb{P})$ that map a binary relation $\mathcal{R} : D \rightarrow D \rightarrow \mathbb{P}$ to

$$\begin{aligned} F_F(\mathcal{R}) x y &:= \forall f. f \in l_F \rightarrow \forall(\mathbf{v} : D^{|f|})(i : \mathbb{F}_{|f|}). \mathcal{R}(f^\mathcal{M} \mathbf{v}[x/i])(f^\mathcal{M} \mathbf{v}[y/i]) \\ F_P(\mathcal{R}) x y &:= \forall P. P \in l_P \rightarrow \forall(\mathbf{v} : D^{|P|})(i : \mathbb{F}_{|P|}). P^\mathcal{M} \mathbf{v}[x/i] \leftrightarrow P^\mathcal{M} \mathbf{v}[y/i]. \end{aligned}$$

Fact 42. *The following results hold for the operator F_F (resp. F_P).*

1. F_F is monotonic, i.e. $\mathcal{R} \subseteq \mathcal{R}' \rightarrow F_F(\mathcal{R}) \subseteq F_F(\mathcal{R}')$;
2. F_F is continuous, i.e. $\bigcap_n F_F(\mathcal{R}_n) \subseteq F_F(\bigcap_n \mathcal{R}_n)$ with decreasing $(\mathcal{R}_n)_{n < \omega}$;
3. F_F preserves reflexivity, i.e. $=_D \subseteq F_F(=_D)$;
4. F_F preserves symmetry, i.e. $F_F^{-1}(\mathcal{R}) \subseteq F_F(\mathcal{R}^{-1})$;
5. F_F preserves transitivity, i.e. $F_F(\mathcal{R}) \circ F_F(\mathcal{R}) \subseteq F_F(\mathcal{R} \circ \mathcal{R})$;
6. F_F preserves decidability, i.e. if \mathcal{R} is decidable then so is $F_F(\mathcal{R})$.

Hence the combination $F(\mathcal{R}) := F_F(\mathcal{R}) \cap F_P(\mathcal{R})$ also preserves these properties.

Proof. The proofs of items (1)-(5) are easy, even without assuming boundedness by l_F/l_P . However, to ensure the preservation of decidability (6), that bound is essential for the quantification over l_F (resp. l_P) to stay finite. Notice that since D is finite then so is $D^{|f|}$ and the remaining quantifications over $\mathbf{v} : D^{|f|}$ and $i : \mathbb{F}_{|f|}$ are finite quantifications again. Hence, they behave as finitary conjunctions and thus preserve decidability. Notice that compared to F_F , the case of F_P is degenerated because it is constant w.r.t. \mathcal{R} . \square

Theorem 43. *First-order indistinguishability \doteq up to l_F/l_P is extensionally equivalent to \equiv_F (Kleene's greatest fixpoint of F), i.e. for any $x, y : D$ we have*

$$x \doteq y \leftrightarrow x \equiv_F y \quad \text{where } x \equiv_F y := \forall n : \mathbb{N}. F^n(\lambda u v. \top) x y.$$

Proof. For the \rightarrow implication, it is enough to show that \doteq is a pre-fixpoint of F , i.e. $\doteq \subseteq F(\doteq)$, and we get this result using suitable substitutions. The converse implication \leftarrow follows from the fact that \equiv_F is a fixpoint of F , hence its is a congruence for every symbol in l_F/l_P , so $x \equiv_F y$ entails that formulas built from l_F/l_P cannot distinguish x from y . \square

¹⁷ that is the greatest fixpoint of an ω -continuous operator.

With \equiv_F we have a more workable characterization of \doteq but still no decidability result for it since the quantification over n in $\forall n : \mathbb{N}. F^n(\lambda u v. \top) x y$ ranges over the infinite domain \mathbb{N} . We now establish that the greatest fixpoint is reached after finitely many iterations of F . Classically one would argue that F operates over the finite domain of binary relations over D and since the sequence $\lambda n. F^n(\lambda u v. \top)$ cannot decrease strictly forever (by the PHP), it must stay constant after at most $n_0 := 2^{d \times d}$ iterations where $d := \text{card } D$.

Unfortunately, no such reasoning is constructively possible since even for the unit type $\mathbb{1}$, there is no list enumerating the predicates $\mathbb{1} \rightarrow \mathbb{P}$. However, there is a notion of weak powerset. Recall that a predicate $p : X \rightarrow \mathbb{P}$ is *weakly decidable* if it satisfies $\forall x. p x \vee \neg p x$.

Lemma 44. (Weak powerset) *For every finite type X , one can compute a list $ll : \mathbb{L}(X \rightarrow \mathbb{P})$ which contains every weakly decidable predicate in $X \rightarrow \mathbb{P}$ up to extensionality, i.e. $\forall p : X \rightarrow \mathbb{P}. (\forall x. p x \vee \neg p x) \rightarrow \exists q. q \in ll \wedge \forall x. p x \leftrightarrow q x$.*

Proof. The list ll is built by induction on the list $l_X : \mathbb{L} X$ enumerating X . If l_X is $[]$ then X is a void type and thus $ll := (\lambda z. \top) :: []$ fits. If l_X is $x :: l$ then we apply the induction hypothesis to l and get ll for the finite sub-type composed of the elements of l and we define $ll' := (\lambda p z. x \neq z \wedge p z) ll + (\lambda p z. x = z \vee p z) ll$. We check that ll' contains every weakly decidable predicate over $x :: l$. Notice that $|ll'| = 2|ll|$ in the induction step, hence one could easily show that $|ll| = 2^{|l_X|}$, recovering the cardinality of the (classical) powerset. \square

Notice that the weak powerset contains all weakly decidable predicates but not every predicate in it is necessarily weakly decidable.¹⁸ Now we show that $\lambda n. F^n(\lambda u v. \top)$ converges after finitely many steps.

Theorem 45. *One can compute $n : \mathbb{N}$ such that \equiv_F is equivalent to $F^n(\lambda u v. \top)$.*

Proof. By a variant of Lemma 44, we compute the weak powerset of $D \rightarrow D \rightarrow \mathbb{P}$,¹⁹ i.e. a list ll containing every weakly decidable binary relation over D , up to extensional equivalence. Since $\lambda u v. \top : D \rightarrow D \rightarrow \mathbb{P}$ is strongly decidable and F preserves (both weak and) strong decidability, the sequence $\lambda n. F^n(\lambda u v. \top)$ is contained in the list ll , up to extensionality. Hence by Theorem 7 (PHP),²⁰ after $|ll|$ steps, there must have been a duplicate, i.e. there exists $a < b \leq |ll|$ such that $F^a(\lambda u v. \top)$ and $F^b(\lambda u v. \top)$ are extensionally equivalent. However the values of a and b are not computed by the PHP but we can still deduce that $F^n(\lambda u v. \top)$ must be stalled after $n = a$, hence a fortiori after $n = |ll|$. Hence $F^{|ll|}(\lambda u v. \top)$ is extensionally equivalent to \equiv_F . \square

Theorem 46. *First-order indistinguishability up to $l_{\mathcal{F}}/l_{\mathcal{P}}$ is a strongly decidable equivalence and a congruence for all the symbols in $l_{\mathcal{F}}/l_{\mathcal{P}}$.*

¹⁸ unless X is moreover discrete.

¹⁹ via $D \rightarrow D \rightarrow \mathbb{P} \simeq D \times D \rightarrow \mathbb{P}$, and finiteness of $D \times D$.

²⁰ and here we really need a finite PHP over non-discrete types.

Proof. Remember that the real difficulty was strong decidability. By Theorem 45, the operator F reaches its fixpoint \equiv_F after finitely many steps, and by Fact 42 item 6, F preserves decidability, hence by an obvious induction, \equiv_F is decidable. By Theorem 43, the equivalent indistinguishability relation \doteq is decidable. \square

As a side remark, notice that we also show that F preserves first-order definability. A relation $\mathcal{R} : D \rightarrow D \rightarrow \mathbb{P}$ is *first-order definable* if there is a formula φ built only from l_F/l_P such that $\forall \rho. \mathcal{R}(\rho x_0)(\rho x_1) \leftrightarrow \mathcal{M} \models_\rho \varphi$. By Theorem 45, there is thus a first-order formula which is equivalent to first-order indistinguishability in \mathcal{M} . Since its semantics does not depend on variables other than x_0 and x_1 , one can remap all other variables to e.g. x_0 hence we can even ensure that the first-order formula defining \doteq contains only two free variables, namely x_0 and x_1 .

C A Model of Hereditarily Finite Sets

Theorem 47. *Given a decidable n -ary relation $R : X^n \rightarrow \mathbb{P}$ over a finite, discrete and inhabited type X , one can compute a finite and discrete type Y equipped with a decidable relation $\in : Y \rightarrow Y \rightarrow \mathbb{P}$, two distinguished elements $d, r : Y$ and a pair of maps $i : X \rightarrow Y$ and $s : Y \rightarrow X$ s.t.*

1. \in is extensional;
2. extensionally equal elements of Y are equal;
3. all n -tuples of members of d exist in Y ;
4. $\forall x : X. i x \in d$;
5. $\forall y : Y. y \in d \rightarrow \exists x. y = i x$;
6. $\forall x : X. s(i x) = x$;
7. $R \mathbf{v}$ iff $i(\mathbf{v})$ is a n -tuple member of r , for any $\mathbf{v} : X^n$.

Proof. The proof of Theorem 27 (numbered 47 here in Appendix C) is quite involved and we only give its outline, referring to the Coq code for details. The type Y is built from the type of hereditarily finite sets based on [13], and when we use the word “set” below, it means hereditarily finite set. The idea is first to construct d as a *transitive set* of which the elements are in bijection i/s with the type X , hence d is the cardinal of X in the set-theoretic meaning. Then the iterated powersets $\mathcal{P}(d), \mathcal{P}^2(d), \dots, \mathcal{P}^k(d)$ are all transitive as well and contain d both as a member and as a subset. Considering $\mathcal{P}^{2n}(d)$ which contains all the n -tuples built from the members of d , we define r as the set of n -tuples collecting the encoding $i(\mathbf{v})$ of vectors $\mathbf{v} : X^n$ such that $R \mathbf{v}$. We show $r \in p$ for p defined as $p := \mathcal{P}^{2n+2}(d)$. Then we define $Y := \{z \mid z \in p\}$,²¹ restrict membership \in to Y and this gives the finite type equipped with all the required properties. \square

²¹ using the Boolean counterpart of $(\cdot) \in p$ for unicity of proofs.