NECESSARY CONDITIONS FOR FINITE DECIDABILITY IN LOCALLY FINITE VARIETIES ADMITTING STRONGLY ABELIAN BEHAVIOR

By

Matthew Smedberg

Dissertation

Submitted to the Faculty of the

Graduate School of Vanderbilt University

in partial fulfillment of the requirements

for the degree of

DOCTOR OF PHILOSOPHY

in

Mathematics

May 2014

Nashville, Tennessee

Approved:

Ralph McKenzie

Constantine Tsinakis

Mark Sapir

Denis Osin

To all those who have believed in me,

espcially my beloved partner, Lauren

personally

and

Ralph, Matt, Ross, Keith, Miklos, and Petar

professionally.

ACKNOWLEDGMENTS

My thanks go first to my advisor, Ralph McKenzie, without whose encouragement and guidance this project would have been dead in the water. My thanks too to Ralph's network of collaborators worldwide, especially Matt Valeriote, Ross Willard, Miklos Maroti, Petar Markovic, Libor Barto, Marcin Kozik, Pawel Idziak, Peter Mayr, Brian Davey, Keith Kearnes, and Agnes Szendrei, who have given me every opportunity to think and talk through research ideas across three continents. It's been a real adventure.

Gratitude is also due to the faculty of Vanderbilt's mathematics department, who have always been willing to hash out problems both mathematical and otherwise. Likewise to my fellow graduate students, who have collectively created a stimulating and welcoming environment during my time here. I want to thank also the faculty of the City College of New York and the CUNY Graduate Center for their role as the springboard to my advanced studies.

I want lastly to acknowledge my debt to my partner, Lauren, who keeps me on an even keel, and who, I hope, got something out of the deal when she and two skeptical cats pulled up stakes to be here in Nashville with me.

TABLE OF CONTENTS

Page

DEDICATION . . . ii

ACKNOWLEDGMENTS . . . iii

Chapter

I. INTRODUCTION . . . 1
- I–A. Results and outline of arguments . . . 2

II. DEFINITIONS AND LOGICAL PRELIMINARIES . . . 4
- II–A. Algebraic basics . . . 4
- II–B. First-order properties and theories . . . 4
- II–C. The decidability and finite decidability problems . . . 5
- II–D. Abelian algebras, solvability, and TCT . . . 7
- II–E. Powers and subpowers . . . 11
- II–F. Multi-sorted structures . . . 12

III. CONGRUENCE GEOMETRY OF THE STRONGLY-SOLVABLE RADICAL . . . 15
- III–A. Coherence . . . 15
- III–B. Proof of Theorem A . . . 18

IV. STRONGLY-SOLVABLE IMPLIES STRONGLY ABELIAN . . . 25
- IV–A. The radical centralizes minimal sets . . . 25
- IV–B. The action of the twin group . . . 29
- IV–C. The strongly-solvable radical is abelian . . . 33
- IV–D. Proof of Theorem B . . . 43

V. RESIDUAL FINITENESS . . . 46

VI. TWO MULTI-SORTED CONSTRUCTIONS . . . 51
- VI–A. The sorted language corresponding to a congruence . . . 51
- VI–B. The sorted language corresponding to a strongly-abelian congruence . . . 54

VII. ARITY BOUNDS . . . 57
- VII–A. Preparatory lemmas . . . 57
- VII–B. Proof of Theorem D . . . 65

VIII. NEXT STEPS . . . 72

BIBLIOGRAPHY . . . 75

CHAPTER I

INTRODUCTION

Among the algorithmic properties most investigated by algebraists is the problem of when a given computably axiomatizable class $\mathcal{K}$ of first-order structures will have computable first-order theory too. This problem was investigated for varieties of groups and rings beginning in the 1950s, with signal contributions from Tarski and his students in the USA ([TMR53], [Szm55]) and from the Russian school of Luzin, Ershov et al. ([Mal65], [Ers72], [Zam76], [Zam78]).

For many but not all interesting classes $\mathcal{K}$, it was shown that not only is $\mathrm{Th}(\mathcal{K})$ undecidable, but $\mathrm{Th}(\mathcal{K}_{\mathrm{fin}})$ will frequently be so as well, where $\mathcal{K}_{\mathrm{fin}}$ denotes the class of all finite structures in $\mathcal{K}$. We will say that $\mathcal{K}$ is *(finitely) decidable* if $\mathrm{Th}(\mathcal{K})$ (resp. $\mathrm{Th}(\mathcal{K}_{\mathrm{fin}})$) is a computable set of sentences.

For example, any variety of groups has decidable theory iff it contains only abelian groups (as is showed in [Szm55] and [Zam78]). Szmielew actually showed more: the theory of abelian groups has effective Δ-elimination for a nice class Δ of definable properties, most of which are trivial in the case of finite groups; one can use this to show without too much difficulty that every prevariety of finite abelian groups has decidable theory. (For an exposition of this in modern notation, consult [EF72].) Together with the famous construction by Olshanskii of a variety of groups whose smallest nonabelian member is infinite ([Ols91]), shows that a variety can be undecidable and simultaneously finitely-decidable. (Zamyatin had given an earlier example of this for varieties of rings in [Zam76].) Kharlampovich and Sapir give a detailed survey of decidability and other algorithmic properties in varieties (mostly associative and Lie varieties) in [KS95].

We restrict our attention in this paper to varieties of abstract algebras in a finite language. The natural questions here are: given a computably axiomatizable variety $\mathcal{V}$ (in particular, a variety of the form $\mathrm{HSP}(\mathbf{A})$ for some finite algebra $\mathbf{A}$), is $\mathrm{Th}(\mathcal{V})$ (resp. $\mathrm{Th}(\mathcal{V}_{\mathrm{fin}})$) a computable set of sentences? One immediately sees that $\mathrm{Th}(\mathcal{V})$ is computably enumerable, so the decidability problem for $\mathcal{V}$ is equivalent to the computable enumerability of the set of sentences refutable in some member of $\mathcal{V}$; on the other hand, it is also clear that the set of sentences *refuted* in some finite member of $\mathcal{V}$ is computably enumerable, while the set of sentences true in all these algebras may not be.

In [MV89], McKenzie and Valeriote showed that locally-finite *decidable* varieties have a very restricted structure theory. Such a variety must decompose as the varietal product of a discriminator variety, a variety of modules, and a strongly-abelian variety. In particular,

Corollary I.1. *If $\mathcal{V}$ is a locally-finite variety with decidable first-order theory, then every strongly-solvable congruence of an algebra in $\mathcal{V}$ is strongly-abelian.*

While the analogues betwen the decidability problem and the finite decidability problem are strong, not all the necessary conditions for decidability transfer down; Corollary I.1 does, however, as we will show in this paper (Theorem B).

One of the properties that does *not* continue to hold is the direct decomposition theorem. In [Idz97], P. Idziak gave a characterization of finitely-decidable locally-finite varieties with modular congruence lattices; this characterization essentially gives a recipe for building a variety with no possible direct decomposition into discriminator and affine varietal factors. One goal of Chapter VI is to suggest a potential reformulation of the direct product criterion to make it work in the finitely-decidable setting.

I–A. Results and outline of arguments

Our overall goal in this paper is to generalize as many properties as possible from the theory of congruence-modular locally-finite finitely-decidable varieties to the non-modular setting. (See Chapter II for definitions and notations.) The following are our main results:

Theorem A. *Let* $\mathbf{S}$ *be a finite, subdirectly-irreducible algebra belonging to a finitely-decidable, locally-finite variety. Then the strongly-solvable radical* $\mathrm{Rad}_u(\mathbf{S})$ *is comparable to every congruence on* $\mathbf{S}$*, and is either* $\top_S$ *or meet-irreducible with boolean-type upper cover.*

Theorem B. *If* $\mathbf{A}$ *is a finite algebra in a locally-finite, finitely-decidable variety, then every strongly-solvable congruence of* $\mathbf{A}$ *is strongly-abelian.*

Theorem C. *If* $\mathcal{V}$ *is a finitely-generated, finitely-decidable variety, then* $\mathcal{V}$ *has a finite residual bound. In particular, every algebra in* $\mathcal{V}$ *is residually finite.*

Theorem D. *Let* $\mathbf{A}$ *be an algebra in a finitely-decidable variety. Let* $\sigma > \bot_A$ *be the strongly-solvable radical of* $\mathbf{A}$*, and let*

$$\check{t}\colon C_1 \times C_2 \times \cdots \times C_\ell \to C_0$$

be a σ*-sorted term operation of* $\mathbf{A}$*. Then the essential arity of* $\check{t}$ *is no greater than the maximal arity of a* σ*-sorted decomposition term on* C_0*.*

(See Definition II.14 for the meaning of σ-sorted map.)

The plan of attack is as follows: We prove Theorem A in Chapter III. Then in Chapter IV, we employ three semantic interpretation constructions of increasing length and complexity to prove Theorem B. (This chapter makes the most difficult reading of the dissertation.) Chapter V culminates in the proof of Theorem C, in which I employ some old methods from the study of congruence-modular varieties and correct a mistake in a paper of McKenzie and Snow.

After this, we veer off into the territory of multi-sorted structures: Chapter VI defines two constructions which take an ordinary single-sorted algebra, especially one which looks like a good candidate for generating a finitely-decidable variety, and return a multi-sorted algebra which better controls the propagation of nonabelian behavior in the variety this algebra generates. Lastly, in Chapter VII, we use the more complicated of these constructions to generalize a theorem of Valeriote. The original theorem provided an obstruction to finite decidability for strongly-abelian varieties based on bounding the arities of term operations; the use of multi-sorted structures allows us to remove the hypothesis of strong abelianness (at the cost of making the statement of the theorem more opaque).

CHAPTER II

DEFINITIONS AND LOGICAL PRELIMINARIES

All numbered environments, such as theorems, definitions, etc will be numbered in a single sequence within their respective chapters. The end of a proof will be marked with a box as usual; the end of a claim will be marked with a numbered turnstile.

II–A. Algebraic basics

Throughout the paper, an *algebra* will be a structure in a first-order language with only operation (and constant) symbols. We denote algebras in bold face $\mathbf{A}$, and their underlying sets in lightface A. $\perp_A$ and $\top_A$ denote respectively the discrete and total equivalence relations on A, which are congruences of every algebraic structure on A.

As usual, a *congruence* of an algebra $\mathbf{A}$ is an equivalence relation preserved by all basic operations of $\mathbf{A}$. The lattice of congruences of $\mathbf{A}$ will be denoted $\mathrm{Con}(\mathbf{A})$. If $\mathrm{Con}(\mathbf{A})$ has a least nontrivial congruence μ, we call $\mathbf{A}$ *subdirectly-irreducible* and μ its *monolith.* More generally, minimal nontrivial congruences are called *atoms.*

A *polynomial operation* is a function

$$x_1, \ldots, x_k \mapsto t^{\mathbf{A}}(x_1, \ldots, x_k, a_{k+1}, \ldots, a_n)$$

for some $\mathcal{L}$-term t and some elements $a_i \in A$. The set of all polynomial operations of k or fewer variables is denoted $\mathsf{Pol}_k(\mathbf{A})$. Unless otherwise specified, all first-order languages in this paper have only finitely many basic symbols, all of which are operations (or constants). (An important exception to this rule is the *non-indexed algebras* described on page 9.)

A class $\mathcal{V}$ of algebras is said to be *residually* κ if for each $\mathbf{A} \in \mathcal{V}$ and each $a \neq b \in A$ there exists a homomorphism from $\mathbf{A}$ onto a some algebra $\mathbf{B}$ with $|B| < \kappa$, separating a from b. "Residually ω" is usually called "residually finite". A *residual bound* for $\mathcal{V}$ is any cardinal κ such that $\mathcal{V}$ is residually κ. If every finitely generated $\mathbf{A} \in \mathcal{V}$ is finite, we say $\mathcal{V}$ is *locally-finite.*

II–B. First-order properties and theories

If $\mathbf{A}$ a structure in the first-order language L, the *theory* $\mathrm{Th}\,(\mathbf{A})$ of $\mathbf{A}$ is the set of all $\mathcal{L}$-sentences true in $\mathbf{A}$. If $\mathcal{K}$ is a class of $\mathcal{L}$-structures, $\mathrm{Th}\,(\mathcal{K})$ is the set of all sentences true in all members of $\mathcal{K}$. We write $\mathcal{K}_{\mathrm{fin}}$ for the class of all finite members of $\mathcal{K}$ and $\mathrm{Th}_{\mathrm{fin}}(\mathcal{K})$ for $\mathrm{Th}\,(\mathcal{K}_{\mathrm{fin}})$.

A class $\mathcal{K}$ of $\mathcal{L}$-algebras is a *variety* if it is axiomatized by some set of *equations*, that is, sentences of the form

$$\forall \vec{v}\; t_1(\vec{v}) = t_2(\vec{v})$$

for some terms of the language. Equivalently, and more usefully for us, $\mathcal{K}$ is a variety iff it is closed under taking direct products, subalgebras, and surjective homomorphic images. (Cf [MMT87], [BS81].) For a given algebra $\mathbf{A}$ (resp. class $\mathcal{K}$ of algebras) we denote the smallest variety containing it by $\mathrm{HSP}(\mathbf{A})$ (resp. $\mathrm{HSP}(\mathcal{K})$). A basic theorem originally due to Birkhoff asserts that if $\mathcal{V}$ is a variety and κ any cardinal, then $\mathcal{V}$ contains a free algebra on κ generators. If $\kappa < \omega$, elements of this algebra are in canonical bijection (up to $\mathrm{Th}(\mathcal{V})$-equivalence) with $\mathcal{L}$-terms in κ variables.

II–C. The decidability and finite decidability problems

For a given finite structure $\mathbf{A}$, it is a trivial matter to determine whether a given first-order sentence holds in $\mathbf{A}$; the same is not true for the problem of determining whether that same sentence holds throughout some variety containing $\mathbf{A}$, such as $\mathrm{HSP}(\mathbf{A})$.

Fact II.1. Let $\mathbf{A}$ be any finite algebra, $\mathcal{V} = \mathrm{HSP}(\mathbf{A})$.

1 $\mathcal{V}$ is locally-finite and computably axiomatizable; it follows that $\mathrm{Th}(\mathcal{V})$ is computably enumerable. We will say that $\mathcal{V}$ is *decidable* if this set of sentences is computable.

2 The *complement* of $\mathrm{Th}_{\mathrm{fin}}(\mathcal{V})$, the set of all sentences *falsified* in some finite member of $\mathcal{V}$, is computably enumerable. We will say that $\mathcal{V}$ is *finitely-decidable* if this set of sentences is computable.

There do exist finite algebras $\mathbf{A}$ such that $\mathrm{HSP}(\mathbf{A})$ is undecidable and/or finitely undecidable. For example, by [Zam78], any non-abelian finite group generates an undecidable variety; for many other instances of undecidable and/or finitely undecidable varieties, see [Mal65], [Ers72], [Zam76], [Idz86], [MV89], [II88], [Idz89a], [Idz89b], and [Jeo99].

As the alert reader has seen in Fact II.1, there is a fundamental asymmetry between decidability and finite decidability, as in the one case it is the set of provable sentences which is easily shown to be enumerable, while in the other it is the refutable sentences. This asymmetry is not just apparent: the two properties are in fact completely independent. Specific examples of the four possibilities are given in [Szm55], [II88], [Ols91], and [Jeo99].

The principal tool this investigation will employ in establishing undecidability is the method of interpretation. The reader is referred to standard texts [Hod93, Chapter 5], [BS81, Section V.5], for full details; our conventions will be as follows:

Definition II.2. Let L_0 be a single- or multi-sorted first-order language, with sort symbols $s_1, \dots, s_\ell$, function symbols $f_i(v_1, \dots, v_{\mathrm{ar}(f_i)})$, and predicate symbols $r_i(v_1, \dots, v_{\mathrm{ar}(R_i)})$. An interpretation of the structure

$$\mathbf{M} = \langle s_1^{\mathbf{M}}, \dots, s_\ell^{\mathbf{M}}; f_i^{\mathbf{M}}, \dots, r_i^{\mathbf{M}}, \dots \rangle$$

into the structure $\mathbf{N}$ (of a possibly different language L_1) will mean a family of definable subsets of $\mathbf{N}$ which jointly produce an isomorphic copy of $\mathbf{M}$. (We allow finitely many parameters from $\mathbf{N}$ in the definitions.) More precisely, the interpretation consists of L_1-formulas

- $\mathrm{WHO}_i(v)$ for each sort symbol s_i;
- $\mathrm{EQ}(v_1, v_2)$
- $\mathrm{F}_i(v_0, v_1, \ldots, v_{\mathrm{ar} f_i})$ for each f_i;
- $\mathrm{R}_i(v_1, \ldots, v_{\mathrm{ar}(r_i)})$ for each r_i;

such that

- the extensions of the WHO_i in $\mathbf{N}$ are nonempty and disjoint;
- $\mathrm{EQ}(x, y)$ holds only if $\mathrm{WHO}_i(x) \wedge \mathrm{WHO}_i(y)$ holds for some $1 \leq i \leq m$, and on each WHO_i, the formula defines an equivalence relation $\approx$;
- for all function symbols f with type signature

$$f(s_{i_1}, \ldots, s_{i_{\mathrm{ar}(f)}}) \to s_{i_0},$$

  corresponding to the formula $F(v_0, \ldots, v_{\mathrm{ar}(f)})$,

  - if $\mathbf{N} \models F(x_0, x_1, \ldots, x_{\mathrm{ar}(f)})$ then $\mathrm{WHO}_{i_0}(x_0), \ldots, \mathrm{WHO}_{i_{\mathrm{ar}(f)}}(x_{\mathrm{ar}(f)})$;
  - for all $x_1 \in \mathrm{WHO}_{i_1}(\mathbf{N}), \ldots, x_{\mathrm{ar}(f)} \in \mathrm{WHO}_{i_{\mathrm{ar}(f)}}(\mathbf{N})$ there exist elements $x'_1 \approx x_1$, …, $x'_{ar(f)} \approx x_{ar(f)}$ and x_0 so that $\mathbf{N} \models F(x_0, x'_1, \ldots, x'_{\mathrm{ar}(f)})$; and
  - if

$$\begin{gathered} \mathbf{N} \models F(x_0, x_1, \ldots, x_{\mathrm{ar}(f)}) \\ \mathbf{N} \models F(y_0, y_1, \ldots, y_{\mathrm{ar}(f)}) \\ \text{and} \\ x_1 \approx y_1, \ldots, x_{\mathrm{ar}(f)} \approx y_{\mathrm{ar}(f)} \end{gathered}$$

    then $x_0 \approx y_0$,

  inducing an operation $\mathrm{WHO}(f)$ on the indicated $\approx$-classes;
- for all predicate symbols r with type signature

$$r(s_{i_1}, \ldots, s_{i_{\mathrm{ar}(r)}})$$

  corresponding to the formula $R(v_1, \ldots, v_{\mathrm{ar}(r)})$,

$$\begin{gathered} \mathbf{N} \models R(x_1, \ldots, x_{\mathrm{ar}(r)}) \\ \text{only if} \\ \mathbf{N} \models \mathrm{WHO}_{i_1}(x_1) \wedge \cdots \wedge \mathrm{WHO}_{i_{\mathrm{ar}(r)}}(x_{\mathrm{ar}(r)}) \end{gathered}$$

inducing the quotient predicate

$$\mathrm{WHO}(r)(x_1/\approx,\ldots,x_{\mathrm{ar}(r)}/\approx) \iff \exists x_1' \approx x_1,\ldots,x_{\mathrm{ar}(r)}' \approx x_{\mathrm{ar}(r)}\ R(x_1',\ldots,x_{\mathrm{ar}(r)}')$$

on $\approx$-classes;

- $$\langle s_1^{\mathbf{M}},\ldots,s_\ell^{\mathbf{M}}; f_i^{\mathbf{M}},\ldots,r_i^{\mathbf{M}},\ldots\rangle \cong \langle \mathrm{WHO}_i/\approx,\ldots,\mathrm{WHO}_\ell/\approx; \mathrm{WHO}(f_i),\ldots,\mathrm{WHO}(r_i),\ldots\rangle$$

Frequently the equivalence relation EQ will be true equality, in which case we will not mention it explicitly.

Observe that if an undecidable class $\mathcal{G}_{\mathrm{fin}}$ of finite structures interprets into $\mathcal{K}_{\mathrm{fin}}$ as above, then not only $\mathcal{K}$ but every class $\mathcal{K}' \supset \mathcal{K}$ of structures in the language is finitely undecidable as well: we say that $\mathcal{K}$ is *hereditarily* finitely undecidable.

The classes we will be interpreting will be the class of graphs and the class $\mathcal{E}_2$, defined below. For this investigation, a *graph* is a first-order structure $\mathbb{G} = \langle V; E\rangle$, where $E^{\mathbb{G}}$ is a symmetric, irreflexive binary relation. We will not enforce the distinction between an ordered edge $\langle x, y\rangle$ and an unordered edge $\{x, y\}$ for symmetric graphs. (It follows from our definition that graphs in our sense do not possess multiple edges between a single pair of vertices.) It was shown by Ershov and Rabin in the 1960s that graphs are both undecidable and finitely undecidable.

$\mathcal{E}_2$ is the class of structures $\mathbf{E} = \langle I; R_0, R_1\rangle$ where each R_i is a binary predicate symbol whose interpretation in the structure is an equivalence relation on I, such that $R_0^{\mathbf{E}} \cap R_1^{\mathbf{E}} = \perp_I$. We will sometimes refer to $\mathrm{Th}(\mathcal{E}_2)$ as the theory of two disjoint equivalence relations. [BS81, Corollary 5.16] shows that the theory of this class is undecidable and finitely undecidable.

(In fact, it can be shown that for each of the above classes, $\mathrm{Th}(\mathcal{K})$ is *computably inseparable* from the set of sentences finitely refutable in $\mathcal{K}$; but we will not need this stronger property.)

II–D. Abelian algebras, solvability, and TCT

Modern investigations in universal algebra are greatly aided by the linked toolboxes of the theory of solvable and strongly-solvable algebras and congruences (see for example [FM87]) and the "tame congruence theory" developed by McKenzie and Hobby in [HM88].

Let $\mathbf{A}$ be any algebra, and α, β, γ be congruences (or more generally, any binary relations) on A. $\mathbf{A}$ is said to satisfy the *term condition* $\mathrm{C}(\alpha, \beta; \gamma)$ if the implication

$$\begin{gathered} t(\vec{a}_1, \vec{b}_1) \equiv_\gamma t(\vec{a}_1, \vec{b}_2) \\ \Downarrow \\ t(\vec{a}_2, \vec{b}_1) \equiv_\gamma t(\vec{a}_2, \vec{b}_2) \end{gathered}$$

is valid for all terms t and all tuples $\vec{a}_1 \equiv_\alpha \vec{a}_2$ and $\vec{b}_1 \equiv_\beta \vec{b}_2$. If $R, S \subset A$, then we will write $\mathrm{C}(R, S; \gamma)$ when we mean $\mathrm{C}(R^2, S^2; \gamma)$. If $\gamma \leq \beta \in \mathrm{Con}(\mathbf{A})$ and $\mathrm{C}(\beta, \beta; \gamma)$, then we say that β is *abelian* over γ. If $\mathrm{C}(\beta, \beta; \perp_A)$ then we say that β is an abelian congruence. If $\mathrm{C}(\top_A, \top_A; \perp_A)$ then we say that $\mathbf{A}$ is an abelian algebra.

We can always transform a failure

$$\begin{gathered} t(\vec{a}_1, \vec{b}_1) = t(\vec{a}_1, \vec{b}_2) \\ \text{but} \\ t(\vec{a}_2, \vec{b}_1) \neq t(\vec{a}_2, \vec{b}_2) \end{gathered}$$

of $\mathrm{C}(\alpha, \beta; \gamma)$ into one

$$\begin{gathered} s(a'_1, \vec{b}'_1) = s(a'_1, \vec{b}'_2) \\ \text{but} \\ s(a'_2, \vec{b}'_1) \neq s(a'_2, \vec{b}'_2) \end{gathered} \tag{II.2.1}$$

where α-shifting occurs in only one variable. The same is not true in general for the β-shifted variables; however, this is possible in the special case where all the elements in $\vec{b}_1, \vec{b}_2$ are taken from some $U \subset A$ such that every operation on U is realized by a polynomial of $\mathbf{A}$. We leave the verification of this to the reader.

Another asymmetry between the roles played by the first two variables of the term condition has to do with congruence generation. If R is a binary relation on A, then $\mathrm{C}(R, \beta; \gamma)$ holds iff $\mathrm{C}(\rho, \beta; \gamma)$ does, where ρ is the least congruence of $\mathbf{A}$ identifying all the pairs in $R \cup \gamma$. By comparison, $\mathrm{C}(\alpha, R; \gamma)$ holds iff $\mathrm{C}(\alpha, \mathbf{S}; \gamma)$, where $\mathbf{S}$ is the reflexive, symmetric subalgebra of $\mathbf{A}^2$ generated by R.

If $\gamma \leq \beta \in \mathrm{Con}(\mathbf{A})$, we say that β satisfies the *strong* term condition over γ, or that β is *strongly-abelian* over γ, if for all terms t and tuples $\vec{a}_1 \equiv_\beta \vec{a}_2$, $\vec{b}_1 \equiv_\beta \vec{b}_2 \equiv_\beta \vec{b}_3$,

$$\begin{gathered} t(\vec{a}_1, \vec{b}_1) \equiv_\gamma t(\vec{a}_2, \vec{b}_2) \\ \Downarrow \\ t(\vec{a}_1, \vec{b}_3) \equiv_\gamma t(\vec{a}_2, \vec{b}_3) \end{gathered}$$

If $\mathrm{C}(\beta, \beta; \gamma)$, this condition is equivalent to the apparently weaker condition

$$\vec{a}_1 \equiv_\beta \vec{a}_2 \;\&\; \vec{b}_1 \equiv_\beta \vec{b}_2 \;\&\; t(\vec{a}_1, \vec{b}_1) \equiv_\gamma t(\vec{a}_2, \vec{b}_2) \;\Rightarrow\; \forall i, j\; t(\vec{a}_1, \vec{b}_1) \equiv_\gamma t(\vec{a}_i, \vec{b}_j)$$

which is easier to use.

If $\mathbf{A}$ is a locally-finite algebra and $\alpha^- < \alpha^+ \in \mathrm{Con}(\mathbf{A})$, we say that α^+ is (strongly) *solvable* over α^- if every chain of congruences

$$\alpha^- = \beta_0 < \beta_1 < \cdots < \beta_{m-1} < \beta_m = \alpha^+$$

admits a refinement

$$\alpha^- = \gamma_0 < \gamma_1 < \cdots < \gamma_{n-1} < \gamma_n = \alpha^+$$

such that each γ_{i+1} is (strongly) abelian over γ_i.

Let $\mathbf{A}$ be an algebra. For any subset $W \subset A$, the *non-indexed algebra* $\mathbf{A}_{|W}$ *induced by* $\mathbf{A}$ *on* W is defined to have underlying set W, and a basic operation $f(v_1, \ldots, t_k)$ for each polynomial $f \in \mathsf{Pol}_k(\mathbf{A})$ such that $f(W^k) \subset W$. We do not usually wish to specify any more parsimonious signature for an induced algebra; even if the signature of $\mathbf{A}$ was finite, $\mathbf{A}_{|W}$ is not in general representable as a first-order structure in any finite language.

If $\mathbf{A}$ is a finite algebra and and $\alpha \prec \beta$ in $\mathrm{Con}(\mathbf{A})$ (that is, β is an upper cover of α in the order-theoretic sense), an (α, β)-*minimal set* $U \subset A$ is an inclusion-minimal polynomial image $e(A)$ of the algebra, where $e \in \mathsf{Pol}_1(\mathbf{A})$ is required to be *idempotent* ($e \circ e = e$) and to preserve the α-inequivalence of some pair $\langle a, b\rangle \in \beta \setminus \alpha$. Clearly, every (α, β) minimal set has at least two elements. If U is (α, β)-minimal, a $\beta_{|U}$-class which properly contains two or more $\alpha_{|U}$-classes is called a *trace*. The union of the traces included in U is called the *body* of U; the remainder is called the *tail*.

Theorem II.3 (Fundamental Theorem of Tame Congruence Theory, [HM88, Theorem 2.8, Theorem 4.7, Lemma 4.8])**.** *Let* $\mathbf{A}$ *be a finite algebra with congruences* $\alpha \prec \beta$.

(1) All (α, β)*-minimal sets* U_1, U_2 *are polynomially isomorphic, in the sense that there exists* $f \in \mathsf{Pol}_1(\mathbf{A})$ *which maps* U_1 *bijectively to* U_2 *in such a way that every induced operation*

$$t_2 \in U_2^{U_2^k}$$

in the signature of $\mathbf{A}_{|U_2}$ *is the* f*-image of an operation*

$$t_1 \in U_1^{U_1^k}$$

in the signature of $\mathbf{A}_{|U_1}$*.*

(2) Let $N \subset U$ *be any trace in an* (α, β)*-minimal set. If* $\mathbf{A}_{|N}/\alpha_{|N}$ *is isomorphic to the two-element boolean algebra, the two-element lattice, or the two-element semilattice, then we say that the covering is of (respectively) boolean type* ($\alpha \overset{3}{\prec} \beta$)*, lattice type* ($\alpha \overset{4}{\prec} \beta$)*, or semilattice type* ($\alpha \overset{5}{\prec} \beta$)*. (This is well-defined by (1).)*

(3) If none of these possibilities occur, then $\mathbf{A}_{|N}/\alpha_{|N}$ *is an abelian algebra, and is either isomorphic to a finite module over some ring, in which case the cover is of affine type* ($\alpha \overset{2}{\prec} \beta$)*; or*

isomorphic to a finite G-set for some finite group G (unary type, $\alpha \overset{1}{\prec} \beta$). In the former case, β is abelian over α but not strongly-abelian; in the latter, β is strongly-abelian over α.

We will write $\operatorname{typ}\{\mathbf{A}\} \subset \{1, 2, 3, 4, 5\}$ for the set of tame congruence types which appear in $\operatorname{Con}(\mathbf{A})$.

Let $i \neq j$ be tame congruence types. We will say that the algebra $\mathbf{A}$ satisfies the (i, j)-*transfer principle* if, for all covering chains

$$\alpha_1 \overset{i}{\prec} \alpha_2 \overset{j}{\prec} \alpha_3$$

there exists

$$\alpha_1 \overset{j}{\prec} \beta_j \leq \alpha_3$$

and likewise

$$\alpha_1 \leq \beta_i \overset{i}{\prec} \alpha_3$$

Fact II.4. Let $\mathcal{V}$ be a finitely-decidable variety.

(1) $\mathcal{V}$ omits the lattice and semilattice tame congruence types.

(2) The (1,2), (2,1), (3,1), and (3,2) transfer principles hold throughout $\mathcal{V}$; in particular,

(3) If $\mathbf{S} \in \mathcal{V}$ is a finite subdirectly-irreducible algebra with boolean-type monolith, then $\operatorname{typ}\{\mathbf{S}\} = \{3\}$. If the monolith is affine, then $\operatorname{typ}\{\mathbf{S}\} \subset \{2, 3\}$, and if the monolith is unary, then $\operatorname{typ}\{\mathbf{S}\} \subset \{1, 3\}$.

(4) If $\mathbf{A} \in \mathcal{V}$ and $\alpha \overset{2,3}{\prec} \beta$, then all (α, β)-minimal sets have no tail. In the boolean case, this means that each minimal set contains just two elements, and every possible operation from this set to itself is realized by a polynomial of the algebra.

Proof. (1) is proved in [HM88, Theorem 11.1]; it is a consequence of the fact that (finite) graphs interpret semantically into each of

$$\mathrm{HSP}\left(\langle\{0, 1\}; \wedge\rangle\right)$$

and

$$\mathrm{HSP}\left(\langle\{0, 1\}; \wedge, \vee\rangle\right)$$

(2) is proved in [VW92] and [Val94]; (3) follows immediately. (4) is also proved in [VW92]. □

It follows by [HM88, Theorem 8.5] that any locally-finite, finitely-decidable variety omitting the unary type is congruence-modular.

The following fact will be of use later in the paper:

Theorem II.5 ([HM88, Chapter 7])**.** *Let* $\mathbf{A}$ *be any finite algebra.*

(1) Each of the relations

$$\alpha \overset{ss}{\sim} \beta \iff \alpha \text{ is connected to } \beta \text{ via covers of type 1}$$

and

$$\alpha \overset{s}{\sim} \beta \iff \alpha \text{ is connected to } \beta \text{ via covers of types 1 and 2}$$

is a lattice congruence of $\mathrm{Con}(\mathbf{A})$*.*

(2) *If* $\alpha \leq \beta$ *and* $\gamma \in \mathrm{Con}(\mathbf{A})$ *is any other congruence, and if the interval from* α *to* β *contains only covers of type 1, then the same is true for each of the intervals* $\gamma \wedge \alpha \leq \gamma \wedge \beta$, $\gamma \vee \alpha \leq \gamma \vee \beta$*.*

It follows that for every finite algebra $\mathbf{A}$, the sets of congruences $\overset{ss}{\sim}$-equivalent (resp. $\overset{s}{\sim}$-equivalent) to $\perp_A$ have largest elements, which we call the strongly-solvable radical $\mathrm{Rad}_u(\mathbf{A})$ and solvable radical $\mathrm{Rad}(\mathbf{A})$ of $\mathbf{A}$.

II–E. Powers and subpowers

If $\mathbf{A}$ is an algebra and I is an index set, the direct power $\mathbf{A}^I$ has its expected meaning. Elements of this power will be denoted in one of two ways:

$$\mathbf{x} = \langle x^i \rangle_{i \in I}$$

$$\mathbf{x} = a_{|I_0} \oplus b_{|I_1} \oplus c_{|\mathrm{else}}$$

where I_0 and I_1 (and any other sets which appear) are of course understood to be disjoint subsets of I.

A subalgebra

$$\mathbf{B} \leq \mathbf{A}^I$$

is a *subpower* of $\mathbf{A}$. A subpower is *subdirect* if, for each $i \in I$ and each $a \in A$ there is some $\mathbf{x} \in B$ with $x^i = a$, and *diagonal* if, for each $a \in A$, the point $\mathbf{a} = a_{|I}$ belongs to B.

If $\mathbf{B} \leq \mathbf{A}^I$, then every congruence of $\mathbf{A}^I$ restricts to a congruence of $\mathbf{B}$; when no confusion can result, we will let it be clear from context whether we are referring to $\alpha \in \mathrm{Con}(\mathbf{A}^I)$ or $\alpha \in \mathrm{Con}(\mathbf{B})$. We allow the same abuse of language for other subsets of $\mathbf{B}$.

If $\mathbf{A}$ is an algebra, $U \subset A$, and $\mathbf{B} \leq \mathbf{A}^I$, we will frequently be interested in subsets of the form $U^I \cap B$. If the meaning is clear from context, we will usually abbreviate this to U^I.

Proposition II.6. *Let* $\mathbf{A}$ *be any algebra, and let* $e \in \mathsf{Pol}_1(\mathbf{A})$ *be idempotent (that is,* $e \circ e = e$*). Then if* $U = e(A)$*, and if* $\mathbf{B} \leq \mathbf{A}^I$ *is any diagonal subpower of* $\mathbf{A}$*, then* $U^I \cap B$ *is an A-definable subset of* $\mathbf{B}$*.*

Proof. Since $\mathbf{B}$ contains the diagonal, the function $\mathbf{e} = e^I$ is realized as a polynomial of $\mathbf{B}$. $U^I \cap B$ is the set of fixed points of this polynomial. □

Indeed, for any such diagonal subpower and for each k, the map

$$\mathsf{Pol}_k(\mathbf{A}) \hookrightarrow \mathsf{Pol}_k(\mathbf{B})$$

$$f(v_1, \ldots, f_k) = t(v_1, \ldots, v_k, a_1, \ldots, a_\ell) \mapsto t(v_1, \ldots, v_k, \mathbf{a}_1, \ldots, \mathbf{a}_\ell) = f^I$$

is an embedding (of clones), which we will make continual use of.

Lemma II.7. *Let $\mathbf{A}_i$, $1 \le i \le p$ be finite algebras with trivial strongly-solvable radical. Then every*

$$\mathbf{B} \le_s \prod_i \mathbf{A}_i$$

has trivial strongly-solvable radical.

Proof. We show the contrapositive: suppose that $\perp_B \overset{1}{\prec} \alpha$ is an atom of $\mathrm{Con}(\mathbf{B})$. Then there is some projection congruence η_j such that $\alpha \vee \eta_j > \eta_j$. By Theorem II.5, since $\perp_B \overset{ss}{\sim} \alpha$, $\eta_j \overset{ss}{\sim} \alpha \vee \eta_j$; it follows that the strongly-solvable radical of $\mathbf{A}_j$ sits above $\alpha \vee \eta_j$. □

If $\mathbf{A}_1, \dots, \mathbf{A}_p, \mathbf{B}$ are as in the previous Lemma, and all belong to some finitely-decidable variety, then we can conclude (via the transfer principles) that in fact $\mathbf{B}$ has no unary-type covers anywhere in its congruence lattice. This remains true if we introduce finitely many constant symbols in such a way that each element of each A_i is named by at least one constant symbol; call these expansions $\langle \mathbf{A}_i; A_i \rangle$. Lemma II.7 implies that $\mathrm{HSP}\left(\{\langle \mathbf{A}_i; A_i \rangle\}_{i=1}^p\right)$ is modular (since all minimal sets will have empty tails), and so has Day (or Gumm) terms.

In particular, if we are considering a fixed finite $\mathbf{B} \le_s \prod_i \mathbf{A}_i$, we may introduce constant symbols for each element of B and interpret them in the $\mathbf{A}_i$ via their coordinate projections. Then $\langle \mathbf{B}; B \rangle$ has Day terms, which become Day polynomials when we reduct back out to the original language. It follows that all the nice properties of congruence-modular varieties, such as most of the theory of commutators, hold for $\mathbf{B}$.

It is an open problem whether the finite decidability of

$$\mathrm{HSP}\left(\mathbf{A}_1, \dots, \mathbf{A}_p\right)$$

implies the finite decidability of

$$\mathrm{HSP}\left(\langle \mathbf{A}_1; A_1 \rangle, \dots, \langle \mathbf{A}_p; A_p \rangle\right)$$

The best we can say is that the latter variety must be ω-structured, in the sense of [MV89].

II–F. Multi-sorted structures

In the latter portion of the paper, we will begin to deal with multi-sorted algebras. The following definitions will be of use there:

Definition II.8. Let

$$X_1 \times X_2 \times \cdots \times X_n \xrightarrow{f} Y$$

be a function. We say that f *depends essentially* on its i^{th} variable if there exist $a \ne a' \in X_i$ and

$b_j \in X_j$ ($j \neq i$) so that

$$f(b_1, b_2, \ldots, b_{i-1}, a, b_{i+1}, \ldots, b_n) \neq f(b_1, b_2, \ldots, b_{i-1}, a', b_{i+1}, \ldots, b_n)$$

(Clearly, if f depends on its i^{th} variable, it follows that $|X_i| > 1$.)

In particular, if f is a term of the (ordinary first-order) algebra $\mathbf{A}$, unless otherwise specified each X_i is A; if $\mathbf{M}$ is a multi-sorted algebra, the default assumption is that each X_i is the entire sort associated to the corresponding input variable of f.

Definition II.9. Let A be a finite set. We say that the operation $d(v_1, \ldots, v_K)$ is a *decomposition operation* on A if

- $d(A, \ldots, A) \subseteq A$;
- the action of d on A depends on all its variables;
- $d(x, \ldots, x) = x$ for all $x \in A$; and
- $$\begin{aligned} &d(d(x_{1,1}, \ldots, x_{1,K}), d(x_{2,1}, \ldots, x_{2,K}), \ldots, d(x_{K,1}, \ldots, x_{K,K})) \\ &\qquad = \\ &d(x_{1,1}, x_{2,2}, \ldots, x_{K,K}) \end{aligned} \tag{II.9.1}$$

  for all $x_{i,j} \in A$.

Typically, we will have in mind an algebraic structure on A or perhaps on some superset of A. If the operation d is a term operation (resp. polynomial operation) of the structure $\mathbf{A}$, we will call it a *decomposition term* (resp. *decomposition polynomial*).

Proposition II.10 ([MV89, Lemma 11.3])**.** *If $\mathbf{A}$ is a strongly-abelian algebra having an idempotent term $t(v_1, \ldots, v_K)$ depending essentially on all its variables, then $\mathbf{A}$ has a decomposition term of arity K.*

It follows that in such an algebra, if t is a term which depends on all its variables and such that $t(x, x, \ldots, x)$ is a permutation, then there is a decomposition term of the same arity as t.

Decomposition operators have a nice description in the case where $\mathbf{A}$ is strongly-abelian:

Proposition II.11 ([MV89, Lemma 11.4])**.** *If $\mathbf{A}$ is a finite strongly-abelian algebra and K the largest arity of a decomposition term d on $\mathbf{A}$, then there exist finite sets $A_1, \ldots, A_K$ and an isomorphism φ from $\mathbf{A}$ to a structure $\mathbf{B}$ with underlying set $A_1 \times \cdots \times A_K$ such that, if we denote*

$$\varphi(a) = \begin{matrix} a^1 \\ \vdots \\ a^K \end{matrix}$$

then

$$d^{\mathbf{B}}(\varphi(a_1), \varphi(a_2), \ldots, \varphi(a_K)) = \varphi\left(d^{\mathbf{A}}\begin{pmatrix} a_1^1 & a_2^1 & \cdots & a_K^1 \\ a_1^2 & a_2^2 & \cdots & a_K^1 \\ \vdots & \vdots & \ddots & \vdots \\ a_1^K & a_2^K & \cdots & a_K^K \end{pmatrix}\right) = \begin{matrix} a_1^1 \\ a_2^2 \\ \vdots \\ a_K^K \end{matrix}$$

In [MV89, Theorem 11.9], McKenzie and Valeriote showed that

Theorem II.12. *If* $\mathbf{A}$ *is strongly-abelian and* K *the largest arity of a decomposition term over* $\mathbf{A}$*, then any other term's depending on more than* K *variables implies that* $\mathrm{Th}(\mathcal{V})$ *and* $\mathrm{Th}_{\mathrm{fin}}(\mathcal{V})$ *are undecidable for any variety* $\mathcal{V}$ *containing* $\mathbf{A}$*.*

Our goal is to generalize this result to algebras $\mathbf{A}$ which are not themselves strongly-abelian, but do contain nontrivial strongly-abelian congruences.

Proposition II.13. *Let* $\mathbf{A}$ *be a finite algebra with a strongly-abelian congruence* τ*. Let* $C \subset A$ *be any* τ*-class; then the non-indexed algebra*

$$\mathbf{A}_{|C} = \langle C;\ \{f_{|C} \colon f \in \mathsf{Pol}(\mathbf{A}),\ f(C, C, \ldots, C) \subseteq C\}\rangle$$

is strongly-abelian.

It would be natural to search for a generalization of Theorem II.12 by looking at polynomials which restrict to decomposition operations on τ-classes; however, we found this approach to have attendant difficulties. Instead, we make the following definition:

Definition II.14. Let $\mathbf{A}$ be an algebra with a congruence α, let $t(v_1, \ldots, v_\ell, v_{\ell+1}, \ldots, v_{\ell+m})$ be a term, and let $C_0, C_1, \ldots, C_{\ell+m}$ be α-blocks such that $t(C_1, \ldots, C_{\ell+m}) \subseteq C_0$ and the action of t on the rectangle $C_1 \times \cdots \times C_{\ell+m}$ depends only on some subset of the first ℓ variables. We will call the restriction of t to the rectangle $C_1 \times \cdots \times C_\ell$ an "α-sorted term operation", and use a symbol $\check{t}$ for such a restricted operation. (Formally, we should specify what subset of the variables of t we are selecting as the domain of $\check{t}$, but this will be clear from context.)

In the course of proving Theorem D, we will need the following definitions:

Definition II.15. Let $\mathbf{A}$ be any algebra.

(1) We say that a term $t(v_1, \ldots, v_n)$ is *left-invertible* at v_i if there exists a term $r(v_0, v_{n+1}, \ldots, v_{n+k})$ such that

$$\mathbf{A} \models v_i = r(t(v_1, \ldots, v_n), v_{n+1}, \ldots, v_k)$$

(2) Likewise we call $t(v_1, \ldots, v_n)$ *right-invertible* if there exist terms $s_i(v_0, \ldots, v_\ell)$, $1 \leq i \leq n$, such that

$$\mathbf{A} \models t(s_1(v_0, \ldots, v_\ell), \ldots, s_n(v_0, \ldots, v_\ell)) = v_0$$

CHAPTER III

CONGRUENCE GEOMETRY OF THE STRONGLY-SOLVABLE RADICAL

Our goal in this section will be the proof of Theorem A.

III–A. Coherence

Before we get directly to matters of decidability, we will need some preliminary technical material of general algebraic interest.

Lemma III.1. *Let* $\mathbf{S}$ *be a subdirectly-irreducible algebra in a finitely-decidable variety with unary-type monolith* μ*. Then the centralizer of* μ*, the greatest congruence* ζ *such that* $\mathrm{C}(\zeta, \mu; \perp)$*, is strongly-solvable.*

Proof. This is proved in [IV01, Theorem 4]. □

Lemma III.2. *Let* $\mathbf{A}$ *be a finite algebra with* $\perp_A \overset{1}{\prec} \delta$*, and let* U *be* $(\perp, \delta)$*-minimal.*

(1) If $D_1, \ldots, D_k$ *are* δ*-classes, then every mapping*

$$f : D_1 \times \cdots \times D_k \to U$$

(where $f \in \mathsf{Pol}_k(\mathbf{A})$*) depends on no more than one of its variables.*

(2) (Maroti's Lemma) *If* $\delta \leq \beta$ *in* $\mathrm{Con}(\mathbf{A})$ *and* $\mathrm{C}(\beta, \delta_{|U}; \perp)$*, and* $B_1, \ldots, B_k$ *are* β*-classes, then for every mapping*

$$f : B_1 \times B_2 \times \cdots \times B_k \to U$$

(where $f \in \mathsf{Pol}_k(\mathbf{A})$*) there exists* $1 \leq j \leq k$ *so that*

$$\vec{x} \equiv_\delta \vec{y} \textit{ and } x_j = y_j \Rightarrow f(\vec{x}) = f(\vec{y})$$

Proof. The second statement is Lemma 7.2 of [IMV09]; the first statement is a special case of the second (or can be proved independently, as in [HM88, Theorem 5.6]). □

Definition III.3 ([Kea93, Definition 4.1])**.** Let $\alpha \prec \beta$ be a congruence cover of the finite algebra $\mathbf{A}$, and let $\gamma \in \mathrm{Con}(\mathbf{A})$. Let T denote the set of all (α, β)-traces in $\mathbf{A}$. We say that (α, β) is γ*-coherent* if

$$\mathop{\&}_{N \in T} \mathrm{C}(\gamma, \beta_{|N}; \alpha) \quad \Longrightarrow \quad \mathrm{C}(\gamma, \beta; \alpha)$$

If $\alpha = \perp$ then we will say that β is γ-coherent. Note, that since all (α, β)-traces are polynomially isomorphic, $\mathrm{C}(\gamma, \beta_{|N}; \alpha)$ holds for all $N \in T$ iff it holds for any such N.

Lemma III.4. *Let* $\mathbf{A}$ *be any finite algebra with congruences* $\perp \overset{1}{\prec} \delta$ *and* $\alpha \overset{3}{\prec} \beta$*, such that* $\beta = \mathrm{Cg}(\langle 0,1\rangle)$ *for some (hence any)* (α,β)*-trace* $\{0,1\}$*. Assume further that* $\neg\mathrm{C}(\beta,\delta;\perp)$*. Then there exists a polynomial* $p(x,y) = p(x,p(x,y))$ *taking values in some* $(\perp,\delta)$*-minimal set* U*, so that*

(1) If δ *is* β*-coherent, then* $p(0,y)$ *collapses traces to points and* $p(1,u) = u$ *for all* $u \in U$*;*

(2) If δ *is* β*-incoherent, then* $p(0,u) = u = p(1,u)$ *for all* $u \in U$*, but for some* $c \in U$*,* $c \equiv_\delta d$*,* $d \notin U$*,*

$$p(0,c) = p(0,d)$$
$$\textit{but}$$
$$p(1,c) \neq p(1,d)$$

witnesses the failure of centralization.

Proof. Suppose first that δ is β-coherent. Then for some $(\perp,\delta)$-trace N included in some minimal set U, we have $\neg\mathrm{C}(\beta,N;\perp)$. Since β is generated by $\{0,1\}$, $\mathrm{C}(\{0,1\},N;\perp)$ must already be false.

Choose a witnessing package

$$t(0,\vec{c}) = t(0,\vec{d})$$
$$\text{but}$$
$$t(1,\vec{c}) \neq t(1,\vec{d})$$

where we may choose t so that its range lies entirely in U. The polynomial mapping $t(1,\vec{y})$ is essentially unary as a mapping from $\vec{N}$ into U; say it depends on y_1, and let $p(x,y) = t(x,y,c_2,c_3,\dots)$. Then $p(1,c_1) = t(1,\vec{c}) \neq t(1,\vec{d}) = p(1,d_1)$ while $p(0,c_1) = p(0,d_1)$. Iterating p in the second variable if necessary, we get a polynomial satisfying the Lemma.

The other case requires a bit more work.

Assume now that for all traces N, we have $\mathrm{C}(\beta,N;\perp)$. As in the first case, $\neg\mathrm{C}(\beta,\delta;\perp)$ implies that $\neg\mathrm{C}(\{0,1\},\delta;\perp)$ already. Take a witnessing package

$$t(0,\vec{c}) = t(0,\vec{d})$$
$$\text{but}$$
$$t(1,\vec{c}) \neq t(1,\vec{d})$$

where we may assume that the image of t is contained in some $(\perp,\delta)$-minimal set U_0. The map $t(0,\vec{y}) : c_1/\delta \times c_2/\delta \times \cdots \to U_0$ depends only one one variable, say y_{k_0}, and likewise $t(1,\vec{y})$ on y_{k_1}.

Claim III.4.1. $k_0 = k_1$

Suppose the Claim were false. Let $q(x,y) = t(x,c_1,\dots,c_{k_1-1},y,c_{k_1+1},\dots)$. Then $q(0,c_{k_1}) = q(0,y)$ for all $y \equiv_\delta c_1$.

Now, since $c_{k_1} \equiv_\delta d_{k_1}$, there exists a sequence

$$c_{k_1} = a_0, a_1, \ldots, a_\ell = d_{k_1} \tag{III.4.2}$$

where each pair $\{a_i, a_{i+1}\}$ belong to a $(\perp, \delta)$-trace N_i $(i < \ell)$ included in a minimal set $U_i = e_i(A)$. Since $q(1, a_0) \neq q(1, a_\ell)$, there must exist some $i < \ell$ such that $q(1, a_i) \neq q(1, a_{i+1})$. But we have already seen that $q(0, a_i) = q(0, a_{i+1})$, contradicting $\mathrm{C}(\{0,1\}, \delta_{|N_i}; \perp)$. This proves the Claim, and we may set $k := k_0 = k_1$. $\dashv_{\text{III.4.1}}$

Let $a_0, a_1, \ldots, a_\ell$ be the sequence defined in (III.4.2); our assumption that $\{0,1\}$ centralizes N_i means that for each $i < \ell$, $q(0, y)$ is injective on N_i iff $p(1, y)$ is.

Let i be the first index for which $q(1, a_i) \neq q(1, a_{i+1})$; then

$$q(0, d_k) = q(0, c_k) = q(0, a_0) = q(0, a_1) = \ldots = q(0, a_i)$$

but

$$q(1, d_k) \neq q(1, c_k) = q(1, a_0) = q(1, a_1) = \ldots = q(1, a_i)$$

Then with $c = a_i$, $d = d_k$, $U = U_i$, and $p(v_0, v_1)$ equalling an iterate of $e_i \circ q(v_0, v_1)$ such that $p(x, p(x, y)) = p(x, y)$ for all $x, y \in A$, the conclusions of the lemma are satisfied. $\square$

Lemma III.5. *If* $\perp_{\mathbf{A}} \overset{1}{\prec} \delta$, $\alpha \overset{3}{\prec} \beta$, $K = \{0,1\}$, *and* U *are as in the statement of Lemma III.4, and* $N \subset U$ *is a trace, then at least one of* $\mathrm{C}(K, N; \perp)$ *and* $\mathrm{C}(N, K; \perp)$ *must fail. In both cases, the failure is witnessed by a binary polynomial which takes* $K \times N$ *into* U.

Proof. In the case where δ is β-coherent, the polynomial p found in that Lemma witnesses $\neg\mathrm{C}(K, N; \perp)$.

So let $\mathrm{C}(\beta, N; \perp)$ for all $(\perp, \delta)$-traces N, and fix witnesses

$$c = p(0, c) = p(0, d)$$

but

$$c = p(1, c) \neq p(1, d)$$

where $c \in U$, $c \equiv_\delta d \notin U$, and the range of p is contained in U. We aim to show that $\mathrm{C}(N, K; \perp)$ fails, and that its failure is witnessed by a binary polynomial of the claimed kind.

Let

$$c = a_0, a_1, \ldots, a_{\ell-1}, a_\ell = d$$

be a walk from c to d through traces (see the discussion following Equation (III.4.2)). Since

$$p(0, a_i) \equiv_\delta p(0, a_0) = c = p(1, a_0) \equiv_\delta p(1, a_i)$$

for all $i \leq \ell$, we know that $p(K, c/\delta) \subset N$. Now let $\{a_j, a_{j+1}\} \subset N_j$ be the first step where

$$p(0, a_j) = p(1, a_j)$$
$$\text{but}$$
$$p(0, a_{j+1}) \neq p(1, a_{j+1})$$

By hypothesis, $j > 0$. It follows that at least one, and hence both, of $p(0, y)$ and $p(1, y)$ are polynomial isomorphisms from N_j to N. Let $q \in \mathsf{Pol}_1(\mathbf{A})$ be the inverse isomorphism to $p(0, y)$, where $q(a) = a_j$ and $q(a') = a_{j+1}$. Then

$$p(0, q(a)) = p(0, a_j) = p(1, a_j) = p(1, q(a))$$
$$\text{but}$$
$$p(0, q(a')) = p(0, a_{j+1}) \neq p(1, a_{j+1}) = p(1, q(a'))$$

so that $p(x, q(y))$ witnesses $\neg\mathrm{C}(N, K; \perp)$ as required. $\square$

III–B. Proof of Theorem A

Lemma III.6. *Let* $\mathbf{A}$ *be a finite algebra,* $\perp \overset{1}{\prec} \delta$ *and* $\alpha \overset{3}{\prec} \beta$*, and let* $K = \{0, 1\}$ *be* (α, β)*-minimal, where* $\beta = \mathrm{Cg}(\langle 0, 1\rangle)$*. If* $\neg\mathrm{C}(K, \delta; \perp)$ *and* $\mathrm{C}(\delta, K; \perp)$*, then* $\mathrm{HSP}\,(\mathbf{A})$ *is hereditarily finitely undecidable.*

In other words, the centralizer of a boolean neighborhood must be disjoint from any of the unary-type atoms (or at least those which that neighborhood does not itself centralize), if $\mathbf{A}$ is to live in a finitely-decidable variety.

Proof. Fix a $(\perp, \mu)$-minimal set U. By Lemma III.5, for any $(\perp, \delta)$-trace $N \subset U$, at least one of $\mathrm{C}(K, N; \perp)$ or $\mathrm{C}(N, K; \perp)$ must fail. But if $\neg\mathrm{C}(N, K; \perp)$ then $\neg\mathrm{C}(\delta, K; \perp)$, contrary to the assumptions of the Lemma.

Hence $\neg\mathrm{C}(K, N; \perp)$. Choose a witnessing package

$$q(0, c) = q(0, d)$$
$$\text{but}$$
$$q(1, c) \neq q(1, d)$$

Without loss of generality, we can assume that $q(1, u) = u$ for all $u \in U$.

Our plan is to semantically interpret the class of graphs with at least three vertices into diagonal subpowers of $\mathbf{S}$. So let $\mathbb{G} = \langle V, E\rangle$ be such a graph, and let $I = V \times \{+, -\} = V^{\pm}$. Define

$\mathbf{D} = \mathbf{D}(\mathbb{G}) \leq \mathbf{A}^I$ to be generated by the diagonal together with the points

$$\begin{aligned}
\chi^\beta_v &:= 1_{|\{v^+,v^-\}} \oplus 0_{|\text{else}} \qquad &(\text{all } v \in V)\\
\chi^\delta_e &:= d_{|\{v^+,w^+\}} \oplus c_{|\text{else}} \qquad &(\text{all } e = \langle v,w\rangle \in E)\\
\chi^\delta_{V^+} &:= d_{|V^+} \oplus c_{|V^-}
\end{aligned}$$

Let $\vec{\chi}^\beta$ and $\vec{\chi}^\delta$ enumerate the respective sets of generators.

Observe that there cannot be any nonconstant polynomial map from N to $\{0,1\}$. This implies that $D \cap \{0,1\}^I$ consists of all points which are constant on each set $\{v^+,v^-\}$; in other words, $\mathbf{D}_{|\{0,1\}^I}$ is canonically isomorphic to the boolean algebra $\mathbf{2}^V$. This subset is definable (Proposition II.6), as is its set of atoms $\{\chi^\beta_v \colon v \in V\}$; by abuse of language, we will allow ourselves to quantify over these atoms by saying things like "there exists a vertex χ^β_v..."

Claim III.6.1. The set of those $\mathbf{x} \in D$ of the form $d_{|\{w_1^+,w_2^+\}} \oplus c_{|\text{else}}$ for two distinct vertices $w_1, w_2 \in V$ is definable (using the parameter $\chi^\delta_{V^+}$).

It is sufficient to show that for $\mathbf{x} \in U^I \cap D$, $\mathbf{x} = d_{|\{w_1^+,w_2^+\}} \oplus c_{|\text{else}}$ iff

$$q\left(\chi^\beta_{w_1} + \chi^\beta_{w_2}, \mathbf{x}\right) = q\left(\chi^\beta_{w_1} + \chi^\beta_{w_2}, \chi^\delta_{V^+}\right) \tag{III.6.2}$$

and

$$q\left((\chi^\beta_{w_1} + \chi^\beta_{w_2})', \mathbf{x}\right) = q\left((\chi^\beta_{w_1} + \chi^\beta_{w_2})', c\right) \tag{III.6.3}$$

(where $+$ is boolean join and $'$ is boolean complement).

The direction ($\Rightarrow$) is a straightforward computation. For the reverse direction,

$$i \in \{w_1^+, w_2^+\} \quad \Longrightarrow \quad x^i = q(1,x^i) = q(1,(\chi^\delta_{V^+})^i) = d$$

from equation (III.6.2), and similarly

$$i \in \{w_1^-, w_2^-\} \quad \Longrightarrow \quad x^i = q(1,x^i) = q(1,(\chi^\delta_{V^+})^i) = c$$

while equation (III.6.3) yields

$$i \notin \{w_1^\pm, w_2^\pm\} \quad \Longrightarrow \quad x^i = p(1,x^i) = p(1,c) = c$$

The proof of the claim is then accomplished by existentially quantifying $\chi^\beta_{w_1}, \chi^\beta_{w_2}$. $\dashv_{\text{III.6.1}}$

Claim III.6.4. If $\mathbf{x} = d_{|\{w_1^+,w_2^+\}} \oplus c_{|\text{else}} \in D$ then $w_1 \overset{E}{—} w_2$.

To see this, let $\mathbf{x} = d_{|\{w_1^+,w_2^+\}} \oplus c_{|\text{else}} = t(\vec{\chi}^\beta, \vec{\chi}^\delta) \in D$ for some polynomial $t \in \mathsf{Pol}_{|V|+|E|+1}(\mathbf{A})$. Without loss of generality, t's image is contained in U. By inspecting the v^- coordinates, we see that for any $v \in V$

$$t(0,\ldots,0,1,0,\ldots,c,\ldots,c) = c$$

(the 1 occurring in the v^{th} place). Fix any $w \in V$; then

$$x^{v^-} = t\left((\vec{\chi}^{\beta})^{w^-}, (\vec{\chi}^{\delta})^{v^-}\right)$$

Moreover, since $(\vec{\chi}^{\beta})^{v^+} = (\vec{\chi}^{\beta})^{v^-}$ for all v and $\mathrm{C}(\delta, \{0,1\}; \perp)$, one has

$$t\left((\vec{\chi}^{\beta})^{v^+}, (\vec{\chi}^{\delta})^{v^-}\right) = t\left((\vec{\chi}^{\beta})^{v^-}, (\vec{\chi}^{\delta})^{v^-}\right) = t\left((\vec{\chi}^{\beta})^{w^-}, (\vec{\chi}^{\delta})^{v^-}\right)$$
$$\Downarrow$$
$$t\left((\vec{\chi}^{\beta})^{v^+}, (\vec{\chi}^{\delta})^{v^+}\right) = t\left((\vec{\chi}^{\beta})^{w^-}, (\vec{\chi}^{\delta})^{v^+}\right)$$

In other words,

$$x^{i} = t\left((\vec{\chi}^{\beta})^{w^-}, (\vec{\chi}^{\delta})^{i}\right)$$

for all $i \in I$.

But as a polynomial on U, $t((\vec{\chi}^{\beta})^{w^-}, \vec{y})$ depends only on one variable, say $t((\vec{\chi}^{\beta})^{w^-}, \vec{y}) = f(y_k)$, with y_k corresponding to a generator $\chi^{\delta}_k \in \vec{\chi}^{\delta}$; k is either an edge of $\mathbb{G}$ or V^+. Since c, d are taken from the same trace and f does not collapse traces to points, we must have that $\mathbf{x} = f(\chi^{\delta}_k)$ and $x^i = x^j$ iff $(\chi^{\delta}_k)^i = (\chi^{\delta}_e)^j$ for all $i, j \in I$; since $|V| > 2$ and $\mathbf{x}$ has d at only has two coordinates (out of at least six), k must be the edge $\langle w_1, w_2 \rangle$. $\dashv_{\text{III.6.4}}$

We can now complete the semantic interpretation: V is defined as the atoms of $\{0,1\}^I \cap D$, and $v \overset{E}{—} w$ iff there exists $\mathbf{x}$ as in Claim III.6.1 such that χ^{β}_v and χ^{β}_w are the two atoms witnessing the truth of the formula in that Claim. □

Lemma III.7. *If* $\mathbf{S}$ *is a finite, subdirectly-irreducible algebra with unary-type monolith, and the strongly-solvable radical of* $\mathbf{S}$ *is incomparable to some congruence of* $\mathbf{S}$*, then* HSP $(\mathbf{S})$ *is hereditarily finitely undecidable.*

Proof. Let $\mathbf{S}$ be subdirectly-irreducible, with unary-type monolith μ; let β be incomparable to the strongly-solvable radical σ. Without loss of generality (see Fact II.4), $\operatorname{typ}\{\mathbf{S}\} = \{1,3\}$, and some lower cover of β is (strictly) below σ. Choose $\beta \wedge \sigma \overset{1}{\prec} \alpha \leq \sigma$; clearly $\beta \wedge \sigma = \alpha \wedge \beta =: \alpha\beta \overset{3}{\prec} \beta$.

Choose an $(\alpha\beta, \beta)$-minimal set, which we may take without loss of generality to be polynomially isomorphic to the two-element boolean algebra $\{0,1\}$; similarly without loss of generality, $\beta = \mathrm{Cg}(\langle 0,1 \rangle)$; also choose a $(\alpha\beta, \alpha)$-minimal set U containing elements $c \equiv_{\alpha \setminus \alpha\beta} d$.

Now, by Lemma III.1, the centralizer of μ is solvable; hence $\neg\mathrm{C}(\{0,1\}, \mu; \perp)$. By Lemma III.6, we may assume that the centralizer of $\{0,1\}$ is the trivial congruence: for any $a_1 \neq a_2$ in $\mathbf{S}$, there exists a polynomial $t(x, \vec{y})$ and tuples $\vec{b}_0, \vec{b}_1$ from $\{0, 1\}$ so that

$$t(a_1, \vec{b}_0) = t(a_1, \vec{b}_1)$$
$$\text{but}$$
$$t(a_2, \vec{b}_0) \neq t(a_2, \vec{b}_1)$$

Since $\mathbf{S}_{|\{0,1\}}$ is a boolean algebra, the discussion after equation (II.2.1) shows that we can transform this package into one using a binary polynomial:

$$s(a_1, 0) = s(a_1, 1)$$
$$\text{but}$$
$$s(a_2, 0) \neq s(a_2, 1)$$

witnessing that $\{a_1, a_2\}$ does not centralize $\{0, 1\}$.

Our strategy is to interpret the class of graphs with at least five vertices into HSP$(\mathbf{S})$, so let $\mathbb{G} = \langle I, E\rangle$ be any graph. Define $\mathbf{D} = \mathbf{D}(\mathbb{G}) \leq \mathbf{S}^I$ to be the subalgebra generated by the constants together with all points

$$\chi_i^\beta := 1_{|i} \oplus 0_{|\text{else}} \qquad (i \in I)$$

and

$$\chi_e^\alpha := d_{|\{i,j\}} \oplus c_{|\text{else}} \qquad (e = \{i, j\} \in E)$$

By the usual arguments, $\{0, 1\}^I \subseteq D$ is a definable subset, as is the set of its atoms.

Let χ_i^β be any atom in $\{0, 1\}^I$. Let $\mathbf{y}, \mathbf{z}$ be any elements of $\mathbf{D}$. Then

$$p(0, \mathbf{y}) = p(\chi_i^\beta, \mathbf{y})$$
$$\Updownarrow$$
$$p(0, \mathbf{z}) = p(\chi_i^\beta, \mathbf{z})$$

for all $p \in \mathsf{Pol}_2(\mathbf{S})$ if and only if y^i and z^i are congruent modulo the centralizer of $\{0, 1\}$, i.e. are equal. But $\mathbf{S}$ only has finitely many binary polynomial operations; hence the above condition is a first-order property $\Phi(\chi_i^\beta, \mathbf{y}, \mathbf{z})$: we have proved

Claim III.7.1. If $s \in S$, $i \in I$, $\mathbf{y} \in D$ then $y^i = s$ iff $\Phi(\chi_i^\beta, \mathbf{y}, s)$.

Or in plainer English: $\mathbf{D}$ knows its own product structure. $\dashv_{\text{III.7.1}}$

In particular: the set of those $\mathbf{x} \in U^I \cap D$ of the form $d_{|\{i_0,i_1\}} \oplus c_{|\text{else}}$ for precisely two vertices i_0, i_1, is a definable subset. The generators χ_e^α belong to this set.

Claim III.7.2. If $i_0 \neq i_1$ and $\mathbf{x} = d_{|\{i_0,i_1\}} \oplus c_{|\text{else}} \in D$ then $i_0 \overset{E}{—} i_1$.

So let

$$\mathbf{x} = d_{|\{i_0,i_1\}} \oplus c_{|\text{else}} = f(\vec{\chi}^\beta, \vec{\chi}^\alpha)$$

belong to D, where $f \in \mathsf{Pol}_{|I|+|E|}(\mathbf{S})$ takes values in U and $\vec{\chi}^{\alpha}$, $\vec{\chi}^{\beta}$ enumerate the two sets of generators.

Let $j \in I$ be any vertex. Then

$$x^j = f\left((\vec{\chi}^{\beta})^j, (\vec{\chi}^{\alpha})^j\right) \equiv_\beta f\left((\vec{\chi}^{\beta})^{i_0}, (\vec{\chi}^{\alpha})^j\right) \equiv_\alpha f\left((\vec{\chi}^{\beta})^{i_0}, (\vec{\chi}^{\alpha})^{i_0}\right) = x^{i_0} \equiv_\alpha x^j$$

Hence

$$\mathbf{x} \equiv_{\alpha\beta} f\left((\vec{\chi}^{\beta})^{i_0}, \vec{\chi}^{\alpha}\right)$$

But considered as a mapping from α-classes into U, $f((\vec{\chi}^{\beta})^{i_0}, \vec{v})$ depends modulo $\alpha\beta$ on no more than one of the edge-variables, say $f((\vec{\chi}^{\beta})^{i_0}, \vec{v}) = g(v_e)$ for some $e = \{j_0, j_1\} \in E$; since $\mathbf{x}$ is not constant modulo $\alpha\beta$, g cannot collapse traces to points, implying that $\mathbf{x} = g(\chi_e^{\beta})$ has the same $\alpha\beta$-equivalence pattern as χ_e^{α}. The two equal coordinates of $\mathbf{x}$ must match two equal coordinates of χ_e^{α} such that all other coordinates have a different value; since $|V| > 4$, the only set of such coordinates is $\{j_0, j_1\}$; but this implies $\mathbf{x} = \chi_e^{\alpha}$, as desired. $\dashv_{\text{III.7.2}}$

We have shown that we can definably recover the edge relation of $\mathbb{G}$ on a definable set in bijection with the vertex set of $\mathbb{G}$. $\square$

The investigations of *congruence modular* finitely-decidable varieties identified quite early how constrained the congruence geometry of such varieties must be. In particular, it was discovered that the congruences above the solvable radical of a subdirectly-irreducible algebra in such a variety were forced to be linearly ordered. Lemma III.7 allows us to remove the hypothesis of modularity:

Lemma III.8. *Let* $\mathbf{S}$ *be a finite subdirectly-irreducible algebra with unary-type monolith. If the congruence interval above the solvable radical of* $\mathbf{S}$ *is not linearly ordered, then* $\mathrm{HSP}(\mathbf{S})$ *is hereditarily finitely undecidable.*

Proof. Due to the transfer prinicples (see Fact II.4), we already know that $\mathrm{typ}\{\mathbf{S}\} \subset \{1, 3\}$; without loss of generality, the solvable radical $\mathrm{Rad}(\mathbf{S}) < \top_S$ and every cover above $\mathrm{Rad}(\mathbf{S})$ has boolean type. If $\mathrm{Rad}(\mathbf{S})$ were to have just one upper cover, then $\mathbf{S}/\mathrm{Rad}(\mathbf{S})$ would be subdirectly-irreducible with boolean monolith; Idziak's characterization ([Idz97]) implies then the whole interval $[\mathrm{Rad}(\mathbf{S}), \top]$ would be a chain. Hence it suffices to show that the radical having at least two upper covers α_0, α_1 leads to a contradiction.

Theorem III.7 implies that every subcover of $\mathrm{Rad}(\mathbf{S})$ is meet-irreducible, so without loss of generality (by passing to a quotient by such a subcover) we may assume that $\bot \overset{1}{\prec} \mathrm{Rad}(\mathbf{S}) =: \mu$. Let $K_a = \{0_a, 1_a\}$ be respectively (μ, α_a)-minimal sets ($a \in \{0, 1\}$).

We know that $\neg\mathrm{C}(K_a, \mu; \bot)$ for $a = 0, 1$, since each of these sets generate a congruence above the centralizer of μ. By Lemma III.6, we may also assume that $\neg\mathrm{C}(\mu, K_a; \bot)$. Let

$$p_0(c, 0_0) = p_0(c, 1_0)$$

but

$$p_0(d, 0_0) \neq p_0(d, 1_0)$$

witness this latter failure. Observe that $p_0(d, 0_0) \equiv_\mu p_0(d, 1_0)$; hence there exists $q \in \mathsf{Pol}_1(\mathbf{S})$ taking K_0 injectively into some $(\perp, \mu)$-trace N. Since $\mu = \mathrm{Cg}_{\mathbf{S}}(\langle q(0_0), q(1_0)\rangle)$, we must have $\neg\mathrm{C}(\{q(0_0), q(1_0)\}, K_1; \perp)$. Choose a witnessing package

$$p_1(q(0_0), 0_1) = p_1(q(0_0), 1_1)$$
$$\text{but}$$
$$p_1(q(1_0), 0_1) \neq p_1(q(1_0), 1_1)$$

Our strategy is to interpret the undecidable class $\mathcal{E}_2$ (see page 7) into the diagonal subpowers of $\mathbf{S}$. So let $\mathbf{E} = \langle I; R_0, R_1\rangle \models \mathcal{E}_2$, and define a diagonal subpower $\mathbf{D} = \mathbf{D}(\mathbf{E}) \leq \mathbf{S}^I$ as the subalgebra consisting of all $\mathbf{x} \in \mathbf{S}^I$ such that $\mathbf{x}$ is α_0-constant on each block of R_1 and α_1-constant on each block of R_0. Note that, since $0_0, 1_0$ are α_0-congruent but not α_1, a point $\mathbf{x} \in K_0^I$ belongs to $\mathbf{D}$ iff it is constant on each R_0-block. We conclude that $\mathbf{D}_{|K_0}$ is canonically isomorphic to the boolean algebra $\mathbf{2}^{I/R_0}$; the corresponding facts hold mutatis mutandis for K_1^I. Furthermore, these two subsets are uniformly definable (by Lemma II.6). Let $\mathrm{AT}_a(v)$ be a formula asserting that v is an atom of the boolean algebra $\mathbf{D}_{|K_a}$, and let H be the (definable) set of pairs $\langle \mathbf{y}, \mathbf{z}\rangle$ such that $\mathbf{y} = 1_{0|B_y} \oplus 0_{0|\text{else}}$ is a K_0-atom and $\mathbf{z} = 1_{1|B_z} \oplus 0_{1|\text{else}}$ is a K_1-atom.

Now, for each pair $\langle \mathbf{y}, \mathbf{z}\rangle \in H$, the blocks coded by the two points are either empty or share one $i \in I$. Write $\mathbf{y} \bowtie \mathbf{z}$ if the intersection is nonempty. It suffices to show that the relation $\mathbf{y} \bowtie \mathbf{z}$ is definable. Why is this so? Since $R_0 \cap R_1$ is trivial, every $i \in I$ corresponds canonically to exactly one $\langle \mathbf{y}_i, \mathbf{z}_i\rangle \in H$, namely $\mathbf{y}_i = (1_0)_{|i/R_0} \oplus (0_0)_{|\text{else}}$ and $\mathbf{z}_i = (1_1)_{|i/R_1} \oplus (0_1)_{|\text{else}}$. These two points are $\bowtie$-related by construction. But if $\bowtie$ is definable, the structure $\mathbf{E}$ can be recovered on the underlying set $\bowtie = \{\langle \mathbf{y}_i, \mathbf{z}_i\rangle : i \in I\}$ using the first-order theory of $\mathbf{D}$, since $\langle i, j\rangle \in R_0$ (resp R_1) iff $\mathbf{y}_i = \mathbf{y}_j$ (resp $\mathbf{z}_i = \mathbf{z}_j$).

To this end, observe: if $\langle \mathbf{y}, \mathbf{z}\rangle \in H$ and $i \in I$,

$$p_1(q(y^i), z^i) \neq p_1(q(y^i), 0_1) \iff z^i \neq 0_1 \text{ and } y^i \neq 0_0$$

It follows that

$$\begin{gathered} p_1(q(\mathbf{y}), \mathbf{z}) \neq p_1(q(\mathbf{y}), 0_1) \\ \Updownarrow \\ p_1(q(y^i), z^i) \neq p_1(q(y^i), 0_1) \text{ for some } i \in I \\ \Updownarrow \\ y^i = 1_0 \text{ and } z^i = 1_1 \text{ for some } i \in I \\ \Updownarrow \\ \mathbf{y} \bowtie \mathbf{z} \end{gathered}$$

□

Proof of Theorem A. Let $\mathbf{S}$ be a finite, subdirectly-irreducible algebra with unary-type monolith, and belonging to some finitely-decidable variety $\mathcal{V}$. If $\mathrm{Rad}_u(\mathbf{S}) = \top_S$ then we are done. If not, then Lemma III.7 tells us that $\mathrm{Rad}_u(\mathbf{S})$ is comparable to every congruence of $\mathbf{S}$. The transfer principles tell us that the affine tame-congruence type does not appear in $\mathrm{Con}(\mathbf{S})$; since $\mathcal{V}$ is finitely-decidable, it follows that every upper cover of $\mathrm{Rad}_u(\mathbf{S})$ is of the boolean type. Lemma III.8 now implies that the radical has only one upper cover. Since $\mathbf{S}/\mathrm{Rad}_u(\mathbf{S})$ is subdirectly-irreducible with boolean monolith and belongs to the finitely-decidable variety $\mathcal{V}$, its congruence lattice is a chain of boolean-type covers. $\square$

CHAPTER IV

STRONGLY-SOLVABLE IMPLIES STRONGLY ABELIAN

The goal of this chapter is to prove Theorem B. The arguments in this chapter are highly technical; we have not at this time found any way to reduce their complexity. However, at a high level, each of the main semantic interpretations found below can be read as asserting some kind of "sparseness" in subalgebra or congruence generation resulting from assumptions of strong solvability.

Denis Osin is fond of saying that group theory is infinitely distorted in mathematics, in the sense that there are theorems about groups whose shortest purely group-theoretic proof requires heroic strength of mind to read (never mind to discover), but whose proofs in The Book pass through other seemingly unrelated fields of mathematics. It may well be the case that something similar holds here.

IV–A. The radical centralizes minimal sets

The semantic interpretations constructed in this chapter (and following ones) depend on our ability to define the strongly-solvable radical uniformly in a variety. The conclusions of the following lemma can be shown to hold for either of the solvable radical, or the strongly-solvable radical, of any finite algebra $\mathbf{A}$; however, the proof of this more general theorem is no more enlightening for our purposes, so we omit it.

Lemma IV.1. *If $\mathbf{A}$ is any finite algebra in a finitely-decidable variety with strongly-solvable radical σ, there exists a first-order formula with parameters from A which defines the congruence σ^I/Θ, uniformly for all $\mathbf{D}/\Theta$, where I is any index set, $\Delta \leq \mathbf{D} \leq \mathbf{A}^I$ is any diagonal subpower, and $\Theta \leq \sigma^I \cap \mathbf{D} \in \mathrm{Con}(\mathbf{D})$.*

Proof. The argument comes from the theory of snags (see [HM88] Chapter 7). Let $E(\mathbf{A})$ denote the collection of all idempotent polynomials with nontrivial range, and for each $e \in E(\mathbf{A})$ choose $p \in \mathsf{Pol}_3(\mathbf{A})$ which is Malcev on the image of e if any such polynomial exists; if none, then let p be second projection. Then we have that a pair $\langle x, y\rangle$ fails to belong to σ iff there is a congruence cover $\alpha \overset{2,3}{\prec} \beta$ below $\mathrm{Cg}(\langle x, y\rangle)$ iff the following first-order formula is satisfied:

$$\bigvee_{e\in E(\mathbf{A})} \bigvee_{f\in \mathsf{Pol}_1(\mathbf{A})} ef(y) = p(ef(y), ef(x), ef(x)) = p(ef(x), ef(x), ef(y))$$
$$\neq p(ef(x), ef(x), ef(x)) = ef(x)$$

The formula is clearly false if every cover below $\langle x, y\rangle$ has type 1, while a cover of boolean or affine type will guarantee the formula's truth, since the minimal sets of that cover have empty tails and hence Malcev polynomials. This proves that the indicated formula defines [the complement of] σ in $\mathbf{A}^1$, and its truth is preserved by factoring out by congruences under σ.

Now since the defining formula is quantifier-free, it is preseved in subpowers. Finally, if $\mathbf{x} \equiv_{\sigma^I} \mathbf{y}$, $e \in E(\mathbf{A})$, $p(v_1, v_2, v_3) = v_2$ and $f \in \mathsf{Pol}_1(\mathbf{A})$,

$$p(ef(\mathbf{y}), ef(\mathbf{x}), ef(\mathbf{x})) = p(ef(\mathbf{x}), ef(\mathbf{x}), ef(\mathbf{y})) = p(ef(\mathbf{x}), ef(\mathbf{x}), ef(\mathbf{x})) = ef(\mathbf{x})$$

which is preserved under factoring out Θ. On the other hand, if $x^i \not\equiv_\sigma y^i$, then the polynomials which witness

$$ef(y^i) \equiv_\theta p(ef(y^i), ef(x^i), ef(x^i)) \not\equiv_\theta ef(x^i)$$

(θ being the projection of Θ into the ith coordinate) also witness it in $\mathbf{D}$. □

Lemma IV.2. *Let $\mathbf{S}$ be a subdirectly-irreducible algebra with unary-type monolith μ and strongly-solvable radical σ which is abelian over μ but not over $\bot_S$. Let $U = e(S)$ be any $(\bot_S, \mu)$-minimal set. If* $\mathrm{C}(\sigma, \mu_{|U}; \bot)$ *fails in* $\mathbf{S}$, *then* HSP $(\mathbf{S})$ *is hereditarily finitely undecidable.*

Proof. Since $\mathrm{C}(\mu, \mu_{|U}; \bot)$ always holds, we may climb the congruence lattice until we get a cover $\mu \leq \theta_0 \overset{1}{\prec} \theta_1 \leq \sigma$ such that $\mathrm{C}(\theta_0, \mu_{|U}; \bot)$ holds and $\mathrm{C}(\theta_1, \mu_{|U}; \bot)$ does not. Fix a (θ_0, θ_1)-minimal set $U' = e'(S)$ with trace N' containing θ_0-inequivalent elements a_0, a_1. Since these elements generate θ_1 over θ_0, already $\neg\mathrm{C}(\mathrm{Cg}(\langle a_0, a_1\rangle), \mu_{|U}; \bot)$, and we may take a witnessing package

$$t(a_0, \vec{b}_0) = t(a_0, \vec{b}_1)$$

$$\text{but}$$

$$t(a_1, \vec{b}_0) \neq t(a_1, \vec{b}_1)$$

There is no loss of generality in assuming that the image of t is contained in U.

Since μ is strongly-abelian, we may assume that $\vec{b}_0$ and $\vec{b}_1$ differ only in one place (say the first), so that for $q(v_1, v_2) = t(v_1, v_2, b^2, \ldots,)$, the polynomial $q(a_0, x)$ is constant on $\mu_{|U}$-blocks while the polynomial $q(a_1, x)$ permutes U. (Observe that $q(x, y) \in U$ for any $x, y \in S$.) Of course we may by iterating q guarantee that for each $u' \in U'$, the operation $q(u', x)$ is idempotent. The same argument shows that for each $u' \in U'$, $q(u', x)$ is either the identity on U (in which case we call u' *permutational*) or else squashes each μ-block of U to a point (at which we call u' *collapsing*). Since $\mathrm{C}(\theta_0, \mu_{|U}; \bot)$, these two properties are invariant under θ_0-congruence.

Let $N \subseteq U$ be any trace; we have that $q(a_0, N) = m_0$ for some $m_0 \in U$. In fact, since σ is abelian over μ, we have

$$q(a_0, m_0) = q(a_1, m_0)$$

$$\Downarrow$$

$$m_0 = q(a_0, u) \equiv_\mu q(a_1, u) = u \qquad \text{for any } u \in N$$

and thus $m_0 \in N$; more generally, we have that the polynomial $v_1 \mapsto q(a_0, v_1)$ retracts each trace down to one of its points. Since N was a trace, there exists some $m_1 \neq m_0$ in N, which we fix for future use.

We want to semantically embed graphs into the diagonal subpowers of $\mathbf{S}$, so let $\mathbb{G} = \langle V, E \rangle$ be a graph. Our index set I will equal $V \sqcup \{\infty\}$. Our subpower $\mathbf{S}[\mathbb{G}]$ will be the subalgebra of $\mathbf{S}^I$ generated by the diagonal together with

- for each vertex $v \in V$, the element

$$\mathbf{g}_v = a_{1|\{v,\infty\}} \oplus a_{0|\text{else}}$$

- for each edge $\{v_1, v_2\} \in E$, the element

$$\mathbf{g}_{v_1 v_2} = a_{1|\{v_1,v_2,\infty\}} \oplus a_{0|\text{else}}$$

  and

- the element

$$\chi_\infty = m_{0|V} \oplus m_{1|\infty}$$

Recall our notational convention (page 11) that for $s \in S$ we will use a boldface $\mathbf{s}$ to denote the corresponding diagonal element; let $\vec{\mathbf{s}}$ be a fixed enumeration of these diagonal elements. Observe that each generator, and hence every element of $\mathbf{S}[\mathbb{G}]$, is constant modulo θ_1; and that χ_∞ is also constant mod θ_0 (indeed, mod μ).

Claim IV.2.1. Every element of $(U')^I \cap \mathbf{S}[\mathbb{G}]$ assumes at most two values (mod θ_0), with one supported either on all of I, or on $\{v, \infty\}$ (for some $v \in V$), or on $\{v_1, v_2, \infty\}$ (for some $v_1 \overset{E}{—} v_2$).

(As on page 11, we will drop the "$\cap\, \mathbf{S}[\mathbb{G}]$" when the context is unambiguous.)

Let $\mathbf{x} = t(\mathbf{g}_v, \ldots, \mathbf{g}_{v_1 v_2}, \ldots, \chi_\infty, \vec{\mathbf{s}})$ represent an arbitrary element of $\mathbf{S}[\mathbb{G}]$ all of whose coordinates lie in U'. Without loss of generality (by precomposing with e') t respects U'; but then this operation is sensitive (mod θ_0) to changes (mod θ_1) in no more than one of its variables. Since all generators are constant (mod θ_1), we conclude that the blocks of I on which $\mathbf{x}$ is constant (mod θ_0) coincide with those of whichever generator sits at the active place. $\dashv_{\text{IV.2.1}}$

We now identify a subset Γ of the universe, definable (using parameters for the diagonal elements and χ_∞) and a definable preorder $\ll$ on Γ.

Set

$$\Gamma = \left\{ \mathbf{x} \in (U')^I : q(\mathbf{x}, \mathbf{m}_0) = \mathbf{m}_0 \;\&\; q(\mathbf{x}, \chi_\infty) = \chi_\infty \right\}$$

and preorder it by

$$\mathbf{x} \ll \mathbf{y} \iff \forall \mathbf{u}, \mathbf{v} \in U^I \; q(\mathbf{x}, \mathbf{u}) = q(\mathbf{x}, \mathbf{v}) \to q(\mathbf{y}, \mathbf{u}) = q(\mathbf{y}, \mathbf{v})$$

Since the sets U^I and $(U')^I$ are definable (Proposition II.6), it follws that $\ll$ and its associated equivalence relation $\sim$ are definable too. Let $\mathrm{EQ}(v_1, v_2)$ be a formula defining the equivalence $\sim$.

The second conjunct defining Γ implies that if $\mathbf{x} \in \Gamma$ then $\mathbf{x}$ is permutational at infinity. (So, for example, Γ contains $\mathbf{a}_1$ but not $\mathbf{a}_0$.) The first implies that any non-permutational factor of $\mathbf{x}$ must collapse N to m_0. If $\mathbf{x} \in \Gamma$, $\mathbf{u}_1, \mathbf{u}_2 \in U^I$, and x^i is not permutational, then $q(\mathbf{x}, \mathbf{u}_1) = q(\mathbf{x}, \mathbf{u}_2)$ implies $u_1^i \equiv_\mu u_2^i$.

Claim IV.2.2. For $\mathbf{x} \in \Gamma$, define

$$\begin{aligned}\mathrm{supp}(\mathbf{x}) &= \{i \in I : x^i \text{ is permutational}\}\\ &= \{i \in I : q(x^i, m_1) = m_1\}\end{aligned}$$

(We already know that each support is either I or one of the sets $\{v_1, v_2, \infty\}$ ($v_1 \overset{E}{—} v_2$) or $\{v, \infty\}$ ($v \in G$).) Then

$$\mathbf{x} \ll \mathbf{y} \iff \mathrm{supp}(\mathbf{x}) \supseteq \mathrm{supp}(\mathbf{y})$$

($\Rightarrow$): If $v \in \mathrm{supp}(\mathbf{y}) \setminus \mathrm{supp}(\mathbf{x})$, take $\mathbf{u} = q(\mathbf{g}_v, \mathbf{m}_1)$. Then

$$q(\mathbf{x}, \mathbf{u}) = \chi_\infty = q(\mathbf{x}, \chi_\infty)$$

but

$$q(\mathbf{y}, \mathbf{u})_{|v} = q(y^v, m_1) = m_1 \neq m_0 = q(\mathbf{y}, \chi_\infty)_{|v}$$

so $\mathbf{x} \not\ll \mathbf{y}$.

($\Leftarrow$): For $\mathbf{t}, \mathbf{u} \in U^I$, $q(\mathbf{x}, \mathbf{t}) = q(\mathbf{x}, \mathbf{u})$ is equivalent to

$$\mathbf{t}_{|\mathrm{supp}(\mathbf{x})} = \mathbf{u}_{|\mathrm{supp}(\mathbf{x})} \text{ and for } v \notin \mathrm{supp}(\mathbf{x}),\ t^v \equiv_\mu u^v$$

which implies

$$\mathbf{t}_{|\mathrm{supp}(\mathbf{y})} = \mathbf{u}_{|\mathrm{supp}(\mathbf{y})} \text{ and for } v \notin \mathrm{supp}(\mathbf{y}),\ t^v \equiv_\mu u^v$$

which is equivalent to $q(\mathbf{y}, \mathbf{t}) = q(\mathbf{y}, \mathbf{u})$. $\dashv_{\mathrm{IV.2.2}}$

As an immediate consequence, we have that every $\mathbf{x} \in \Gamma$ is $\sim$ to exactly one of $\{\mathbf{a}_1\} \cup \{\mathbf{g}_{v_1 v_2} : v_1 \overset{E}{—} v_2\} \cup \{\mathbf{g}_v : v \in V\}$. The quotient partial order on $\Gamma/\sim$ has height two, with $\mathbf{a}_1$ at level zero, all the edges at level one and all the vertices at level two.

Let $\mathrm{WHO}(v_1)$ be a formula asserting that $v_1 \in \Gamma$ and v_1 is at $\ll$-level two. We have just observed that the map $w \mapsto \mathbf{g}_w/\sim$ is a bijection of V with the extension of $\mathrm{WHO}(v_1)$ modulo $\sim$ (which was already found to be a definable equivalence relation). Let $\mathrm{EDGE}(v_1, v_2)$ be a formula asserting that there exists $\mathbf{y} \in \Gamma$ at $\ll$-level one such that $\mathbf{y} \ll v_1$ & $\mathbf{y} \ll v_2$. Then these formulas recover the structure of $\mathbb{G}$. $\square$

IV–B. The action of the twin group

Definition IV.3. Let $\mathbf{A}$ be any algebra, $U \subseteq A$, and σ be the strongly-solvable radical of $\mathbf{A}$. We write

(1) $\mathfrak{S}_U^{\mathbf{A}} := \mathsf{Pol}_1(\mathbf{A}_{|U}) \cap \mathfrak{S}(U)$ for the group of permutations of U realized as polynomials of $\mathbf{A}$, and

(2) $\mathrm{T}_U^{\mathbf{A}}$ for the subgroup consisting of those $f \in \mathfrak{S}_U^{\mathbf{A}}$ such that for some term $t(v_0, \dots, v_n)$ and some $\vec{d} \equiv_\sigma \vec{e}$ we have

$$\mathbf{A}_{|U} \models v_0 = t(v_0, \vec{e}) \ \&\ f(v_0) = t(v_0, \vec{d})$$

(Such a permutation is known as a σ-twin of the identity.)

A straightforward computation shows that $\mathrm{T}_U^{\mathbf{A}}$ is normal in $\mathfrak{S}_U^{\mathbf{A}}$.

Note that there is nothing special about the radical in this context; we can define α-twins for any congruence α, but since we will be exclusively concerned with σ-twins in this investigation, we will leave the definition more specialized so as to avoid needing a third parameter in the symbol $\mathrm{T}_U^{\mathbf{A}}$.

Proposition IV.4. *Let $\mathbf{A}$ be a finite algebra. If $\perp_A \overset{1}{\prec} \mu$ in $\mathrm{Con}(\mathbf{A})$ and U is $(\perp, \mu)$-minimal, then*

(1) $\mathfrak{S}_U^{\mathbf{A}}$ acts transitively by polynomial isomorphisms on the set of traces inside U;

(2) the action of $\mathfrak{S}_U^{\mathbf{A}}$ on the body of U has at most two orbits;

(3) if some $f \in \mathfrak{S}_U^{\mathbf{A}}$ nontrivially permutes some trace, then $\mathfrak{S}_U^{\mathbf{A}}$ acts transitively on the body of U.

Proof. That $\mathfrak{S}_U^{\mathbf{A}}$ acts on traces is an easy consequence of the fact that μ is a congruence of the algebra.

To transitivity: μ is generated by any of its nontrivial pairs, so let $N_i \subseteq U$ $(i = 1, 2)$ be traces containing elements $a_i \neq b_i$. Then we can string a chain of elements

$$a_2 = u_0 \neq u_1 \neq \cdots \neq u_{m+1} = b_2$$

where $\{u_j, u_{j+1}\} = \{f_j(a_1), f_j(b_1)\}$ for some $f_j \in \mathfrak{S}_U^{\mathbf{A}}$. Then $f_m(N_1) = N_2$. This argument actually shows that $b_2 \in \mathfrak{S}_U^{\mathbf{A}}(a_1) \cup \mathfrak{S}_U^{\mathbf{A}}(b_1)$, which proves the second and third statements. □

Lemma IV.5. *Let $\mathbf{S}$ be a finite subdirectly-irreducible algebra with type-1 monolith μ and strongly-solvable radical σ satisfying $\mathrm{C}(\sigma, \sigma; \mu)$. Let $U = e(S)$ be a $(\perp, \mu)$-minimal set. If $\mathrm{T}_U^{\mathbf{S}}$ nontrivially permutes some trace, then* $\mathrm{HSP}\,(\mathbf{S})$ *is hereditarily finitely undecidable.*

Proof. The last statement of Proposition IV.4 ensures that $\mathfrak{S}_U^{\mathbf{S}}$ acts transitively on the body of U; the same may not be true of the induced action of $\mathrm{T}_U^{\mathbf{S}}$, but elementary group theory shows that

$\mathfrak{S}^{\mathbf{S}}_U/\mathrm{T}^{\mathbf{S}}_U$ acts in a well-defined and transitive way on the orbits of the action by $\mathrm{T}^{\mathbf{S}}_U$. Since the action of $\mathfrak{S}^{\mathbf{S}}_U$ is transitive, we will use the symbol $\mathcal{O}(a)$ exclusively to refer to the orbit of the element $a \in \mathrm{Body}(U)$ under the action by $\mathrm{T}^{\mathbf{S}}_U$.

Claim IV.5.1. For each $c \in \mathrm{Body}(U)$,

$$|\mathcal{O}(c) \cap N| > 1$$

where N is the trace containing c.

Let $g(a) = b \equiv_{\mu\setminus\perp} a$ and $f(c) = a$, where $g \in \mathrm{T}^{\mathbf{S}}_U$ is the hypothesized nontrivial permutation of a trace and $f \in \mathfrak{S}^{\mathbf{S}}_U$. Then $f^{-1} \circ g \circ f(c) \equiv_{\mu\setminus\perp} c$, which proves the claim. $\dashv_{\text{IV.5.1}}$

By Lemma IV.2, we may assume that $\mathrm{C}(\sigma, \mu_{|U}; \perp)$. This immediately implies that if $t(v_0, \ldots, v_n)$ is any term and $\vec{c} \equiv_\sigma \vec{d}$, and if $t(U, \vec{c}), t(U, \vec{d}) \subseteq U$ then these two polynomials are either both permutations of U or both collapse traces into points.

Our plan is a bit more complicated this time around. Instead of semantically embedding graphs into diagonal subpowers of $\mathbf{S}$, we will embed them into algebras $\mathbf{C}[\mathbb{G}] = \mathbf{D}(\mathbb{G})/\Theta$, where $\mathbf{D}(\mathbb{G}) \leq \mathbf{S}^I$ is a diagonal subpower of $\mathbf{S}$ and $\Theta \leq \sigma^I$. We will not attempt to show that Θ is a definable congruence, uniformly or otherwise.

Fix your favorite graph $\mathbb{G} = \langle V, E\rangle$. Define $V^{\pm} = \{v^+, v^- : v \in V\}$ (the disjoint union of two copies of V), and set $I = V^{\pm} \sqcup \{\infty\}$; each of the sets $\{v^+, v^-\}$ as well as $\{\infty\}$ will be called a "vertex block" or "V-block". Let $\mathbf{D} = \mathbf{D}(\mathbb{G}) \leq \mathbf{S}^I$ be the diagonal subpower generated by the set Γ which is the union of the following three disjoint sets:

- Γ_0 is the set of those $\mathbf{x} \in U^I$ which are constant on each V-block and constant $(\mathrm{mod}\ \sigma)$ on all of I.
- Γ_V is the set of those $\mathbf{x} \in \mathbf{S}^I$ such that for some $a \in \mathrm{Body}(U)$, $x^i \in (a/\sigma) \cap \mathcal{O}(a)$ for all $i \in I$, and for one $v \in V$, $x^{v^+} \equiv_{\mu\setminus\perp} x^{v^-}$, while for all $w \neq v$, $x^{w^+} = x^{w^-}$. For convenience, if $\mathbf{x}$ and v are as just described, we write $\mathrm{Label}(\mathbf{x}) = \langle v, x^{v^+}\rangle$.
- Γ_E is like Γ_V; but instead of having one nonconstant vertex block, each point will have two, at the blocks of v and w, where $v \overset{E}{\text{—}} w$, and write $\mathrm{Label}(\mathbf{x}) = \langle v, x^{v^+}, w, x^{w^+}\rangle$.

We will refer to the non-constant vertex blocks as "spikes".

Observe that since each generator is constant modulo σ, every element of $\mathbf{D}$ is too.

Claim IV.5.2. $\mathbf{D} \cap U^I \subset \Gamma$, and for every polynomial $\mathbf{f} \in \mathfrak{S}^{\mathbf{D}}_U$ and every $v \in V$, the v^+ component of $\mathbf{f}$ is the same function as the v^- component.

(As is our convention, $\mathfrak{S}^{\mathbf{D}}_U$ should really more precisely be $\mathfrak{S}^{\mathbf{D}}_{U^I \cap D}$, but that would be cumbersome.)

Both parts of the claim are consequences of Maroti's Lemma. To the first: let $\mathbf{y} = et(\Gamma_0, \Gamma_V, \Gamma_E)$ be a typical element of $\mathbf{D} \cap U^I$. There is one special input place where this term is sensitive to changes by μ; at all other places, flatten out all the spikes so that $\mathbf{y} = et'(\Gamma_0, \mathbf{x})$ where $\mathbf{x} \in \Gamma$ is the element at the special place. Then if $\mathbf{x} \in \Gamma_0$ or if $et'(\Gamma_0, \cdot)$ is not injective on U^I then at each coordinate $t'(\Gamma_{0|i}, \cdot)$ collapses μ into points; under those hypotheses, $\mathbf{y} \in \Gamma_0$.

On the other hand, if $et'(\Gamma_0, \cdot)$ permutes U^I, then $\mathbf{y}$ has the same spike pattern that $\mathbf{x}$ had (since every element of Γ_0 is constant on V-blocks); furthermore, if $\mathbf{x} \in \Gamma_V \cup \Gamma_E$, we can conclude that $\mathbf{y}$ takes all its values from one $\mathrm{T}_U^{\mathbf{S}}$-orbit, since all the coordinatewise polynomials $et'(\Gamma_{0|i}, \cdot)$ are σ-twins, and hence all in the same coset mod $\mathrm{T}_U^{\mathbf{S}}$. But this means that $\mathbf{y} \in \Gamma$ already.

Similarly for the second part of the claim: let $f(v_0) = et(v_0, \Gamma) \in \mathfrak{S}_U^{\mathbf{D}}$; then it is not possible for the special variable to be anything except the first. The claim follows, since the other parameters only vary up to μ on vertex blocks. $\dashv_{\text{IV.5.2}}$

In fact, let $T(v_0, \ldots, v_n)$ be a universal term for $\mathrm{T}_U^{\mathbf{S}}$, i.e. there exist pairwise-σ tuples $\{\vec{d}_g : g \in \mathrm{T}_U^{\mathbf{S}}\}$ so that $g(v_0) = T(v_0, \vec{d}_g)$ for all g. (We leave it to the reader to verify that such a term exists.) Then this term allows us to realize the full product ${\mathrm{T}_U^{\mathbf{S}}}^{V \sqcup \{\infty\}}$ as polynomial permutations of U^I; it follows that $\mathfrak{S}_U^{\mathbf{D}}$ is isomorphic to the inverse image of the diagonal subgroup under the canonical projection

$$
{\mathfrak{S}_U^{\mathbf{S}}}^{V \sqcup \{\infty\}} \longrightarrow (\mathfrak{S}_U^{\mathbf{S}} / \mathrm{T}_U^{\mathbf{S}})^{V \sqcup \{\infty\}}
$$

We still have to define the congruence Θ. This is done as follows: Θ will be generated by identifying those pairs $\langle \mathbf{x}, \mathbf{y} \rangle$ such that

- $\mathbf{x}, \mathbf{y} \in \Gamma_0$ and for all $i \in I$, $x^i \equiv_\mu y^i$, or
- $\mathbf{x}, \mathbf{y} \in \Gamma_V$, $\mathrm{Label}(\mathbf{x}) = \mathrm{Label}(\mathbf{y}) = \langle v, a \rangle$ and for all $i \neq v^+$, $x^i \equiv_\mu y^i$, or
- $\mathbf{x}, \mathbf{y} \in \Gamma_E$, $\mathrm{Label}(\mathbf{x}) = \mathrm{Label}(\mathbf{y}) = \langle v, a, w, b \rangle$, and for all $i \neq v^+, w^+$, $x^i \equiv_\mu y^i$.

and we set $\mathbf{C} = \mathbf{C}[\mathbb{G}] = \mathbf{D}/\Theta$. We will usually write, e.g., Γ instead of Γ/Θ when context makes it unambiguous.

Claim IV.5.3. $\Theta_{|\Gamma}$ consists of just the generating pairs and no more.

To see this, let $\langle \mathbf{x}, \mathbf{y} \rangle$ be a generating pair and $\mathbf{f} \in \mathsf{Pol}_1(\mathbf{D}_{|U})$. Then $\langle \mathbf{f}(\mathbf{x}), \mathbf{f}(\mathbf{y}) \rangle$ is clearly a generating pair if $\mathbf{f}$ collapses μ to points, or if $\mathbf{x}$ and $\mathbf{y}$ belong to Γ_0, so let $\mathbf{f} \in \mathfrak{S}_U^{\mathbf{D}}$. Then if $\mathbf{x}, \mathbf{y} \in \Gamma_V$ with $\mathrm{Label}(\mathbf{x}) = \mathrm{Label}(\mathbf{y}) = \langle v, a \rangle$ then

$$
\begin{gathered}
a = x^{v^+} = y^{v^+} \Rightarrow f^v(a) = f^v(x^{v^+}) = f^v(y^{v^+}) \\
x^i \equiv_\mu y^i \Rightarrow f^i(x^i) \equiv_\mu f^i(y^i)
\end{gathered}
$$

so $\langle \mathbf{f}(\mathbf{x}), \mathbf{f}(\mathbf{y}) \rangle$ is again a generating pair. The proof for generating pairs from Γ_E is identical.

By Lemma IV.1, σ^I is a uniformly definable congruence; it follows that quantification over any of the groups $\mathrm{T}_\Gamma^{\mathbf{D}}, \mathfrak{S}_\Gamma^{\mathbf{D}}, \mathrm{T}_\Gamma^{\mathbf{C}}, \mathfrak{S}_\Gamma^{\mathbf{C}}$ is uniformly first-order in the respective algebra. Of course, we also have that Γ (respectively U^I/Θ) is a definable subset of both algebras as well, since it consists of precisely the fixed points of the polynomial retraction e.

Claim IV.5.4. (i) If $g \in \mathrm{T}^{\mathbf{S}}_U$, $a \equiv_\sigma b$ and $g(a) \equiv_\mu a$ then $g(b) \equiv_\mu b$.

(ii) $\Gamma_0 \cap U^I$ is uniformly definable (using at most $n \cdot |\mathrm{T}^{\mathbf{S}}_U|$ parameters) in $\mathbf{C}$.

The first part is true because $\mathrm{C}(\sigma, \sigma; \mu)$:

$$\begin{gathered} a = T(a, \vec{d}_{\mathrm{id}}) \equiv_\mu T(a, \vec{d}_g) = g(a) \\ \Downarrow \\ b = T(b, \vec{d}_{\mathrm{id}}) \equiv_\mu T(b, \vec{d}_g) = g(b) \end{gathered}$$

where $T(v_0, \ldots, v_n)$ is the universal term for $\mathrm{T}^{\mathbf{S}}_U$ defined above. To the second: for each $g \in \mathrm{T}^{\mathbf{S}}_U$ let $\vec{\mathbf{c}}_g$ be constants so that

$$T(\cdot, \vec{\mathbf{c}}_g) = \mathrm{id}_{|V} \oplus g_{|\infty}$$

Then for $\mathbf{x} \in U^I$ we have

$$T(\mathbf{x}, \vec{\mathbf{c}}_g) \equiv_\Theta \mathrm{x} \iff g(x^\infty) \equiv_\mu x^\infty$$

Hence

$$\mathbf{x} \in \Gamma_0 \Rightarrow \forall g \in \mathrm{T}^{\mathbf{S}}_U \left(T(\mathbf{x}, \vec{\mathbf{c}}_g) \equiv_\Theta \mathbf{x} \to T(\mathbf{x}, \vec{\mathbf{d}}_g) \equiv_\Theta \mathbf{x} \right)$$

(where $\vec{\mathbf{d}}_g$ are the obvious diagonal elements), while if $\mathbf{x} \in \Gamma_V$ (resp. Γ_E) with label $\langle v, a \rangle$ (resp. $\langle v, a, w, b \rangle$) and $g(a) \equiv_{\mu \setminus \perp} a$ then $g(x^\infty) \equiv_\mu x^\infty$ so

$$T(\mathbf{x}, \vec{\mathbf{c}}_g) \equiv_\Theta \mathbf{x} \text{ and } T(\mathbf{x}, \vec{\mathbf{d}}_g) \not\equiv_\Theta \mathbf{x}$$

$\dashv_{\text{IV.5.4}}$

We are almost done: for the last step, define a preorder $\ll$ on $\Gamma \setminus \Gamma_0$ by

$$\begin{aligned} \mathbf{x} \ll \mathbf{y} \iff & \exists f \in \mathfrak{S}^{\mathbf{S}}_U \, \exists \mathbf{g} \in \mathrm{T}^{\mathbf{C}}_\Gamma \, \big[\mathbf{g} f(\mathbf{x}) \equiv_\sigma \mathbf{y} \, \& \\ & \forall \mathbf{h} \in \mathrm{T}^{\mathbf{C}}_\Gamma \, [\mathbf{h}\mathbf{g} f(\mathbf{x}) \not\equiv_\Theta \mathbf{g} f(\mathbf{x}) \to \mathbf{h}(\mathbf{y}) \not\equiv_\Theta \mathbf{y}] \big] \end{aligned}$$

Claim IV.5.5. (i) If $\mathbf{x}, \mathbf{y} \in \Gamma_V$ (resp. Γ_E) are labeled by the same vertex (resp. edge), they are $\ll$-equivalent.

(ii) If $\mathbf{x}, \mathbf{y} \in \Gamma_V$ (resp. Γ_E) are labeled by different vertices (resp. edges), they are $\ll$-incomparable.

(iii) If $\mathbf{x} \in \Gamma_E$ and $\mathbf{y} \in \Gamma_V$ then $\mathbf{x} \not\ll \mathbf{y}$.

(iv) If $\mathbf{x} \in \Gamma_V, \mathbf{y} \in \Gamma_E$, then $\mathbf{x} \ll \mathbf{y}$ iff $\mathbf{x}$ is labeled by one of the endpoints of the edge which labels $\mathbf{y}$.

All parts of this claim are straightforward:

(i) Say $\mathbf{x}, \mathbf{y} \in \Gamma_E$, $\mathrm{Label}(\mathbf{x}) = \langle v, a_1, w, b_1 \rangle$ and $\mathrm{Label}(\mathbf{y}) = \langle v, a_2, w, b_2 \rangle$. Then $\mathcal{O}(a_j) = \mathcal{O}(b_j)$ ($j \in \{1, 2\}$), and we can choose $f \in \mathfrak{S}^{\mathbf{S}}_U$ so that $f\mathcal{O}(a_1) = \mathcal{O}(a_2)$. Then we can choose

$\{g^i\}_{i\in I} \in \mathrm{T}_U^{\mathbf{S}}$ so that $g^v f(a_1) = a_2$, $g^w f(b_1) = b_2$, and $g^i f(x^i) \equiv_\mu y^i$ for all other i, and set $\mathbf{g} = \bigoplus g^i$. Then in fact $\mathbf{g}f(\mathbf{x}) \equiv_\Theta \mathbf{y}$ so $\mathbf{x} \ll \mathbf{y}$ holds automatically. The proof is the same for Γ_V except easier.

(ii) Say $\mathbf{x}$ has a spike at a V-block where $\mathbf{y}$ does not, say at v. Then for every $f \in \mathfrak{S}_U^{\mathbf{S}}$ and every $\mathbf{g} \in \mathrm{T}_\Gamma^{\mathbf{C}}$, $\mathbf{g}f(\mathbf{x})$ has a spike at v, which $\mathbf{y}$ does not. Assume $\mathbf{g}f(\mathbf{x}) \equiv_\sigma \mathbf{y}$. Choose $h \in \mathrm{T}_U^{\mathbf{S}}$ such that $hg^v f(x^{v^+}) \equiv_{\mu\setminus\perp} g^v f(x^{v^+})$; then $h(y^{v^+}) \equiv_\mu y^{v^+}$. Let $\mathbf{h} \in \mathrm{T}_\Gamma^{\mathbf{C}}$ be h on $\{v^\pm\}$ and the identity on all other vertex blocks; then

$$\mathbf{hg}f(\mathbf{x}) \not\equiv_\Theta \mathbf{g}f(\mathbf{x}) \;\&\; \mathbf{h}(\mathbf{y}) \equiv_\Theta \mathbf{y}$$

(iii) The same as in (ii).

(iv) The direction ($\Rightarrow$) is the same as in (ii). For ($\Leftarrow$), assume that $\mathrm{Label}(\mathbf{x}) = \langle v, a_1\rangle$, $\mathrm{Label}(\mathbf{y}) = \langle v, a_2, w, b\rangle$. Choose $f \in \mathfrak{S}_U^{\mathbf{S}}$ with $f(a_1) = a_2$, and for $i \neq v^\pm$ choose $g^i \in \mathrm{T}_U^{\mathbf{S}}$ so that $g^i f(x^i) \equiv_\mu y^i$, $g^v = \mathrm{id}$, $\mathbf{g} = \bigoplus_i g^i$; then we have $\mathbf{z} := \mathbf{g}f(\mathbf{x}) \equiv_\mu \mathbf{y}$ and $z^{v^+} = y^{v^+}$. Consequently, if $\mathbf{h}(\mathbf{z}) \not\equiv_\Theta \mathbf{z}$ then either

$$h^i(z^i) \not\equiv_\mu z^i$$

for some $i \in I$, in which case $h^i(y^i) \not\equiv_\mu y^i$, or

$$h^v(y^{v^+}) = h^v(z^{v^+}) \neq z^{v^+} = y^{v^+}$$

so in either case $\mathbf{h}(\mathbf{y}) \not\equiv_\Theta \mathbf{y}$. $\dashv_{\text{IV.5.5}}$

This completes the proof of the lemma, since up to $\ll$-biequivalence, vertices of $\mathbb{G}$ correspond precisely to $\ll$-classes at level zero, edges to classes at level one, and two vertices are joined iff there is a class properly dominating both. $\square$

IV–C. The strongly-solvable radical is abelian

Lemma IV.6. *Let $\mathbf{S}$ be a finite subdirectly-irreducible algebra with unary-type monolith μ and strongly-solvable radical σ satisfying $\mathrm{C}(\sigma,\sigma;\mu)$, $\mathrm{C}(\sigma,\mu;\perp)$, and $\mathrm{C}(\mu,\sigma;\perp)$ but not $\mathrm{C}(\sigma,\sigma;\perp)$. Then $\mathrm{HSP}(\mathbf{S})$ is hereditarily finitely undecidable.*

Proof. Choose a package

$$c = t_0(a_0, \vec{b}_0) = t_0(a_0, \vec{b}_1)$$
but
$$m_0 = t_0(a_1, \vec{b}_0) \neq t_0(a_1, \vec{b}_1) = m_1$$

witnessing $\neg\mathrm{C}(\sigma,\sigma;\perp)$, where $a_0/\sigma = a_1/\sigma =: A$ and $\vec{b}_0 \equiv_\sigma \vec{b}_1$. Since $\mathrm{C}(\sigma,\sigma;\mu)$, $m_0 \equiv_\mu m_1$, and we may suppose that the range of $t_0(v_0,\ldots,v_\ell)$ is included in a $(\perp,\mu)$-minimal set U. Denote the trace containing the m_j by M.

We will be working with diagonal subpowers $\mathbf{X} \leq \mathbf{S}^I$ and their quotients $\mathbf{Y} = \mathbf{X}/\Theta$, where $\Theta \leq \sigma = \sigma^I \cap X^2 \in \mathrm{Con}(\mathbf{X})$. Lemma IV.1 once again implies that σ is a definable congruence in all such $\mathbf{Y}$.

We will wherever possible refer to elements of $\mathbf{Y}$ with $\mathbf{x}$ rather than $\mathbf{x}/\Theta$, with the understanding that $\mathbf{x} \in S^I$ is one representative. (Of course, this will necessitate showing that certain properties are well-defined.)

For such algebras $\mathbf{Y}$, and $\vec{\mathbf{y}}_1, \vec{\mathbf{y}}_2 \equiv_\sigma \vec{\mathbf{b}}_0$ define

$$\mathrm{E}^{\mathbf{Y}}(\vec{\mathbf{y}}_1, \vec{\mathbf{y}}_2) = \{\mathbf{x} \equiv_\sigma \mathbf{a}_0 \colon \mathbf{Y} \models t_0(\mathbf{x}, \vec{\mathbf{y}}_1) = t_0(\mathbf{x}, \vec{\mathbf{y}}_2)\}$$

In particular, we have

$$\mathrm{E}^{\mathbf{S}}(\vec{b}_0, \vec{b}_1) \subsetneq A$$

and there is no loss of generality in assuming that the $\vec{b}_j$ are chosen so that their equalizer set is maximal for being properly included in A.

We will be using ℓ-tuples extensively, so to avoid a proliferation of vector notation we will reserve the letters b, y, z for ℓ-tuples and a, x for single elements.

The plan is as follows: We want to interpret the class of graphs with at least three vertices into $\mathrm{HSP}\,(\mathbf{S})$. Given such a graph $\mathbb{G} = \langle V, E\rangle$, we will choose an index set I and a diagonal subpower $\mathbf{D} \leq \mathbf{S}^I$, which will depend only on V, and then a congruence $\Theta \in \mathrm{Con}(\mathbf{D})$ below σ^I (in fact, below μ^I), which will depend on both V and E, and set $\mathbf{C} = \mathbf{D}/\Theta$. Θ will be sparse in a sense we will make precise. Then we will define a set $\mathcal{B} \subset C^\ell$, and show that a preorder $\ll$ recovering the index set I is definable there; vertices will interpret as unions of two $\ll$-biequivalence classes, and the edge relation from $\mathbb{G}$ will be first-order definable on these vertices. Here "definable" will include reference to $|A| + 1$ parameters (in addition to the diagonal).

We begin with a graph $\mathbb{G} = \langle V, E\rangle$, and set $I = V^\pm \sqcup \{\infty\}$ as in Lemma IV.5. Define $\mathbf{D} \leq \mathbf{S}^I$ to be the subalgebra consisting of all elements which are constant modulo σ. By the same logic applied in Claim IV.5.2, $\mathfrak{S}_U^{\mathbf{D}}$ consists of those $\mathbf{f} \in (\mathfrak{S}_U^{\mathbf{S}})^I$ such that all f^i belong to the same coset modulo $\mathrm{T}_U^{\mathbf{S}}$. (Here the coordinate functions f^{v^+}, f^{v^-} may be different.) The relation $\mathrm{C}(\sigma, \mu; \perp)$ implies that a polynomial $\mathbf{f}(v_0) = t(v_0, \vec{\mathbf{d}})$ whose image is contained in U^I is either a permuation of U at all coordinates or collapses traces to points at all coordinates. We note for future reference that

Claim IV.6.1. if $f_1, f_2 \in \mathfrak{S}_U^{\mathbf{S}}$ belong to the same coset modulo $\mathrm{T}_U^{\mathbf{S}}$, and if $f_1(M) = M = f_2(M)$ then by Lemma IV.5 $f_{1|M} = f_{2|M}$

In particular, this is true if these are the coordinate functions of some $\mathbf{f} \in \mathfrak{S}_U^{\mathbf{D}}$. $\dashv_{\text{IV.6.1}}$

Let $\mathbf{C} = \mathbf{D}/\Theta$, where Θ is the congruence on $\mathbf{D}$ generated by identifying

$$m_{1|v^+} \oplus m_{0|I\setminus\{v^+\}} \equiv_\Theta m_{1|v^-} \oplus m_{0|I\setminus\{v^-\}} \qquad (v \in V)$$
$$m_{1|\{v^+,w^+\}} \oplus m_{0|I\setminus\{v^+,w^+\}} \equiv_\Theta m_{1|\{v^-,w^-\}} \oplus m_{0|I\setminus\{v^-,w^-\}} \qquad (v \overset{E}{—} w)$$

Claim IV.6.2. (i) $\Theta \leq \mu^I$, and if $\mathbf{x}_1 \equiv_\Theta \mathbf{x}_2$ then $x_1^\infty = x_2^\infty$.

(ii) $\Theta_{|U^I}$ has blocks of cardinality 1 and 2 only.

(iii) If $\mathbf{x}_1, \mathbf{x}_2 \in U^I$ and $\mathbf{x}_1 \equiv_\Theta \mathbf{x}_2$, then the set of coordinates where they differ is either empty, one V-block $\{v^+, v^-\}$, or two V-blocks $\{v^+, v^-, w^+, w^-\}$ where $v \overset{E}{—} w$.

The first statement is clear. To see (ii), first observe that if $\mathbf{f} \in \mathsf{Pol}_1(\mathbf{D}_{|U}) \setminus \mathfrak{S}_U^{\mathbf{D}}$ then

$$\mathbf{f}(m_{1|\{v^+\}} \oplus m_{0|I\setminus\{v^+\}}) = \mathbf{f}(m_{0|I}) = \mathbf{f}(m_{1|\{v^-\}} \oplus m_{0|I\setminus\{v^-\}})$$

so it suffices to consider images of generating pairs under permutations $\mathbf{f} \in \mathfrak{S}_U^{\mathbf{D}}$. Next, since $\mathfrak{S}_U^{\mathbf{S}}/\mathrm{T}_U^{\mathbf{S}}$ acts on orbits and since we may assume that $\mathcal{O}(m_0) \neq \mathcal{O}(m_1)$, we may conclude that any image $\mathbf{f}(m_{1|\{v^+\}} \oplus m_{0|I\setminus\{v^+\}})$ takes values in one orbit at all coordinates except v^+ and in a different orbit there, and similarly for the other elements involved in the generating pairs. We prove the claim for generators of the vertex type; the edge-type argument is no different.

Given any putative $\Theta_{|U}$-block of more than two elements, we can find a subset of three elements of the form

$$\mathbf{x}_1 = \mathbf{f}_1(m_{1|\{v^+\}} \oplus m_{0|I\setminus\{v^+\}}) = \mathbf{f}_2(m_{1|\{v^+\}} \oplus m_{0|I\setminus\{v^+\}}) = \mathbf{x}_2$$
$$\mathbf{y}_1 = \mathbf{f}_1(m_{1|\{v^-\}} \oplus m_{0|I\setminus\{v^-\}}) \overset{?}{=} \mathbf{f}_2(m_{1|\{v^-\}} \oplus m_{0|I\setminus\{v^-\}}) = \mathbf{y}_2$$

or vice versa. The first line shows that $\mathbf{f}_2^{-1} \circ \mathbf{f}_1(M^I) = M^I$; but since $\mathbf{f}_2^{-1} \circ \mathbf{f}_1 \in \mathrm{T}_U^{\mathbf{D}}$, it must fix M^I pointwise. Hence $\mathbf{y}_1 = \mathbf{y}_2$.

Looking a little more closely at the argument, we see that in fact a pair of unequal elements $\mathbf{x}_1, \mathbf{x}_2 \in U^I$ are Θ-related iff they are the image of a generating pair under some $\mathbf{f} \in \mathfrak{S}_U^{\mathbf{D}}$. Claim (iii) follows immediately. $\dashv_{\text{IV.6.2}}$

With this claim in hand, it is well-defined to speak of x^∞ for $\mathbf{x} \in C$. Furthermore, by Claim IV.6.1, the image of any member of a generating pair under $\mathbf{f} \in \mathfrak{S}_U^{\mathbf{D}}$ cannot be a constant element. (In other words, the constant elements of U^I are isolated modulo Θ.)

Throughout the remainder of the proof, any ℓ-tuple $\mathbf{y}$ or $\mathbf{z}$ will be assumed to be σ-congruent to $\mathbf{b}_0$, and to satisfy the condition

$$\mathbf{c} = t_0(\mathbf{a}_0, \mathbf{b}_0) \equiv_\Theta t_0(\mathbf{a}_0, \mathbf{y}) \tag{IV.6.3}$$

(which is clearly first-order in $\mathbf{C}$). Since $\mathbf{c}$ is isolated, this is in fact an equality. (For instance, every ℓ-tuple from $\{b_0, b_1\}^I$ satisfies this condition, and our life would be much easier if we could work with just that set. The following can be read as a way of coming as close to this as feasible.)

Claim IV.6.4. Define a parameter $\mathfrak{b} = b_{1|\infty} \oplus b_{0|I\setminus\{\infty\}}$ which will be fixed throughout the remainder of the proof. The predicates

$$\mathrm{E}^{\mathbf{S}}(y^\infty, b_0) = A$$

and

$$\mathrm{E}^{\mathbf{S}}(y^\infty, b_1) = A$$

(in the free variable $\mathbf{y}$) are definable using $\mathfrak{b}$ together with $|A|$ other parameters.

This is because

$$\mathrm{E}^{\mathbf{S}}(y^\infty, b_0) = A \iff \bigwedge_{a \in A} a_{|\infty} \oplus a_{0|I\setminus\{\infty\}} \in \mathrm{E}^{\mathbf{C}}(\mathbf{b}_0, \mathbf{y})$$
$$\mathrm{E}^{\mathbf{S}}(y^\infty, b_1) = A \iff \bigwedge_{a \in A} a_{|\infty} \oplus a_{0|I\setminus\{\infty\}} \in \mathrm{E}^{\mathbf{C}}(\mathbf{b}_1, \mathbf{y})$$

We will not name or even make explicit mention of the parameters $a_{|\infty} \oplus a_{0|I\setminus\{\infty\}}$ any more, but they are implicitly present in all that follows. $\dashv_{\text{IV.6.4}}$

The next claim does most of the heavy lifting in this lemma.

Claim IV.6.5. Suppose $\mathbf{y}$ satisfies condition (IV.6.3) and that $\mathrm{E}^{\mathbf{S}}(y^\infty, b_1) = A$. Then the set

$$P(\mathbf{y}) := \left(\bigoplus_{i \neq \infty} \mathrm{E}^{\mathbf{S}}(b_0, y^i) \oplus \left(A \setminus \mathrm{E}^{\mathbf{S}}(b_0, b_1)\right) \right) / \Theta$$

is a definable subset of $\mathbf{C}$.

To show this, we will need one auxiliary definition which will be repeatedly useful:

Definition. If $\mathrm{E}^{\mathbf{S}}(y_1^\infty, b_0) = A = \mathrm{E}^{\mathbf{S}}(y_2^\infty, b_1)$, write $\mathbf{y}_1 \propto \mathbf{y}_2$ if the following equivalent conditions are satisfied:

(1) $\mathrm{E}^{\mathbf{S}}(y_1^i, y_2^i) = A$ for all $i \neq \infty$

(2) $\mathrm{E}^{\mathbf{C}}(\mathbf{b}_0, \mathfrak{b}) \subseteq \mathrm{E}^{\mathbf{C}}(\mathbf{y}_1, \mathbf{y}_2)$

To see that these conditions are in fact equivalent, in the direction (1) $\Rightarrow$ (2), if $t_0(\mathbf{x}, \mathbf{b}_0) \equiv_\Theta t_0(\mathbf{x}, \mathfrak{b})$, then

$$t_0(x^\infty, y_1^\infty) = t_0(x^\infty, b_0) = t_0(x^\infty, b_1) = t_0(x^\infty, y_2^\infty)$$

so that $t_0(\mathbf{x}, \mathbf{y}_1)$ is in fact equal to $t_0(\mathbf{x}, \mathbf{y}_2)$. Conversely, fix $i \neq \infty$ and $a \in A$. We know that

$$\mathbf{c} = t_0(\mathbf{a}_0, \mathbf{y}_1) = t_0(\mathbf{a}_0, \mathbf{y}_2)$$
$$\text{and}$$
$$t_0(a_{|i} \oplus a_{0|I \setminus \{i\}}, \mathbf{b}_0) = t_0(a_{|i} \oplus a_{0|I \setminus \{i\}}, \mathfrak{b})$$
$$\text{hence}$$
$$t_0(a_{|i} \oplus a_{0|I \setminus \{i\}}, \mathbf{y}_1) \equiv_\Theta t_0(a_{|i} \oplus a_{0|I \setminus \{i\}}, \mathbf{y}_2)$$

and these elements do not differ except possibly at i; hence they are in fact equal, showing that

$$t_0(a, y_1^i) = t_0(a, y_2^i)$$

Note that condition (2) is clearly first-order.

Now to the proof of Claim IV.6.5: let $\mathbf{y}$ be as in the statement, and let $\mathbf{z}$ be the tuple which agrees with b_0 at ∞ and with $\mathbf{y}$ everywhere else, so $\mathbf{z} \propto \mathbf{y}$.

Now assume further that $\mathbf{x} \in P(\mathbf{y})$. Then

$$\begin{aligned} t_0(\mathbf{x}, \mathbf{b}_0) &= t_0(\mathbf{x}, \mathbf{z}) \text{ and} \\ t_0(\mathbf{x}, \mathfrak{b}) &= t_0(\mathbf{x}, \mathbf{y}) \text{ and} \\ t_0(\mathbf{x}, \mathbf{b}_0) &\not\equiv_\Theta t_0(\mathbf{x}, \mathfrak{b}) \end{aligned}$$

We have shown

$$\begin{aligned} \mathbf{x} \in P(\mathbf{y}) \Rightarrow \exists \mathbf{z} \equiv_\sigma \mathbf{b}_0 \ \mathrm{E}^{\mathbf{S}}(z^\infty, b_0) = A \text{ and } \mathbf{z} \propto \mathbf{y} \text{ and}& \\ \mathbf{x} \in \mathrm{E}^{\mathbf{C}}(\mathbf{b}_0, \mathbf{z}) \cap \mathrm{E}^{\mathbf{C}}(\mathfrak{b}, \mathbf{y}) \text{ and}& \\ \mathbf{x} \notin \mathrm{E}^{\mathbf{C}}(\mathbf{b}_0, \mathfrak{b})& \end{aligned}$$

Next, we show that the converse holds as well.

Assume the following:

$$\mathbf{x} \in A^I \text{ but not in } P(\mathbf{y}) \tag{IV.6.6}$$
$$\mathbf{z} \equiv_\sigma \mathbf{b}_0 \text{ with } \mathrm{E}^{\mathbf{S}}(z^\infty, b_0) = A \tag{IV.6.7}$$
$$\mathbf{z} \propto \mathbf{y} \tag{IV.6.8}$$
$$\mathbf{x} \in \mathrm{E}^{\mathbf{C}}(\mathbf{b}_0, \mathbf{z}) \cap \mathrm{E}^{\mathbf{C}}(\mathfrak{b}, \mathbf{y}) \tag{IV.6.9}$$

We must show that $\mathbf{x} \in \mathrm{E}^{\mathbf{C}}(\mathbf{b}_0, \mathfrak{b})$.

By (IV.6.6), we know that for some $i \neq \infty$, $x^i \notin \mathrm{E}^{\mathbf{S}}(b_0, y^i)$. By (IV.6.8), $\mathrm{E}^{\mathbf{S}}(y^i, z^i) = A$ for all $i \neq \infty$.

Working in $\mathbf{D}$, define elements

$$\begin{aligned} \mathbf{u}_{00} = t_0(\mathbf{x}, \mathbf{b}_0) &\quad t_0(\mathbf{x}, \mathbf{z}) = \mathbf{u}_{01} \\ \mathbf{u}_{10} = t_0(\mathbf{x}, \mathfrak{b}) &\quad t_0(\mathbf{x}, \mathbf{y}) = \mathbf{u}_{11} \end{aligned}$$

Our assumptions imply the following:

$$u_{10}^{\infty} = u_{11}^{\infty} \qquad \text{since } \mathrm{E}^{\mathbf{S}}(y^{\infty}, b_1) = A \tag{IV.6.10}$$

$$u_{00}^{\infty} = u_{01}^{\infty} \qquad \text{by (IV.6.7)} \tag{IV.6.11}$$

$$i \neq \infty \Rightarrow u_{00}^{i} = u_{10}^{i} \qquad \text{(obvious)} \tag{IV.6.12}$$

$$i \neq \infty \Rightarrow u_{01}^{i} = u_{11}^{i} \qquad \text{by (IV.6.8)} \tag{IV.6.13}$$

$$\mathbf{u}_{00} \equiv_{\Theta} \mathbf{u}_{01} \qquad \text{by (IV.6.9)} \tag{IV.6.14}$$

$$\mathbf{u}_{10} \equiv_{\Theta} \mathbf{u}_{11} \qquad \text{by (IV.6.9)} \tag{IV.6.15}$$

$$\mathbf{u}_{10} \neq \mathbf{u}_{11} \qquad \text{by (IV.6.6)} \tag{IV.6.16}$$

Together, these imply that $\mathbf{u}_{00} \neq \mathbf{u}_{01}$ also.

Choose $\mathbf{f} \in \mathfrak{S}_U^{\mathbf{D}}$ so that $\{\mathbf{f}(\mathbf{u}_{10}), \mathbf{f}(\mathbf{u}_{11})\}$ is a generating pair for Θ, and let $\mathbf{w}_{ij} = \mathbf{f}(\mathbf{u}_{ij})$. Then (IV.6.10)-(IV.6.16) are still true of the $\mathbf{w}_{ij}$. By definition, $\mathbf{w}_{10}, \mathbf{w}_{11} \in M^I$; the same is true of $\mathbf{w}_{00}, \mathbf{w}_{01}$, which is shown as follows: for $i \neq \infty$, $w_{0j}^i = w_{1j}^i \in M$, while at ∞ we can use $\mathrm{C}(\sigma, \sigma; \mu)$ to get

$$\begin{gathered} f^{\infty} t_0(a_0, b_0) = f^{\infty} t_0(a_0, b_1) \\ \Downarrow \\ w_{01}^{\infty} = w_{00}^{\infty} = f^{\infty} t_0(x^{\infty}, b_0) \equiv_{\mu} f^{\infty} t_0(x^{\infty}, b_1) = w_{10}^{\infty} \in M \end{gathered}$$

Similarly, we may choose $\mathbf{g} \in \mathfrak{S}_U^{\mathbf{D}}$ so that $\{\mathbf{g}(\mathbf{w}_{00}), \mathbf{g}(\mathbf{w}_{01})\}$ is a generating pair for Θ, whose nontriviality is guaranteed by (IV.6.16). But we have $g^i(M) = M$ for all $i \in I$, so we may assume (by Claim IV.6.1) that $g^i = g^j = g$ for all $i, j \in I$.

Now: since $\{\mathbf{w}_{10}, \mathbf{w}_{11}\}$ form a generating pair for Θ and since $|V| \geq 3$, there exists $v \in V$ so that $w_{10}^{v^+} = w_{11}^{v^+}$. This value cannot be m_1, so we have

$$w_{00}^{v^+} = w_{10}^{v^+} = m_0 = w_{11}^{v^+} = w_{01}^{v^+}$$

Hence

$$g(w_{00}^{v^+}) = g(m_0) = g(w_{01}^{v^+})$$

which implies $g(m_0) = m_0$ (since $\{\mathbf{g}(\mathbf{w}_{00}), \mathbf{g}(\mathbf{w}_{01})\}$ are a generating pair). But then

$$\begin{gathered}(\mathbf{g}(\mathbf{w}_{00}))^\infty = m_0 = (\mathbf{g}(\mathbf{w}_{01}))^\infty \\ \Downarrow \\ w_{00}^\infty = m_0 = w_{01}^\infty = w_{10}^\infty = w_{11}^\infty \\ \Downarrow \\ \mathbf{w}_{00} = \mathbf{w}_{10} \\ \Downarrow \\ \mathbf{u}_{00} = \mathbf{u}_{10} \\ \Downarrow \\ \mathbf{x} \in \mathrm{E}^{\mathbf{C}}(\mathbf{b}_0, \mathfrak{b})\end{gathered}$$

This completes the proof of Claim IV.6.5. $\dashv_{\text{IV.6.5}}$

The foregoing claim implies that the mapping

$$\mathbf{y} \mapsto \mathrm{E}^{\mathbf{S}}(b_0, y^i)$$

on the set of those points $\mathbf{y} \equiv_\sigma \mathbf{b}_0$ such that

$$\mathrm{E}^{\mathbf{S}}(y^\infty, b_1) = A$$

is invariant modulo Θ. Let $\mathbf{y}$ be such a point. For any $a \in A$, $\mathbf{a} \in \mathrm{E}^{\mathbf{C}}(\mathfrak{b}, \mathbf{y})$ iff a belongs to all the factor sets $\mathrm{E}^{\mathbf{S}}(b_0, y^i)$ $(i \neq \infty)$. It follows that the set $\mathcal{B}$ of those $\mathbf{y}$ such that

$$\mathrm{E}^{\mathbf{S}}(b_0, b_1) \subseteq \mathrm{E}^{\mathbf{S}}(b_0, y^i) \text{ for all } i \neq \infty \text{ and } \mathrm{E}^{\mathbf{S}}(y^\infty, b_1) = A$$

that is, those $\mathbf{y}$ such that

$$\mathrm{E}^{\mathbf{S}}(b_0, y^i) \in \left\{\mathrm{E}^{\mathbf{S}}(b_0, b_1), A\right\} \text{ for all } i \neq \infty$$

is definable (by asserting that $\mathbf{a} \in \mathrm{E}^{\mathbf{C}}(\mathfrak{b}, \mathbf{y})$ for each $a \in \mathrm{E}^{\mathbf{S}}(b_0, b_1)$). We may define a preorder on $\mathcal{B}$ by

$$\mathbf{y}_1 \ll \mathbf{y}_2 \iff P(\mathbf{y}_2) \subseteq P(\mathbf{y}_1)$$

(Note the reverse inclusion.) Because we chose $\mathrm{E}^{\mathbf{S}}(b_0, b_1)$ maximal, the associated partial order is isomorphic to the boolean algebra with $2|V|$ atoms. Indeed, each tuple $b_{1|i,\infty} \oplus b_{0|I\setminus\{i,\infty\}}$ sits at $\ll$-level 1; we denote the elements at $\ll$-levels one and two by $\mathcal{B}_1$ and $\mathcal{B}_2$ respectively. Let $\mathrm{WHO}(v_0)$ be a formula (in the parameters we have already mentioned) asserting that $v_0 \in \mathcal{B}_1$.

For $\mathbf{y} \in \mathcal{B}_1$, let $\chi(\mathbf{y})$ denote the (unique) coordinate $i \neq \infty$ such that $\mathrm{E}^{\mathbf{S}}(b_0, y^i) = \mathrm{E}^{\mathbf{S}}(b_0, b_1) \subsetneq A$. If $\chi(\mathbf{y}) \in \{v^+, v^-\}$ we set $|\chi|(\mathbf{y}) = v$.

Assume that $|\chi|(\mathbf{y}_1) = |\chi|(\mathbf{y}_2)$. Then either $\chi(\mathbf{y}_1) = \chi(\mathbf{y}_2)$, which we know to be definable, or for some $v \in V$ we have $\chi(\mathbf{y}_1) = v^+$ and $\chi(\mathbf{y}_2) = v^-$ (or vice versa). Define

$$\mathbf{b}^+ = b_{1|v^+,\infty} \oplus b_{0|\text{else}} \qquad \mathbf{b}^- = b_{1|v^-,\infty} \oplus b_{0|\text{else}}$$

Then $\mathbf{b}^+, \mathbf{b}^- \in \mathcal{B}_1$, $\chi(\mathbf{y}_1) = \chi(\mathbf{b}^+)$, and $\chi(\mathbf{y}_2) = \chi(\mathbf{b}^-)$. Next define

$$\mathbf{z}^+ = b_{1|v^+} \oplus b_{0|\text{else}} \qquad \mathbf{z}^- = b_{1|v^-} \oplus b_{0|\text{else}}$$

Then $\mathbf{z}^+ \propto \mathbf{b}^+$, $\mathbf{z}^- \propto \mathbf{b}^-$, and

$$t_0(\mathbf{a}_1, \mathbf{z}^+) = m_{1|v^+} \oplus m_{0|\text{else}} \equiv_\Theta m_{1|v^-} \oplus m_{0|\text{else}} = t_0(\mathbf{a}_1, \mathbf{z}^-)$$

We have shown that for $\mathbf{y}_1, \mathbf{y}_2 \in \mathcal{B}_1$,

$$\begin{aligned} |\chi|(\mathbf{y}_1) = |\chi|(\mathbf{y}_2) \Rightarrow \mathbf{C} \models \chi(\mathbf{y}_1) = \chi(\mathbf{y}_2) \text{ or}& \\ \exists v_3, v_4, v_5, v_6, \mathrm{WHO}(v_3) \,\&\, \mathrm{WHO}(v_4) \,\&& \\ \chi(\mathbf{y}_1) = \chi(v_3) \,\&\, \chi(\mathbf{y}_2) = \chi(v_4) \,\&& \\ v_5 \propto v_3 \,\&\, v_6 \propto v_4 \,\&& \\ t_0(a_1, v_5) = t_0(a_1, v_6)& \end{aligned}$$

Let the last formula be denoted $\mathrm{EQ}(\mathbf{y}_1, \mathbf{y}_2)$, with the understanding that the variables v_3 through v_6 are really ℓ-tuples.

Claim IV.6.17. The converse holds too; that is, the formula $\mathrm{EQ}(v_1, v_2)$ defines the equivalence relation $|\chi|(v_1) = |\chi|(v_2)$ on $\mathcal{B}_1$.

To show this, let $\chi(\mathbf{y}_1) = v^+$, say, and $\chi(\mathbf{y}_2) \notin \{v^+, v^-\}$; we must show $\neg\mathrm{EQ}(\mathbf{y}_1, \mathbf{y}_2)$. To this end, let $\mathbf{y}_3, \mathbf{y}_4 \in \mathcal{B}_1$ with $\chi(\mathbf{y}_3) = \chi(\mathbf{y}_1)$, $\chi(\mathbf{y}_4) = \chi(\mathbf{y}_2)$, and let $\mathbf{z}_5 \propto \mathbf{y}_3$, $\mathbf{z}_6 \propto \mathbf{y}_4$. Then for $i \neq \infty$, $t_0(a_1, z_5^i) = t_0(a_1, y_3^i)$ and $t_0(a_1, z_6^i) = t_0(a_1, y_4^i)$ by the definition of the relation $\propto$. Hence

$$\begin{aligned} \text{If } i = \infty \text{ then } & t_0(a_1, z_5^i) = t_0(a_1, b_0) = t_0(a_1, z_6^i) \\ \text{If } i = v^+ \text{ then } & t_0(a_1, z_5^i) = t_0(a_1, y_3^i) \neq t_0(a_1, b_0) = t_0(a_1, y_4^i) = t_0(a_1, z_6^i) \\ \text{If } i = \chi(\mathbf{y}_2) \text{ then } & t_0(a_1, z_5^i) = t_0(a_1, y_3^i) = t_0(a_1, b_0) \neq t_0(a_1, y_4^i) = t_0(a_1, z_6^i) \\ \text{Otherwise } & t_0(a_1, z_5^i) = t_0(a_1, y_3^i) = t_0(a_1, b_0) = t_0(a_1, y_4^i) = t_0(a_1, z_6^i) \end{aligned}$$

We have that $t_0(\mathbf{a}_1, \mathbf{z}_5)$ differs from $t_0(\mathbf{a}_1, \mathbf{z}_6)$ in exactly two coordinates, which do not form a V-block; hence these two elements are not Θ-congruent. This proves the claim. $\dashv_{\text{IV.6.17}}$

All that remains is to show that the edge relation is recoverable, so suppose $v \overset{E}{—} w$, $|\chi|(\mathbf{y}_1) = v$

and $|\chi|(\mathbf{y}_2) = w$. Let $\chi(\mathbf{y}_v^+) = v^+, \chi(\mathbf{y}_v^-) = v^-, \chi(\mathbf{y}_w^+) = w^+, \chi(\mathbf{y}_w^-) = w^-$, and define

$$\mathbf{b}_{vw}^+ = b_{1|v^+,w^+,\infty} \oplus b_{0|\text{else}} \qquad \mathbf{b}_{vw}^- = b_{1|v^-,w^-,\infty} \oplus b_{0|\text{else}}$$

We have $\mathbf{b}_{vw}^+, \mathbf{b}_{vw}^- \in \mathcal{B}_2$, $\mathbf{y}_v^+, \mathbf{y}_w^+ \ll \mathbf{b}_{vw}^+$, and $\mathbf{y}_v^-, \mathbf{y}_w^- \ll \mathbf{b}_{vw}^-$. Next define

$$\mathbf{z}_{vw}^+ = b_{1|v^+,w^+} \oplus b_{0|\text{else}} \qquad \mathbf{z}_{vw}^- = b_{1|v^-,w^-} \oplus b_{0|\text{else}}$$

Then $\mathbf{z}_{vw}^+ \propto \mathbf{b}_{vw}^+$, $\mathbf{z}_{vw}^- \propto \mathbf{b}_{vw}^-$, and

$$t_0(\mathbf{a}_1, \mathbf{z}_{vw}^+) = m_{1|v^+,w^+} \oplus m_{0|\text{else}} \equiv_\Theta m_{1|v^-,w^-} \oplus m_{0|\text{else}} = t_0(\mathbf{a}_1, \mathbf{z}_{vw}^-)$$

We have shown that for $\mathbf{y}_1, \mathbf{y}_2 \in \mathcal{B}_1$,

$$\begin{aligned} |\chi|(\mathbf{y}_1) \overset{E}{\text{—}} |\chi|(\mathbf{y}_2) \;\Rightarrow \exists v_3, \ldots, v_{10} \; & \bigwedge_{3\le j\le 6} v_j \in \mathcal{B}_1 \;\&\; \bigwedge_{7\le j\le 8} v_j \in \mathcal{B}_2 \;\& \\ & |\chi|(v_3) = |\chi|(v_4) = |\chi|(\mathbf{y}_1) \neq |\chi|(\mathbf{y}_2) = |\chi|(v_5) = |\chi|(v_6) \;\& \\ & \chi(v_3) \neq \chi(v_4) \;\&\; \chi(v_5) \neq \chi(v_6) \;\& \\ & v_3, v_5 \ll v_7 \;\&\; v_4, v_6 \ll v_8 \;\& \\ & v_9 \propto v_7 \;\&\; v_{10} \propto v_8 \;\&\; t_0(a_1, v_9) = t_0(a_1, v_{10}) \end{aligned}$$

Call this formula EDGE$(\mathbf{y}_1, \mathbf{y}_2)$ (again all variables v_3 through v_{10} are secretly ℓ-tuples).

Claim IV.6.18. The converse holds too; that is, the formula EDGE(v_1, v_2) recovers the edge relation of $\mathbb{G}$ on $\mathcal{B}_1/|\chi|$.

The proof is similar to the last claim's. Assume $|\chi|(\mathbf{y}_1) \neq |\chi|(\mathbf{y}_2)$ and $|\chi|(\mathbf{y}_1) \overset{E}{\not\text{—}} |\chi|(\mathbf{y}_2)$. Let $\mathbf{y}_3, \ldots, \mathbf{y}_8, \mathbf{z}_9, \mathbf{z}_{10}$ be as in the statement. Then since $\mathbf{z}_9 \propto \mathbf{y}_7$ and $\mathbf{z}_{10} \propto \mathbf{y}_8$, for all $i \neq \infty$ we have

$$t_0(a_1, y_7^i) = t_0(a_1, z_9^i) \quad t_0(a_1, y_8^i) = t_0(a_1, z_{10}^i)$$

By assumption, $\mathrm{E}^{\mathbf{S}}(b_0, z_9^\infty) = A = \mathrm{E}^{\mathbf{S}}(b_0, z_{10}^\infty)$, so in particular

$$t_0(a_1, z_9^\infty) = t_0(a_1, b_0) = t_0(a_1, z_{10}^\infty)$$

Now for $i \in V^{\pm}$

If $i \in \{\chi(\mathbf{y}_3), \chi(\mathbf{y}_5)\}$

$$\text{then } t_0(a_1, z_9^i) = t_0(a_1, y_7^i) \neq t_0(a_1, b_0) = t_0(a_1, y_8^i) = t_0(a_1, z_{10}^i)$$

If $i \in \{\chi(\mathbf{y}_4), \chi(\mathbf{y}_6)\}$

$$\text{then } t_0(a_1, z_9^i) = t_0(a_1, y_7^i) = t_0(a_1, b_0) \neq t_0(a_1, y_8^i) = t_0(a_1, z_{10}^i)$$

Otherwise

$$t_0(a_1, z_9^i) = t_0(a_1, y_7^i) = t_0(a_1, b_0) = t_0(a_1, y_8^i) = t_0(a_1, z_{10}^i)$$

Hence $t_0(\mathbf{a}_1, \mathbf{z}_9)$ differs from $t_0(\mathbf{a}_1, \mathbf{z}_{10})$ on a set of precisely four coordinates $\{v^+, v^-, w^+, w^-\}$ where $v \overset{E}{\not\frown} w$. It follows that

$$t_0(\mathbf{a}_1, \mathbf{z}_9) \not\equiv_\Theta t_0(\mathbf{a}_1, \mathbf{z}_{10})$$

which proves the Claim. $\dashv_{\text{IV.6.18}}$

We have shown that the hereditarily undecidable class of finite graphs can (with possibly finitely many exceptions) be interpreted uniformly into finite members of HSP $(\mathbf{S})$; hence this variety is hereditarily finitely undecidable. $\square$

Lemma IV.7. *The strongly-solvable radical of every finite algebra lying in a finitely-decidable variety is abelian.*

Proof. Let $\mathbf{S}$ be a counterexample of minimum possible cardinality, with strongly-solvable radical σ. We aim for a contradiction.

Claim IV.7.1. $\mathbf{S}$ is subdirectly-irreducible.

To see this, let

$$t(a_1, \vec{b}_1) = t(a_1, \vec{b}_2)$$
$$\text{but}$$
$$t(a_2, \vec{b}_1) \neq t(a_2, \vec{b}_2)$$

witness $\neg\mathrm{C}(\sigma, \sigma; \perp_S)$, and let α be maximal for separating $t(a_2, \vec{b}_1)$ from $t(a_2, \vec{b}_2)$. Then α is meet-irreducible and $\alpha \vee \sigma$ is strongly-solvable over α. Then the strongly-solvable radical of $\mathbf{S}/\alpha$ is not abelian, so if $\alpha > \perp_S$ then $\mathbf{S}/\alpha$ would be a smaller counterexample. $\dashv_{\text{IV.7.1}}$

Let μ denote the monolith of $\mathbf{S}$. Again by minimality, we also have that $\mathrm{C}(\sigma, \sigma; \mu)$. Of course, since σ is nontrivial, the monolith has unary type. By Lemma III.1, the centralizer of μ is a strongly-solvable congruence. We have that $\mathbf{S}$ satisfies all the hypotheses of Lemma IV.2, but by assumption, HSP $(\mathbf{S})$ is not finitely undecidable; hence we must have that for all $(\perp_S, \mu)$-minimal sets U, $\mathrm{C}(\sigma, \mu_{|U}; \perp)$.

Now by Lemma IV.5, we have that for any $(\perp, \mu)$-minimal set U, the action of $\mathrm{T}_U^{\mathbf{S}}$ inside any trace $N \subset U$ is trivial.

Claim IV.7.2. $\mathrm{C}(\mu, \sigma; \perp)$; equivalently, $[\mu, \sigma] = \perp$.

Suppose otherwise. Choose a witnessing package

$$t(a_1, \vec{b}_1) = t(a_1, \vec{b}_2)$$
$$\text{but}$$
$$t(a_2, \vec{b}_1) \neq t(a_2, \vec{b}_2)$$

such that a_1, a_2 belong to some trace N inside a $(\perp, \mu)$-minimal set U and the polynomial $t(v_0, \dots, v_k)$ respects U. Then it is not possible for either of the functions

$$f_i(v_0) = t(v_0, \vec{b}_i)$$

$(i = 1, 2)$ to collapse traces to points; hence these two functions are twin elements of $\mathfrak{S}_U^{\mathbf{S}}$.

But then the first line (equality) says that $f_2^{-1} \circ f_1(a_1) = a_1$, implying that $f_2^{-1} \circ f_1(N) = N$; but the second line yields $f_2^{-1} \circ f_1(a_2) \neq a_2$. This contradiction proves the claim. $\dashv_{\text{IV.7.2}}$

By Theorem 4.5 of [Kea93], Claim IV.7.2 implies that μ is σ-coherent. We have already shown that the hypothesis of the coherence property, $\&_N \, \mathrm{C}(\sigma, \mu_{|N}; \perp)$, holds; hence we have both $\mathrm{C}(\mu, \sigma; \perp)$ and $\mathrm{C}(\sigma, \mu; \perp)$.

This shows that $\mathbf{S}$ satisfies all the hypotheses of Lemma IV.6. Since our assumption was that HSP $(\mathbf{S})$ is not finitely undecidable, we must have $\mathrm{C}(\sigma, \sigma; \perp)$. But this contradicts our choice of $\mathbf{S}$ as a counterexample. □

IV–D. Proof of Theorem B

We have spent considerable effort to show that failures of abelianness in strongly-solvable congruences are bad. But could it not be the case that a strongly-solvable congruence could be abelian but not strongly-abelian? It turns out the answer is no:

Lemma IV.8 ([MV89, Theorem 7.4])**.** *If $\mathbf{F}$ is a finite algebra with a strongly-solvable congruence which is abelian but not strongly-abelian, then $HS(\mathbf{F}^2)$ contains an algebra with a strongly-solvable congruence which is not abelian.*

Proof. Let $\sigma \in \mathrm{Con}(\mathbf{F})$ be strongly-solvable and abelian, but not strongly-abelian. If σ is not abelian over some congruence beneath it, we are done; so without loss of generality σ is strongly-abelian over every nontrivial congruence $\perp < \alpha \leq \sigma$. (Else pass from $\mathbf{F}$ to its quotient by a congruence maximal for σ not being strongly-abelian over it.) We have that

$$\sigma \times \sigma = \eta_1^{-1}(\sigma) \wedge \eta_2^{-1}(\sigma)$$

is a strongly-solvable congruence of $\mathbf{F}^2$.

Let

$$c_1 = t(a_1, \vec{b}_1) \neq t(a_1, \vec{b}_2) = c_3$$
$$c_2 = t(a_2, \vec{b}_1) \neq t(a_2, \vec{b}_2) = c_1$$

witness the failure of strong abelian-ness of σ over $\perp_F$. Since σ is strongly-abelian over every nontrivial $\alpha \leq \sigma$, it follows that $c_1 \equiv_\alpha c_2 \equiv_\alpha c_3$ for all such α; in particular, there is only one congruence atom $\mu = \mathrm{Cg}(\langle c_1, c_2\rangle) = \mathrm{Cg}(\langle c_1, c_3\rangle)$ below σ.

Since $\mathrm{C}(\sigma, \sigma; \perp)$, for any polynomial $p(x) \in \mathsf{Pol}_1(\mathbf{F})$ we have

$$\begin{gathered} p(c_1) = p(t(a_1, \vec{b}_1)) = p(t(a_2, \vec{b}_1)) = p(c_2) \\ \Updownarrow \\ p(c_3) = p(t(a_1, \vec{b}_2)) = p(t(a_2, \vec{b}_2)) = p(c_1) \end{gathered} \tag{IV.8.1}$$

Our proof will proceed somewhat differently depending on whether a_1, a_2 could be chosen μ-equivalent. If this is not possible, then for all polynomials s and all $m_1 \equiv_\mu m_2$ and $\vec{u}_1 \equiv_\sigma \vec{u}_2$,

$$s(m_1, \vec{u}_1) = s(m_2, \vec{u}_2) \;\Rightarrow\; s(m_1, \vec{u}_1) = s(m_1, \vec{u}_2) = s(m_2, \vec{u}_1) \tag{IV.8.2}$$

In both cases, let $\mathbf{C} \leq \mathbf{F}^2$ be the subalgebra generated by the diagonal together with $\binom{a_1}{a_2}$. Then as subalgebras, $\mathbf{C} \leq \sigma$, and if $a_1 \equiv_\mu a_2$ then $\mathbf{C} \leq \mu$. Let $\beta \in \mathrm{Con}(\mathbf{C})$ be generated by identifying $\binom{c_1}{c_2} \equiv_\beta \binom{c_3}{c_1}$. We will show that $\sigma \times \sigma$ is not abelian over β.

Claim IV.8.3. $\binom{c_1}{c_1}$ is isolated mod β; that is, there do not exist $f \in \mathsf{Pol}_1(\mathbf{F})$ and $\binom{e_{1i}}{e_{2i}} \in \mathbf{C}$ such that

$$\begin{pmatrix} c_1 \\ c_1 \end{pmatrix} = \begin{pmatrix} f(c_1, \vec{e}_1) \\ f(c_2, \vec{e}_2) \end{pmatrix} \neq \begin{pmatrix} f(c_3, \vec{e}_1) \\ f(c_1, \vec{e}_2) \end{pmatrix} = \begin{pmatrix} d_1 \\ d_2 \end{pmatrix}$$

Suppose first that a_1 could not be chosen μ-congruent to a_2. By equation (IV.8.2), $c_1 = f(c_2, \vec{e}_1) = f(c_1, \vec{e}_2) = d_2$; it follows by equation (IV.8.1) that $c_1 = f(c_1, \vec{e}_1) = f(c_3, \vec{e}_1) = d_1$. This contradiction proves the first case of the claim.

In the other case, assume that $a_1 \equiv_\mu a_2$, so that $\mathbf{C}$ is a subalgebra of μ, which is a strongly-abelian congruence. The equality $f(c_1, \vec{e}_1) = f(c_2, \vec{e}_2)$ implies that

$$c_1 = f(c_2, \vec{e}_2) = f(c_2, \vec{e}_1) = f(c_1, \vec{e}_2) = d_2$$

Equation (IV.8.1) implies that $f(c_3, \vec{e}_1) = c_1$ too. $\dashv_{\text{IV.8.3}}$

With the previous claim in place, the following failure of the term condition

$$\begin{pmatrix} c_1 \\ c_2 \end{pmatrix} = t\left(\begin{pmatrix} a_1 \\ a_2 \end{pmatrix}, \begin{pmatrix} \vec{b_1} \\ \vec{b_1} \end{pmatrix}\right) \equiv_\beta t\left(\begin{pmatrix} a_1 \\ a_2 \end{pmatrix}, \begin{pmatrix} \vec{b_2} \\ \vec{b_2} \end{pmatrix}\right) = \begin{pmatrix} c_3 \\ c_1 \end{pmatrix}$$
$$\begin{pmatrix} c_2 \\ c_2 \end{pmatrix} = t\left(\begin{pmatrix} a_2 \\ a_2 \end{pmatrix}, \begin{pmatrix} \vec{b_1} \\ \vec{b_1} \end{pmatrix}\right) \not\equiv_\beta t\left(\begin{pmatrix} a_2 \\ a_2 \end{pmatrix}, \begin{pmatrix} \vec{b_2} \\ \vec{b_2} \end{pmatrix}\right) = \begin{pmatrix} c_1 \\ c_1 \end{pmatrix}$$

shows that $\sigma \times \sigma$ is not abelian over β. □

Proof of Theorem B. By Lemma IV.8, if $\mathbf{A}$ is any finite algebra whose strongly-solvable radical is not strongly-abelian, then $\operatorname{HSP}(\mathbf{A})$ contains a finite algebra whose strongly-solvable radical is nonabelian. By Lemma IV.7, such an algebra cannot lie in any finitely-decidable variety. □

CHAPTER V

RESIDUAL FINITENESS

The goal of this chapter is to prove Theorem C. For the remainder of this section, fix a finitely generated, finitely-decidable variety $\mathcal{V}$, say $\mathcal{V} = \mathrm{HSP}\,(\mathcal{K})$, where $\mathcal{K}$ is a finite set of finite algebras.

Lemma V.1. *$\mathcal{V}$ contains only finitely many subdirectly-irreducible finite algebras whose monolith is of boolean type.*

Proof. We will show that in fact every finite subdirectly-irreducible

$$\mathbf{S} \in \mathrm{HSP}\,(\mathcal{K})$$

with boolean-type monolith already belongs to $\mathrm{HS}(\mathcal{K})$.

So let $\mathbf{S}$ be a quotient of

$$\mathbf{B} \leq \prod_{i=1}^{p} A_i$$

where each $\mathbf{A}_i \in \mathcal{K}$ and p is the smallest number of factors for which such a representation exists; say $\mathbf{S} \cong \mathbf{B}/\pi$, where π is meet-irreducible, with upper cover μ such that $\mathrm{typ}(\pi, \mu) = 3$. The minimality of p implies that each $\hat{\eta}_i = \bigwedge_{j \neq i} \eta_j$ has no congruence θ above it such that $\mathbf{B}/\theta \cong \mathbf{S}$; in particular, for each i, $\hat{\eta}_i \vee \pi \geq \mu$.

Choose some (π, μ)-minimal set $U = e(B)$. Then U has empty tail and only one trace, so $U = \{\mathbf{x}, \mathbf{y}\}$. Let $\beta = \mathrm{Cg}(\langle \mathbf{x}, \mathbf{y} \rangle)$, and observe that $\mu = \pi \vee \beta$.

Claim V.1.1. $\mathrm{Con}(\mathbf{B}) = I[\perp, \pi] \sqcup I[\beta, \top]$.

The disjointness is obvious. Let $\theta \not\leq \pi$. Then $\theta \vee \pi \geq \mu$, and in particular identifies $\mathbf{x}$ and $\mathbf{y}$. String a chain of elements between them:

$$\mathbf{x} \equiv_\theta \mathbf{z}_1 \equiv_\pi \mathbf{z}_2 \equiv_\theta \cdots \equiv_\pi \mathbf{z}_n \equiv_\theta \mathbf{y}$$

and hit this chain with e:

$$\mathbf{x} = e(\mathbf{x}) \equiv_\theta e(\mathbf{z}_1) \equiv_\pi e(\mathbf{z}_2) \equiv_\theta \cdots \equiv_\pi e(\mathbf{z}_n) \equiv_\theta e(\mathbf{y}) = \mathbf{y}$$

The resulting chain is in U, so the π-links are trivial, implying that $\mathbf{x} \equiv_\theta \mathbf{y}$, as claimed. $\dashv_{\text{V.1.1}}$

We have already seen that $\hat{\eta}_i \not\leq \pi$ for any $1 \leq i \leq p$; by the claim, each $\hat{\eta}_i$ identifies $\mathbf{x}$ and $\mathbf{y}$. But now observe that if p were to be greater than 1, we would have

$$\langle \mathbf{x}, \mathbf{y} \rangle \in \hat{\eta}_1 \cap \hat{\eta}_2 = \perp$$

which would be absurd. Hence $p = 1$ and the theorem follows. $\square$

Lemma V.2. *$\mathcal{V}$ contains only finitely many subdirectly-irreducible finite algebras whose monolith is of affine type.*

The proof adapts from, but corrects an error in, [MS05] Section 12.

Proof. Let $\mathbf{S} \in \mathrm{HSP}(\mathcal{K})$ be subdirectly-irreducible with affine monolith; say $\mathbf{S} = \mathbf{B}/\pi$, where

$$\mathbf{B} \leq_s \prod_{i=1}^{p} \mathbf{A}_i \qquad (\mathbf{A}_i \in \mathcal{K})$$

Without loss of generality $\mathcal{K} = \mathrm{HS}(\mathcal{K})$, and the representation is minimal in the sense that $\mathbf{S}$ is not representable in this way by fewer than p factors from $\mathcal{K}$, and moreover if $\beta_i \in \mathrm{Con}(\mathbf{A}_i)$ and $\mathbf{S}$ is a quotient of a subalgebra of $\prod_i \mathbf{A}_i/\beta_i$ then all β_i are trivial.

Claim V.2.1. Let σ_i denote the strongly-solvable radical of $\mathbf{A}_i$, and $\sigma_1 \times \cdots \times \sigma_p = \sigma \in \mathrm{Con}(\mathbf{B})$. Then $\sigma \leq \pi$.

Suppose this were false. Let

$$\perp_B \leq \alpha^- \overset{1}{\prec} \alpha^+ \leq \sigma$$

such that $\alpha^- \leq \pi$ but $\pi \leq \beta^- \overset{2,3}{\prec} \beta^+ = \alpha^+ \vee \pi$. Then the covers $\alpha^- \overset{1}{\prec} \alpha^+$ and $\beta^- \overset{2,3}{\prec} \beta^+$ are projective, which is absurd (cf. Theorem II.5). $\dashv_{\mathrm{V.2.1}}$

Our minimality assumption implies now that each $\mathbf{A}_i$ in the representation of $\mathbf{B}$ has trivial strongly-solvable radical. By Lemma II.7, $\mathbf{B}$ has Day polynomials; hence the term condition on congruences of $\mathbf{B}$ is symmetric in the first two variables.

It follows by Theorem 10.1 of [FM87] that $\mathbf{S}/\zeta \in \mathrm{HS}(\mathcal{K})$, where ζ denotes the centralizer of the monolith μ; in particular,

$$|\mathbf{S}/\zeta| \leq \max\{|A| \colon \mathbf{A} \in \mathcal{K}\}$$

We will be done if we can show that there is also a bound on the number of elements of each ζ-block. From now on we will forget about $\mathbf{B}$ and work only in $\mathbf{S}$. Let $\{C_i = r_i/\zeta : 1 \leq i \leq \ell\}$ be an injective enumeration (with fixed representatives) of the ζ-classes, C any fixed one of them, and U a $(\perp_S, \mu)$-minimal set containing a monolith pair $\{0, a\}$.

As before, we have a Malcev polynomial $m(v_1, v_2, v_3)$ on U; furthermore, if $Q \subseteq U$ denotes the ζ-class of 0 in U, then m respects Q. Since the tail of U is empty, $\mathbf{S}_{|U}$ is then an abelian Malcev algebra. By a standard argument, the operation $m(x, y, z) = x - y + z$ defines an abelian group operation on Q under which 0 is the identity element.

Claim V.2.2. The set of polynomial functions

$$R = \{f(v) \in \mathsf{Pol}_1(\mathbf{S}_{|Q}) : f(0) = 0\}$$

is a ring of endomorphisms of $\mathbf{Q}$ (under pointwise addition and function composition), and the size of R is bounded independent of $\mathbf{S}$.

The only nontrivial piece of the first part is that each such f respects addition:

$$f(x) = f(x - y + y) + f(0) = f(y - y + 0) + f(x) = f(x)$$
$$\Downarrow$$
$$f(x + y) = f(x - 0 + y) + f(0) = f(y - 0 + 0) + f(x) = f(y) + f(x)$$

The second comes from the fact that each $f \in R$ is given by an $\ell + 1$-ary term operation in a uniform way: if $f(x) = t(x, \vec{s})$ then

$$0 = t(0, \vec{s}) - t(0, \vec{s}) = t(0, \vec{r}) - t(0, \vec{r})$$
$$\Downarrow$$
$$t(x, \vec{s}) = t(x, \vec{s}) - t(0, \vec{s}) = t(x, \vec{r}) - t(0, \vec{r})$$

where $\vec{r}$ denotes the chosen representatives of the ζ-classes. Hence $|R| \leq |\mathbf{F}_{\mathcal{V}}(1 + \ell)|$. $\dashv_{\text{V.2.2}}$

Now: for any $s_1 \neq s_2 \in S$, there exists a polynomial $f(v_0) = t(v_0, \vec{s})$ so that $t(s_1, \vec{s}) = 0$ and $t(s_2, \vec{s}) = a$. In particular, if $s_1 = 0, s_2 \in Q$ then we may take $f \in R$.

What this shows is that $\mathbf{Q}$ is subdirectly-irreducible as an R-module. By Theorem 1 of [Kea91], $|Q| \leq |R|$.

Now we are almost done: we have already noted that for each $c, d \in C$ there exists a term $t(v_0, \ldots, v_{|S|})$ with $t(c, \vec{s}) = 0$, $t(d, \vec{s}) = a$. One has

$$et(d, \vec{s}) - et(d, \vec{s}) = et(d, \vec{r}) - et(d, \vec{r})$$
$$\Downarrow$$
$$a = et(d, \vec{s}) - et(c, \vec{s}) = et(d, \vec{r}) - et(c, \vec{r})$$

where all these values must lie in Q. Hence the map

$$\begin{aligned} &C \to {}^{\mathbf{F}_{\mathcal{V}}(1+\ell)}Q \\ &x \mapsto \langle et(x, \vec{r}) \colon t \in \mathbf{F}_{\mathcal{V}}(1 + \ell) \rangle \end{aligned}$$

is injective.

We have shown

$$|C| \leq |Q|^{|\mathbf{F}_{\mathcal{V}}(1+\ell)|} \leq |R|^{|\mathbf{F}_{\mathcal{V}}(1+\ell)|} \leq |\mathbf{F}_{\mathcal{V}}(1 + \ell)|^{|\mathbf{F}_{\mathcal{V}}(1+\ell)|}$$

which, combined with the fact that

$$|S| \leq |C| \cdot \max_{\mathbf{A} \in \mathcal{K}}(|A|)$$

completes the proof. $\square$

We will need the following technical lemma limiting the number of variables which can be independent (modulo a strongly-abelian congruence) in a polynomial operation.

Lemma V.3. *Let* $\mathbf{A}$ *be a finite algebra in a locally-finite variety* $\mathcal{V}$*, and* β *a strongly-abelian congruence on* $\mathbf{A}$*, and* $t(v_0, \vec{v}_1, \ldots, \vec{v}_\ell)$ *be any polynomial operation of* $\mathbf{A}$*. Let* $M = \log|\mathbf{F}_\mathcal{V}(\ell+2)|$*. Then there exist subsets* $\breve{v}_i \subset \vec{v}_i$ *of size no more than* M*, such that for any* β*-blocks* $B_1, \ldots, B_\ell$ *the mapping*

$$\begin{aligned} A \times \vec{B}_1 \times \cdots \times \vec{B}_\ell &\to A \\ \langle a, \vec{b}_1, \ldots, \vec{b}_\ell \rangle &\mapsto t(a, \vec{b}_1, \ldots, \vec{b}_\ell) \end{aligned} \tag{V.3.1}$$

depends only on the variables v_0 *and* $\breve{v}_i$*.*

Proof. For simplicity, we show the case $\ell = 2$. Let $t(v_0, v_1^1, \ldots, v_1^{k_1}, v_2^1, \ldots, v_2^{k_2})$ be our term, and let $2^{k_1} > |\mathbf{F}_\mathcal{V}(4)|$.

For $S \subset \{1, \ldots, k_1\}$ let $p_S(v_0, x, y, v_2)$ be the substitution instance of t obtained by identifying all v_2^i to the single variable v_2, and substituting x for v_1^i if $i \in S$ and y if not. Then by Pigeonhole, there exist $S \neq S'$ so that $\mathcal{V} \models p_S = p_{S'}$. Say $k_1 \in S$ but not S'; we claim that no mapping as in (V.3.1) can depend on $v_1^{k_1}$.

To see this, let $a \in A$, $b, c \in B_1$, and $d \in B_2$. Let $q_S(v_0, x, y, v_1^{k_1}, v_2)$ be like p_S, except that $v_1^{k_1}$ is left unsubstituted, and likewise for $q_{S'}$. Then

$$q_S(a, b, c, b, d) = q_{S'}(a, b, c, c, d)$$

But now since β is strongly-abelian, if $\vec{x} \equiv_\beta b$ and $\vec{y} \equiv_\beta d$, the strong term condition gives that

$$t(a, \vec{x}, b, \vec{y}) = t(a, \vec{x}, c, \vec{y})$$

so t is insensitive to changes modulo β in the $v_1^{k_1}$ coordinate. Similarly, if $2^{k_2} > |\mathbf{F}_\mathcal{V}(4)|$ then t is insensitive to changes mod β in some coordinate v_2^i. The general result now follows by a downward induction. ☐

Lemma V.4. $\mathcal{V}$ *contains only finitely many subdirectly-irreducible finite algebras whose monolith is of unary type.*

Proof. Let $\mathbf{S} \in \mathcal{V}$ be subdirectly-irreducible with unary-type monolith

$$\mu = \mathrm{Cg}_\mathbf{S}(\langle c, d \rangle)$$

We already know that typ $\{\mathbf{S}\} \subset \{1, 3\}$. By Theorem B, the strongly-solvable radical σ is a strongly-abelian congruence. Theorem A tells us that either $\sigma = \top_S$ or is meet-irreducible with upper cover of boolean type. In either case,

$$\ell := |\mathbf{S}/\sigma| \leq M_{\mathrm{bool}}$$

(where M_{bool} denotes the maximum cardinality of a finite SI in $\mathcal{V}$ with boolean-type monolith). Fix some enumeration $\langle \vec{s}_1, \ldots, \vec{s}_\ell \rangle$ of S with each σ-block B_i enumerated together. We must now put a uniform bound on the size of σ-blocks.

Let B be any σ-block. Since any unequal pair of elements generates a congruence above μ, we have that for any $b \neq b' \in B$, there exists a unary polynomial $p(v_0) = t(v_0, \vec{s}_1, \ldots, \vec{s}_\ell)$ such that $p(b) = c$ iff $p(b') \neq c$. By Lemma V.3, these terms depend (up to changes mod σ) on v_0 and subsets $\breve{s}_i$, $1 \leq i \leq \ell$, each of size no more than $M := \log(\mathbf{F}_{\mathcal{V}}(\ell + 2))$. Let $P = B_1^M \times \cdots \times B_\ell^M$.

For $b \in B$, we define a subset $G(b) \subset \mathbf{F}_{\mathcal{V}}(1 + \ell M)$ to consist of those terms $t(x, \vec{y})$ such that for some $\vec{p} \in P$, $t(b, \vec{p}) = c$. Observe that for any $b_1 \neq b_2$, at least one of $G(b_1), G(b_2)$ is nonempty.

Claim V.4.1. The mapping $b \mapsto G(b)$ is injective.

Let $b_1 \neq b_2$, and assume towards a contradiction that $G(b_1) = G(b_2)$. Choose a term t and a $\vec{p}_1 \in \Sigma$ so that $c = t(b_1, \vec{p}_1) \neq t(b_2, \vec{p}_1)$. Then $t \in G(b_1) = G(b_2)$, so we can choose $\vec{p}_2 \in \Sigma$ so that $t(b_2, \vec{p}_2) = c$. Hence we have a failure

$$\begin{array}{ll} c = t(b_1, \vec{p}_1) & t(b_1, \vec{p}_2) \\ c \neq t(b_2, \vec{p}_1) & t(b_2, \vec{p}_2) = c \end{array}$$

of the strong term condition, since the entries are equal along the diagonal but not along the rows and columns. This contradicts the strong abelianness of σ. $\dashv_{\text{V.4.1}}$

We have just shown that

$$|B| \leq 2^{|\mathbf{F}_{\mathcal{V}}(1+\ell M)|}$$

which is uniformly bounded in $\mathcal{V}$. This completes the proof. □

Proof of Theorem C. Since $\mathcal{V}$ is locally-finite, it suffices to prove that $\mathcal{V}$ contains only finitely many finite subdirectly-irreducible algebras. (It is a well-known result, originally due to Quackenbush, that an infinite SI algebra in a locally-finite variety has arbitrarily large finite SI subalgebras generated by a monolith pair together with other elements.) Since $\mathcal{V}$ is finitely-decidable, it omits the semilattice and lattice types altogether; and Lemmas V.1, V.2, and V.4 combine to show that there are only finitely many SIs in $\mathcal{V}$ with monoliths of the boolean, affine, or unary types. □

CHAPTER VI

TWO MULTI-SORTED CONSTRUCTIONS

We will be building two multi-sorted languages from which to effect an interpretation. While it is possible to formalize multi-sorted model theory entirely in a usual first-order setting, this formalization takes away much of the naturality of the multi-sorted definition. In particular, the first-order formalization "gets wrong" the structural operations of direct product and substructure; these are key for us, since we will be constructing varieties in our sorted model classes.

Definition VI.1. For our purposes, if L is a finite multi-sorted first-order language, every sort must have nonempty extension in every L-structure. It follows (cf [ARV12]) that the Birkhoff variety theorem holds without modification for L-structures.

Notation VI.2. Every atomic formula $\Phi(v_1, v_2, \ldots)$ of a multi-sorted language must implicitly or explicitly determine what sort each variable must be assigned from. We call this the *type signature* of the formula. In particular, for a term t we write

$$t(S_1, S_2, \ldots) \to S_0$$

to denote that the formula

$$t(x_1, x_2, \ldots) = x_0$$

is meaningful only if $x_0 \in S_0$, $x_1 \in S_1$, $x_2 \in S_2$, and so forth.

For the remainder of this and the next section, fix a finite (single-sorted) algebraic language L and a finite L-algebra $\mathbf{A}$ with a congruence α whose congruence classes are $C_1, \ldots, C_M$.

VI–A. The sorted language L^α corresponding to a congruence α

Definition VI.3. The multi-sorted first-order language L^α will have the following nonlogical symbols:

For each $1 \le i \le M$, the language will have a sort symbol $\langle i \rangle$.

For each basic operation symbol $f(v_1, \ldots, v_n)$ of L and all indices $1 \le i_1, \ldots, i_n \le M$, L^α will have a basic operations symbol $f_{i_1 \cdots i_n}$ of type signature

$$f_{i_1 \cdots i_n}(\langle i_1 \rangle, \langle i_2 \rangle, \ldots, \langle i_n \rangle) \to \langle i_0 \rangle$$

where

$$C_{i_1} \times \cdots \times C_{i_n} \xrightarrow{t^{\mathbf{A}}} C_{i_0}.$$

Construction VI.4. (1) We define an L^α-structure $\mathbf{A}^\alpha$ in the natural way: each sort

$$\langle i \rangle^{\mathbf{A}^\alpha} = C_i$$

and if $x_k \in C_{i_k}$ for $1 \leq k \leq n$,

$$f^{\mathbf{A}^\alpha}_{i_1 \cdots i_n}(x_1, \ldots, x_n) = f^{\mathbf{A}}(x_1, \ldots, x_n).$$

(2) More generally, let $\mathbf{B}$ be any L-structure such with a congruence $\alpha^{\mathbf{B}}$ such that there exists an isomorphism $\varphi\colon \mathbf{A}/\alpha \to \mathbf{B}/\alpha^{\mathbf{B}}$. Define an L^α-structure $\mathbf{B}^\alpha$ by declaring

$$\langle i \rangle^{\mathbf{B}^\alpha} = \varphi(C_i)$$

and defining the basic operations

$$f^{\mathbf{B}^\alpha}_{i_1 \cdots i_n}(x_1, \ldots, x_n) = f^{\mathbf{B}}(x_1, \ldots, x_n)$$

for any $x_k \in \varphi(C_{i_k})$. Note that the isomorphism φ will usually be clear in practice, so we do not include it as a visible parameter in the symbol $\mathbf{B}^\alpha$. Similarly, we will usually refer to the distinguished congruence of $\mathbf{B}$ as α rather than $\alpha^{\mathbf{B}}$.

Proposition VI.5. *Let $\mathbf{M} = \mathbf{B}^\alpha$ and $\mathbf{N} = \mathbf{C}^\alpha$.*

(1) Let $\mathbf{D} \leq \mathbf{B}$ have nonempty intersection with each α-class; then $\mathbf{D}$ satisfies the hypotheses of Construction VI.4(2), and $\mathbf{D}^\alpha$ is a substructure of $\mathbf{M}$. Moreover, every substructure of $\mathbf{M}$ is obtained in this way.

(2) Let $\theta \leq \alpha$ be a congruence on $\mathbf{B}$; then $\mathbf{B}/\theta$ satisfies the hypotheses of Construction VI.4(2), and $(\mathbf{B}/\theta)^\alpha$ is a homomorphic image of $\mathbf{M}$. Moreover, every homomorphic image of $\mathbf{M}$ is obtained in this way.

(3) Let $\mathbf{D} \leq \mathbf{B} \times \mathbf{C}$ be the subalgebra consisting of all pairs $\langle b, c \rangle$ such that $\varphi^{-1}(b/\alpha) = \varphi^{-1}(c/\alpha)$. Then $\mathbf{D}$ satisfies the hypotheses of Construction VI.4(2), and $\mathbf{D}^\alpha$ is the product of $\mathbf{M}$ and $\mathbf{N}$ in the sense of L^α. (This generalizes to any number of factors.)

As mentioned above, the classical proof that a class is equationally axiomatizable iff it is closed under taking products, substructures, and homomorphic images is valid for multi-sorted algebras, so it makes sense to talk about the variety $\mathcal{V}(\mathbf{A}^\alpha) = \mathrm{HSP}\,(\mathbf{A}^\alpha)$. A representation of the free algebras in this variety as subalgebras of a direct power of $\mathbf{A}^\alpha$, where the index set is itself a power of $\mathbf{A}^\alpha$, does exist; but is not straightforward to write down, and one is better off thinking of free algebras as algebras of terms. Note that the trivial algebra in this variety is the one where each sort is a singleton, i.e. $(\mathbf{A}/\alpha)^\alpha$.

Lemma VI.6. *(1) The sorted structure $\mathbf{B}^\alpha$ is abelian (resp. strongly-abelian) if and only if the congruence α was a (strongly) abelian congruence of $\mathbf{B}$.*

(2) If $\mathbf{A}$ belongs to a finitely-decidable variety and α is a (strongly) solvable congruence, then $\mathrm{HSP}\,(\mathbf{A}^\alpha)$ *is a (strongly) abelian variety.*

Proof. (1): A failure of the (strong) term condition $\mathrm{C}(\alpha, \alpha; \perp)$ in $\mathbf{B}$ is readily convertible into a failure of the corresponding condition $\mathrm{C}(\top, \top; \perp)$ in $\mathbf{B}^\alpha$, and vice versa.

(2): By Theorem B, (strongly) solvable congruences in $\mathrm{HSP}(\mathbf{A})$ are (strongly) abelian.

If $\mathrm{HSP}(\mathbf{A}^\alpha)$ were to fail to be (strongly) abelian, this failure would be witnessed in a finitely generated, and hence finite, structure $\mathbf{M}$. We may suppose $\mathbf{M} = \mathbf{N}/\vartheta$, where $\mathbf{N}$ is a substructure of a direct power $(\mathbf{A}^\alpha)^X$.

As we saw in Lemma VI.5, this direct power is the image under $\bullet^\alpha$ of the subalgebra $\mathbf{P}$ of $\mathbf{A}^X$ consisting of all α-constant tuples. Since any failure of (strong) abelianness would project to a failure at some coordinate,

$$\alpha^{\mathbf{P}} = \alpha^X \cap (P \times P)$$

is (strongly) abelian. Hence $(\mathbf{A}^\alpha)^X$ is (strongly) abelian.

We know that $\mathbf{N} = \mathbf{B}^\alpha$ for some $\mathbf{B} \leq \mathbf{P}$, and moreover that

$$\alpha^{\mathbf{B}} = \alpha^{\mathbf{P}} \cap (B \times B);$$

it follows any failure of (strong) abelianness in $\mathbf{B}$ would have represented one in $\mathbf{P}$ already. Hence $\mathbf{N}$ is (strongly) abelian.

Finally, we have that there must exist $\theta \in \mathrm{Con}(\mathbf{B})$ such that $(\mathbf{B}/\theta)^\alpha = \mathbf{N}/\vartheta = \mathbf{M}$. But since α is (strongly) abelian in $\mathbf{B}$, θ is (strongly) solvable, and hence (strongly) abelian as well; and just as in (1) any witness to the failure of the (strong) term condition $\mathrm{C}(\top, \top; \vartheta)$ in $\mathbf{N}$ would give rise to a failure of the corresponding condition $\mathrm{C}(\alpha, \alpha; \theta)$ in $\mathbf{B}$. $\square$

Corollary VI.7. *If* $\mathbf{A}$ *belongs to any finitely-decidable variety and* α *is either the solvable radical or the strongly-solvable radical of* $\mathbf{A}$*, then* $\mathrm{HSP}(\mathbf{A}^\alpha)$ *semantically interprets into* $\mathrm{HSP}(\mathbf{A})$*.*

Proof. The key observation is that, by Lemma IV.1, each of the congruences in the statement of the theorem is uniformly definable in $\mathrm{HSP}(\mathbf{A})$, and our construction guarantees that $\alpha^{\mathbf{B}}$ is the (strongly) solvable radical of $\mathbf{B}$ whenever α was of $\mathbf{A}$.

Let $c_1, \ldots, c_M$ be new constant symbols. Take any $\mathbf{M} = \mathbf{B}^\alpha \in \mathrm{HSP}(\mathbf{A}^\alpha)$, where $\mathbf{M}$ and $\mathbf{B}$ can be taken to be on the same underlying set. First, assign c_i to an arbitrary element of $\varphi(C_i)$ for each i. Then one can recover the sort of x by asserting that x and c_i are congruent modulo the radical; likewise the assertion $f_{i_1 \cdots i_n}(x_1, \ldots, x_n) = x_0$ is true in $\mathbf{M}$ iff each $x_k \equiv_{\mathrm{Rad}(\mathbf{B})} c_{i_k}$ and $f(x_1, \ldots, x_n) = x_0$ in $\mathbf{B}$. $\square$

It follows that whenever $\mathrm{HSP}(\mathbf{A}^\alpha)$ is finitely axiomatizable (which happens, for instance, when the variety is strongly-abelian), then the (finite) undecidability of $\mathrm{HSP}(\mathbf{A}^\alpha)$ implies the (finite) undecidability $\mathrm{HSP}(\mathbf{A})$ too.

VI–B. The sorted language $\mathbf{A}^{\alpha\flat}$ corresponding to a strongly-abelian congruence α

The construction in the previous section required no assumptions about α. If, however, α is strongly-abelian, then we can introduce a further sorted construction, generalizing that effected by McKenzie and Valeriote in [MV89, Chapter 11]. For the remainder of this section, we add the assumption that α is strongly-abelian.

Recall (Proposition II.13) that each induced algebra

$$\mathbf{A}_{|C_i} = \langle C_i\ ;\ \{f \in \mathsf{Pol}(\mathbf{A})\colon f(C_i, \ldots, C_i) \subseteq C_i\}\rangle$$

is a strongly-abelian algebra. For each $1 \leq i \leq M$, let K_i be the greatest arity of a decomposition α-sorted term operation on C_i. (See Definition II.9 for a review of decomposition operators.) Fix operators

$$d_i(v_1, \ldots, v_{K_i}) = D_i(v_1, \ldots, v_n, \vec{a})$$

witnessing this; that is, d_i is a K_i-ary decomposition operator on C_i and $D_i(\vec{x}, \vec{a}) = D_i(\vec{x}, \vec{a}')$ whenever $\vec{x} \in C_i$ and $\vec{a} \equiv_\alpha \vec{a}'$. By Proposition II.11, this determines a product decomposition

$$C_i = C_{i,1} \times \cdots \times C_{i,K_i}.$$

Definition VI.8. The multi-sorted first-order language $L^{\alpha\flat}$ will have the following nonlogical symbols:

For each $1 \leq i \leq M$ and each $1 \leq j \leq K_i$, the language will have a sort symbol $\langle i, j\rangle$.

For each α-sorted term operation

$$f(v_1, \ldots, v_n) = t(v_1, \ldots, v_n, \vec{a})\colon C_{i_1} \times \cdots C_{i_n} \to C_{i_0}$$

($\vec{a} \in C_{i_{n+1}} \times \cdots \times C_{i_{n'}}$) and each $1 \leq j \leq K_{i_0}$ the language $L^{\alpha\flat}$ will have a basic operation of type declaration

$$t^j_{i_1 \cdots i_n i_{n+1} \cdots i_{n'}} \begin{pmatrix} \langle i_1, 1\rangle & \langle i_2, 1\rangle & \cdots & \langle i_n, 1\rangle \\ \langle i_1, 2\rangle & \langle i_2, 2\rangle & \cdots & \langle i_n, 2\rangle \\ \vdots & \vdots & \ddots & \vdots \\ \langle i_1, K_{i_1}\rangle & \langle i_2, K_{i_2}\rangle & \cdots & \langle i_n, K_{i_n}\rangle \end{pmatrix} \to \langle i_0, j\rangle\,.$$

Note that every term $t(v_1, \ldots, v_n)$ of $\mathbf{A}$ is automatically a α-sorted term operation when restricted to any product of n α-classes, so the entire atomic diagram of $\mathbf{A}$ is encoded in that of $\mathbf{A}^{\alpha\flat}$. We will see in a moment that $\mathbf{A}^{\alpha\flat}$ is strongly-abelian; it follows that the language $L^{\alpha\flat}$ may be taken to be finite.

Proposition VI.9. *Every term in the language $L^{\alpha\flat}$ is obtained from one of the basic operations $t^j_{i_1 \cdots i_n}$ by possibly identifying some variables of the same sort.*

The proof (by induction) is left to the reader.

Construction VI.10. (1) We define an $L^{\alpha\flat}$-structure $\mathbf{A}^{\alpha\flat}$ analogously to our definition of $\mathbf{A}^{\alpha}$ in Construction VI.4(1): each sort

$$\langle i,j\rangle^{\mathbf{A}^{\alpha\flat}} = C_{i,j}$$

Now if $t^{j}_{i_1\cdots i_n i_{n+1}\cdots i_{n'}}$ is a basic operation symbol and $x_{k,j} \in C_{i_k,j}$ for $1 \le k \le n$ and $1 \le j \le K_{i_k}$, set

$$x_k = \begin{pmatrix} x_{k,1} \\ x_{k,2} \\ \vdots \\ x_{k,K_{i_k}} \end{pmatrix} \qquad (1 \le k \le n)$$

and choose any $\vec{a} \in C_{i_{n+1}} \times \cdots \times C_{i_{n'}}$. Let

$$t^{\mathbf{A}}(x_1,\dots,x_n,\vec{a}) = x_0 = \begin{pmatrix} x_{0,1} \\ x_{0,2} \\ \vdots \\ x_{0,K_{i_0}} \end{pmatrix}$$

It now makes sense to define

$$t^{j}_{i_1\cdots i_n i_{n+1}\cdots i_{n'}} \begin{pmatrix} x_{1,1} & x_{2,1} & \cdots & x_{n,1} \\ x_{1,2} & x_{2,2} & \cdots & x_{n,2} \\ \vdots & \vdots & \ddots & \vdots \\ x_{1,K_{i_1}} & x_{2,K_{i_2}} & \cdots & x_{n,K_{i_n}} \end{pmatrix} = x_{0,j}$$

(2) The foregoing construction generalizes to any L-structure $\mathbf{B}$ having a congruence $\alpha^{\mathbf{B}}$ such that there exists an isomorphism $\varphi\colon \mathbf{A}/\alpha \to \mathbf{B}/\alpha^{\mathbf{B}}$, and such that the same terms $D_i(v_1,\dots,v_{K_i},\dots,v_{n'})$ define decomposition α-sorted terms on the classes $\varphi(C_i)$, with constants taken from the same classes $\varphi(C_{i_{n+1}}),\dots,\varphi(C_{i_{n'}})$. (We do not require that no decomposition operator on $\varphi(C_i)$ have larger arity.)

Under these hypotheses, each $\alpha^{\mathbf{B}}$ class $\varphi(C_i)$ decomposes into a product of K_i factors as above, and the analogous definition produces a well-defined $L^{\alpha\flat}$-structure $\mathbf{B}^{\alpha\flat}$.

We state without proof the analogues of the lemmata of Section VI–A., since all the proofs differ only in the bookkeeping:

Proposition VI.11. *Let* $\mathbf{M} = \mathbf{B}^{\alpha\flat}$ *and* $\mathbf{N} = \mathbf{C}^{\alpha\flat}$.

(1) Let $\mathbf{D} \le \mathbf{B}$ *have nonempty intersection with each* α*-class; then* $\mathbf{D}$ *satisfies the hypotheses of Construction VI.10(2), and* $\mathbf{D}^{\alpha\flat}$ *is a substructure of* $\mathbf{M}$*. Moreover, every substructure of* $\mathbf{M}$ *is obtained in this way.*

(2) Let $\theta \leq \alpha$ be a congruence on $\mathbf{B}$; then $\mathbf{B}/\theta$ satisfies the hypotheses of Construction VI.10(2), and $(\mathbf{B}/\theta)^{\alpha\flat}$ is a homomorphic image of $\mathbf{M}$. Moreover, every homomorphic image of $\mathbf{M}$ is obtained in this way.

(3) Let $\mathbf{D} \leq \mathbf{B} \times \mathbf{C}$ be the subalgebra consisting of all pairs $\langle b, c\rangle$ such that $\varphi^{-1}(b/\alpha) = \varphi^{-1}(c/\alpha)$. Then $\mathbf{D}$ satisfies the hypotheses of Construction VI.10(2), and $\mathbf{D}^{\alpha\flat}$ is the product of $\mathbf{M}$ and $\mathbf{N}$ in the sense of L^{α}. (This generalizes to any number of factors.)

Lemma VI.12. *(1) The smallest equationally axiomatizable class containing $\mathbf{A}^{\alpha\flat}$ is the closure of $\left\{\mathbf{A}^{\alpha\flat}\right\}$ under* HSP*; this class is axiomatized by the set of all equations which hold in $\mathbf{A}^{\alpha\flat}$. This variety is finitely axiomatizable.*

(2) The sorted structure $\mathbf{B}^{\alpha\flat}$ is abelian (resp. strongly-abelian) if and only if the congruence α was a (strongly) abelian congruence of $\mathbf{B}$.

(3) If $\mathbf{A}$ belongs to a finitely-decidable variety and α is a (strongly) solvable congruence, then HSP $\left(\mathbf{A}^{\alpha\flat}\right)$ *is a (strongly) abelian variety.*

(4) If $\mathbf{A}$ belongs to any finitely-decidable variety and α is the strongly-solvable radical of $\mathbf{A}$, then HSP $\left(\mathbf{A}^{\alpha\flat}\right)$ *semantically interprets into* HSP $(\mathbf{A})$.

Proof. The only new statement here is that HSP $\left(\mathbf{A}^{\alpha\flat}\right)$ is finitely axiomatizable.

It is well known (e.g. [MV89, Theorem 0.17]) that an (ordinary single-sorted) algebra $\mathbf{X}$ is strongly-abelian if and only if for each term $t(v_1, \ldots, v_n)$ there exist equivalence relations $E_1, \ldots, E_n$ on X such that for all $x_1, y_1 \ldots, x_n, y_n \in X$,

$$t(x_1, \ldots, x_n) = t(y_1, \ldots, y_n) \iff \langle x_1, y_1\rangle \in E_1, \ldots, \langle x_n, y_n\rangle \in E_n.$$

Likewise, a congruence α is strongly-abelian iff for each term t and all α-classes

$$C_{i_1} \times \cdots \times C_{i_n} \xrightarrow{t} C_0$$

there exist equivalence relations E_k on C_{i_k} such that for all $x_k, y_k \in C_{i_k}$,

$$t(x_1, \ldots, x_n) = t(y_1, \ldots, y_n) \iff \langle x_1, y_1\rangle \in E_1, \ldots, \langle x_n, y_n\rangle \in E_n.$$

It follows that such a term action cannot depend on more than $\log_2(|C_{i_0}|)$ of its variables; in HSP $\left(\mathbf{A}^{\alpha\flat}\right)$, this means that the basic operation $t^{j}_{i_1 \cdots i_n}$ can only depend essentially on at most $\log_2(|C_{i_0}|) {\cdot} \max_i K_i$ variables. Since there are only finitely many equations using this many variables, and since HSP $\left(\mathbf{A}^{\alpha\flat}\right)$ is axiomatized by the subset of these which are true in $\mathbf{A}^{\alpha\flat}$, we are done. $\square$

CHAPTER VII

ARITY BOUNDS IN HSP $\left(\mathbf{A}^{\sigma\flat}\right)$

This chapter is devoted to the proof of Theorem D. Most of the technical work is done by the following theorem, whose proof will occupy the first two sections:

Theorem VII.1. *Let* $\mathbf{A}$ *be a finite algebra in a variety where every strongly-solvable congruence is strongly-abelian. Let* σ *be the strongly-solvable radical of* $\mathbf{A}$*,*

$$C_{i_1} \times \cdots \times C_{i_n} \xrightarrow{t} C_{i_0} \tag{VII.1.1}$$

be any σ*-sorted term operation, and let* K_{i_0} *be the greatest arity of a* σ*-sorted decomposition term on* C_{i_0}*. If the map in* (VII.1.1) *depends essentially on more than* K_{i_0} *variables, then the class of bipartite graphs interprets semantically into* HSP $\left(\mathbf{A}^{\sigma\flat}\right)$*.*

The construction and proof that follows is based on a construction first developed by Matt Valeriote for his thesis, and subsequently used in [MV89] and [HV91].

VII–A. Preparatory lemmas

In this section, let $\mathbf{A}$ be a fixed finite algebra satisfying the hypotheses of Theorem VII.1. As before, we choose a fixed enumeration $C_1, \ldots, C_M$ of the σ-classes. Fix σ-sorted decomposition terms

$$d_i(v_1, \ldots, v_{K_i}) = D_i(v_1, \ldots, v_{K_i}, \vec{a})\colon C_i^{K_i} \to C_i$$

of maximal arity.

Proposition VII.2. *The algebra* $\mathbf{A}^{\sigma\flat}$ *is essentially unary if and only if every* σ*-sorted term operation*

$$C_{i_1} \times \cdots \times C_{i_n} \xrightarrow{t} C_{i_0} \tag{VII.2.1}$$

depends on at most K_{i_0} *variables.*

Proof. We prove each contrapositive.

($\Rightarrow$): Let the action of $t(v_1, \ldots, v_{K_{i_0}+1}, \ldots)$ on the box in Equation (VII.2.1) depend essentially on at least the first $K_{i_0}+1$ variables. Choose a witnessing assignment

$$t(a, b_2, \ldots, b_n) \neq t(a', b_2, \ldots, b_n)$$

in the first variable: then for some $1 \leq j \leq K_{i_0}$,

$$t(a, b_2, \ldots, b_n) \not\sim_j t(a', b_2, \ldots, b_n)$$

where

$$x \sim_j y \iff x = \begin{matrix} x^1 \\ x^2 \\ \vdots \\ x^K \end{matrix}, y = \begin{matrix} y^1 \\ y^2 \\ \vdots \\ y^K \end{matrix} \text{ and } x^j = y^j \tag{VII.2.1}$$

For this j, the term $t^j_{i_1 \cdots i_n}$ depends on one of the variables in its first column. Similarly, for each of the variables $v_2, \ldots, v_{K_{i_0}+1}$ one of the terms $t^j_{i_1 \cdots i_n}$ depends on a variable in the corresponding column. Now use the pigeonhole principle to get one of the $t^j_{i_1 \cdots i_n}$ depending on at least two variables.

($\Leftarrow$):

Claim VII.2.2. If $t^j_{i_1 \cdots i_n}$ depends in $\mathbf{A}^{\sigma\flat}$ on the variable in column c and row r, then in $\mathbf{A}$ the operation

$$d_{i_0}\left(y_1, \ldots, y_{j-1}, t\left(d_{i_1}\left(x_1^1, \ldots, x_1^{K_{i_1}}\right), \ldots, d_{i_n}\left(x_n^1, \ldots, x_n^{K_{i_n}}\right)\right), \ldots, y_{K_{i_0}}\right)$$

depends on x^r_c (as well as on each of the ys).

To see this, pick a witnessing package

$$t^j_{i_1 \cdots i_n}\begin{pmatrix} b_1^1 & b_2^1 & \cdots & b_n^1 \\ b_1^2 & b_2^2 & \cdots & b_n^2 \\ \vdots & \vdots & a & \vdots \\ b_1^{K_{i_1}} & b_2^{K_{i_2}} & \cdots & b_n^{K_{i_n}} \end{pmatrix} \neq t^j_{i_1 \cdots i_n}\begin{pmatrix} b_1^1 & b_2^1 & \cdots & b_n^1 \\ b_1^2 & b_2^2 & \cdots & b_n^2 \\ \vdots & \vdots & a' & \vdots \\ b_1^{K_{i_1}} & b_2^{K_{i_2}} & \cdots & b_n^{K_{i_n}} \end{pmatrix}$$

Upstairs in $\mathbf{A}$ this becomes

$$t(b_1, \ldots, \hat{a}, \ldots, b_n) \not\sim_j t(b_1, \ldots, \hat{a}', \ldots, b_n)$$

(see Equation VII.2.1) which is what we need. $\dashv_{\text{VII.2.2}}$

Now, let s be any term of $L^{\sigma\flat}$ which depends in $\mathbf{A}^{\sigma\flat}$ on two of its variables. Without loss of generality, we may take s to be equal to $t^j_{i_1 \cdots i_n}$, since identification of variables can never increase essential arity. Let s depend on $v^r_c, v^{r'}_{c'}$; then the term

$$d_{i_0}\left(y_1, \ldots, y_{j-1}, t\left(d_{i_1}\left(x_1^1, \ldots, x_1^{K_{i_1}}\right), \ldots, d_{i_n}\left(x_n^1, \ldots, x_n^{K_{i_n}}\right)\right), \ldots, y_{K_{i_0}}\right)$$

depends on all the y variables and $x^r_c, x^{r'}_{c'}$. $\square$

Lemma VII.3. *Let $t(v_1, v_2, \ldots, v_n)$ be an $L^{\sigma\flat}$-term.*

(1) If the formula $t(v_1, v_1, v_3, \ldots, v_n) = v_1$ is well-formed and holds universally in $\mathbf{A}^{\sigma\flat}$, then t is essentially unary in $\mathbf{A}^{\sigma\flat}$.

(2) *If for some terms* $s_k(v_1, v_2, \ldots)$, *the formula*

$$t(s_1(\vec{v}), s_2(\vec{v}), \ldots, s_n(\vec{v})) = v_1$$

is well-formed and holds universally in $\mathbf{A}^{\sigma\flat}$ *(in which case we call* t *right-invertible) then* t *is essentially unary in* $\mathbf{A}^{\sigma\flat}$.

Proof. (1) For any $y_3, y_3', \ldots, y_n, y_n'$ in the appropriate sorts, the ranges of the polynomials

$$t(v_1, v_2, \vec{y}), \qquad t(v_1, v_2, \vec{y'})$$

are not disjoint. Since $\mathbf{A}^{\sigma\flat}$ is strongly-abelian, all such polynomials must in fact be equal. Let t be a specialization of $s^j_{i_1 i_2 \cdots}$ for some term $s(x_1, x_2, \ldots)$ in L. Since v_1, v_2 have the same sort as t, we may as well assume that v_1 represents the j coordinate of x_1, and similarly for v_2. The operation

$$\begin{gathered} g(y_1, y_2, \ldots, j_{j-1}, x_1, x_2, \ldots, y_{j+1}, \ldots, y_{K_{i_0}}) \\ = \\ d_{i_0}(y_1, y_2, \ldots, y_{j-1}, s(x_1, x_2, \ldots), y_{j+1}, \ldots, y_{K_{i_0}}) \end{gathered}$$

then depends only on the variables shown (i.e. not on $x_3, \ldots$) as a function on

$$\underbrace{C_{i_0} \times \cdots \times C_{i_0}}_{j-1} \times C_{i_0} \times C_{i_0} \times C_{i_3} \times \cdots \times C_{i_\ell} \times \underbrace{C_{i_0} \times \cdots \times C_{i_0}}_{n-j} \to C_{i_0}$$

and is idempotent on the variables in sort $\langle i_0, j \rangle$. Hence $\mathbf{A}$ has a σ-sorted term

$$g(y_1, \ldots, y_{j-1}, x_1, x_2, y_{j+1}, \ldots, y_{K_{i_0}})$$

which is an idempotent operation on C_{i_0} and depends on all the y_k. By maximality this operation cannot depend on both x_1 and x_2, implying that t did not depend on both v_1 and v_2 to begin with.

(2) Let $v_1^1, \ldots, v_1^n$ be variables of the first input sort of s. By part (1), the term

$$t(s_1(v_1^1, v_2, \ldots), s_2(v_1^1, v_2, \ldots), \ldots, s_{n-1}(v_1^1, v_2, \ldots), s_n(v_1^n, v_2, \ldots))$$

depends on none of $v_2, \ldots, v_n$ and on only one of v_1^1, v_1^n. Proceeding inductively, we see that

$$\hat{t}(v_1^1, v_1^2, \ldots, v_1^n, v_2, \ldots) = t(s_1(v_1^1, v_2, \ldots), s_2(v_1^2, v_2, \ldots), \ldots, s_n(v_1^n, v_2, \ldots))$$

depends on just one variable, say v_1^1, and in fact

$$\hat{t}(v_1^1, v_1^2, \ldots, v_1^n, v_2, \ldots) = v_1^1.$$

We claim that t depends only on its first variable. To see this, let $a_1, a_2, a_2', \ldots, a_n, a_n', \vec{b}$ be any elements of the appropriate sorts. Define

$$\begin{aligned} u &= t(a_1, a_2, \ldots, a_n) \\ u' &= t(a_1, a_2', \ldots, a_n') \\ q_2 &= s_2(a_2, \vec{b}) \\ q_2' &= s_2(a_2', \vec{b}) \\ &\vdots \\ q_n' &= s_n(a_n', \vec{b}) \end{aligned}$$

Then since the ranges of $t(v_1, a_2, \ldots, a_n)$ and $t(v_1, q_2, \ldots, q_n)$ both contain u, these two polynomials must be equal; likewise the polynomials $t(v_1, a_1', \ldots, a_n')$ and $t(v_1, q_2', \ldots, q_n')$. But

$$\begin{gathered} u = t(s_1(u, \vec{b}), q_2, \ldots, q_n) = t(s_1(u, \vec{b}), q_2', \ldots, q_n') \\ \Downarrow \\ t(v_1, q_2, \ldots, q_n) = t(v_1, q_2', \ldots, q_n') \end{gathered}$$

which shows that

$$\begin{aligned} t(v_1, a_2, \ldots, a_n) &= t(v_1, q_2, \ldots, q_n) \\ &= t(v_1, q_2', \ldots, q_n') \\ &= t(v_1, a_2', \ldots, a_n') \end{aligned}$$

Since a_k, a_k' were arbitrary, we are done.

$\square$

Lemma VII.4. *If $\mathbf{A}^{\sigma\flat}$ is not essentially unary, then there is an $L^{\sigma\flat}$-term depending essentially in $\mathbf{A}^{\sigma\flat}$ on at least two variables and not left-invertible at either.*

Proof. We show how to take a term depending essentially on v_1, v_2 and invertible at v_1, and produce a new term depending essentially on v_2 and at another variable v_0 (possibly of a different sort than v_1) and not invertible at v_0. We will then show that if we started with a term which was not left-invertible at v_2, then the new term we construct still has this property.

Assume that $t(v_1, v_2, \ldots, v_n)$ depends essentially on v_1 and v_2, and that

$$s(t(v_1, \ldots, v_n), v_{n+1}, \ldots) = v_1 \tag{VII.4.1}$$

The same logic used in part (1) of Lemma VII.3 guarantees that s cannot depend on any variable except the first, so we will write $s(x)$ as if it were a unary term.

Let

$$\hat{t}(v_0, v_2, \ldots, v_n) = t(s(v_0), v_2, \ldots, v_n)$$

Since s maps the output sort of t onto the sort of v_1 in $\mathbf{A}^{\sigma\flat}$, this new term $\hat{t}$ must depend essentially on v_0 and v_2.

Claim VII.4.2. $\hat{t}$ is not left-invertible at v_0.

Suppose elsewise: let

$$r(\hat{t}(v_0, v_2, \ldots, v_n)) = v_0$$

Define another term

$$q(v_0, v_2, v_2', \vec{w}) = \hat{t}(\hat{t}(v_0, v_2, \vec{w}), v_2', \vec{w})$$

(where $\vec{w} = v_3, \ldots, v_n$). Then on the one hand

$$\begin{aligned} \hat{t}(v_0, v_2, \vec{w}) &= r(\hat{t}(\hat{t}(v_0, v_2, \vec{w}), v_2', \vec{w})) \\ &= r(q(v_0, v_2, v_2', \vec{w})) \end{aligned}$$

so q must depend essentially on v_2. But on the other hand

$$\begin{aligned} q(v_0, v_2, v_2', \vec{w}) &= \hat{t}(\hat{t}(v_0, v_2, \vec{w}), v_2', \vec{w}) \\ &= t(s(\hat{t}(v_0, v_2, \vec{w})), v_2', \vec{w}) \\ &= t(s(t(s(v_0), v_2, \vec{w})), v_2', \vec{w}) \\ &= t(s(v_0), v_2', \vec{w}) \end{aligned}$$

which does *not* depend on v_2. $\dashv_{\text{VII.4.2}}$

Lastly, we must show that if $\hat{t}$ were left-invertible at v_2 then t would already have been. This is not hard: suppose

$$v_2 = r_2(\hat{t}(v_0, v_2, \ldots, v_n)) = r(t(s(v_0), v_2, \ldots, v_n))$$

Again using the logic of part (1) of lemma VII.3, the term

$$r(t(s(v_0), v_2, \ldots, v_n))$$

can only depend on v_2; since by Equation (VII.4.1), $v_1 \in \operatorname{ran}(s)$ (considered as elements of the free algebra $\mathbf{F}_{\mathcal{V}(\mathbf{A}^{\sigma\flat})}(v_0, v_1, v_2, \ldots)$), we must have that r_2 inverts t as well. $\square$

Construction VII.5. Let X be any sorted family of generators for a free algebra $\mathbf{F} = \mathbf{F}(X)$ in $\mathcal{V}(\mathbf{A}^{\sigma\flat})$. Let f_0 be an arbitrary fixed element of F, and let $\mathbf{F}' = \mathbf{F}(X \cup \{z\})$, where z is a new free generator of the same sort as f_0.

Generate a congruence $\theta \in \mathrm{Con}(\mathbf{F}')$ from all pairs

$$\langle t(f_0, \vec{u}), t(z, \vec{u})\rangle$$

such that $\vec{u} \in F$ and $t(v_0, \vec{v})$ is not left-invertible at v_0. (Observe that if a term $g \in F$ occurs as the second member $t(z, \vec{u})$ of such a pair, by freeness we get that t does not depend on its first variable, so that the pair is in fact trivial.)

Lemma VII.6. *Let $\mathbf{F}, \mathbf{F}'$, and θ be as in Construction VII.5. If $a \in F$ and $a \equiv_\theta b$, then either $a = b$ or $\langle a, b\rangle$ is a generating pair.*

Proof. Suppose we have basic nontrivial θ-links a—c—b, where

$$\langle a, c\rangle = \langle t_1(f_0, \vec{u}_1), t_1(z, \vec{u}_1)\rangle$$

Case 1:

$$\langle c, b\rangle = \langle p_2(t_2(f_0, \vec{u}_2)), p_2(t_2(z, \vec{u}_2))\rangle$$

where $p_2(v_0) = g_2(v_0, z, \vec{w}_2) \in \mathsf{Pol}_1(\mathbf{A}^{\sigma\flat})$ for some terms $g, \vec{w} \in F$.

We have

$$c = t_1(z, \vec{u}_1) = g_2(t_2(f_0, \vec{u}_2), z, \vec{w}_2)$$

and since z is a free generator, we may substitute any term for z in the above equation. In particular,

$$\begin{aligned} a = t_1(f_0, \vec{u}_1) &= g_2(t_2(f_0, \vec{u}_2), f_0, \vec{w}_2) \\ b &= g_2(t_2(z, \vec{u}_2), z, \vec{w}_2) \end{aligned} \tag{VII.6.1}$$

We will be done with Case 1 if we can establish

Claim VII.6.2. $g_2(t_2(v_0, \vec{u}_2), v_0, \vec{w}_2)$ is not left-invertible at v_0.

Suppose the contrary, say

$$r(g_2(t_2(v_0, \vec{u}_2), v_0, \vec{w}_2)) = v_0 \tag{VII.6.3}$$

By Lemma VII.3, the term

$$r(g_2(t_2(v_0, \vec{u}_2), v_1, \vec{w}_2))$$

must depend only on v_0 or v_1, and because of Equation (VII.6.3) must project to the active variable. Moreover, it cannot be v_0, since then this would be a left-inversion of $t_2(v_0, \vec{u}_2)$. But if v_1 were the active variable, we would have

$$v_1 = r(g_2(t_2(v_0, \vec{u}_2), v_1, \vec{w}_2)) = r(g_2(t_2(f_0, \vec{u}_2), v_1, \vec{w}_2)) = r(t_1(v_1, \vec{u}_1))$$

contradicting our assumption that $t_1(v_1, \vec{u}_1)$ was not invertible. $\dashv_{\text{VII.6.2}}$

Now Equation (VII.6.1) shows that $\langle a, b\rangle$ is a generating pair.

Case 2: As before,

$$\langle a, c\rangle = \langle t_1(f_0, \vec{u}_1), t_1(z, \vec{u}_1)\rangle$$

but now

$$\langle c, b\rangle = \langle p_2(t_2(z, \vec{u})), p_2(t_2(f_0, \vec{u}_2))\rangle$$

with p_2 a polynomial as before. Since

$$c = t_1(z, \vec{u}_1) = g_2(t_2(z, \vec{u}_2), z, \vec{w}_2)$$

and z is a free generator, the same equation holds under any substitution for z:

$$\begin{aligned} a = t_1(f_0, \vec{u}_1) &= g_2(t_2(f_0, \vec{u}_2), f_0, \vec{w}_2) \\ b &= g_2(t_2(f_0, \vec{u}_2), z, \vec{w}_2) \end{aligned}$$

As before, the following claim suffices:

Claim VII.6.4. $g_2(t_2(f_0, \vec{u}_2), v_0, \vec{w}_2)$ is not left-invertible at v_0.

If it were, so

$$r(g_2(t_2(f_0, \vec{u}_2), v_0, \vec{w}_2)) = v_0$$

then the range of this polynomial contains the whole sort of f_0. In particular,

$$r(c) \in \operatorname{ran}\left(r(g_2(t_2(z, \vec{u}_2), \bullet, \vec{w}_2))\right) \cap \operatorname{ran}\left(r(g_2(t_2(f_0, \vec{u}_2), \bullet, \vec{w}_2))\right)$$

By strong abelianness, the two polynomials in the above equation should be equal, contradicting our original assumptions. $\dashv_{\text{VII.6.4}}$

□

Proposition VII.7. *Let* $\mathbf{F}, \mathbf{F}'$, *and* θ *be as in Construction VII.5. Then* z *is isolated* (mod θ).

Proof. Let $\{z, x\} = \{p(t(f_0, \vec{u})), p(t(z, \vec{u}))\}$ be a basic θ-pair, where $p(v_0) = g(v_0, z, \vec{w})$ as in the previous lemma.

First suppose

$$z = p(t(z, \vec{u})) = g(t(z, \vec{u}), z, \vec{w})$$

Then by Lemma VII.3, $(g(t(v_0, \vec{u}), v_1, \vec{w})$ depends only on one variable, either v_0 or v_1. Moreover, v_0 is not a possibility, since then t would be left-invertible. We conclude that $g(t(v_0, \vec{u}), v_1, \vec{w}) = v_1$ throughout HSP $\left(\mathbf{A}^{\sigma\flat}\right)$.

Next suppose

$$z = p(t(f_0, \vec{u})) = g(t(f_0, \vec{u}), z, \vec{w})$$

Then $g(t(v_0, \vec{u}), v_1, \vec{w})$ is right-invertible; invoking Lemma VII.3 again, this term is essentially unary, and since $f_0 \in F$ and z is not, the dependency must be on v_1; hence

$$g(t(v_0, \vec{u}), v_1, \vec{w}) = v_1$$

is valid in $\mathrm{HSP}\left(\mathbf{A}^{\sigma\flat}\right)$.

In either case, we conclude that $z = x$. □

The content of the previous two lemmas is that, for $\mathbf{F}, \mathbf{F}', f_0$, and θ defined in this way, and for $\mathbf{C} = \mathbf{F}'/\theta$, we have that $\mathbf{F}$ is an isomorphic substructure of $\mathbf{C}$, and f_0 and z are indistinguishable by the action of non-left-invertible terms $t(\bullet, \vec{u})$ taken from F.

Recall that since $\mathbf{A}^{\sigma\flat}$ is strongly-abelian, there is an upper bound on the essential arity of terms over this algebra. (For example, $|A| \cdot \max_i K_i$ would work.) Let T be a finite set of $L^{\sigma\flat}$ terms such that every term operation of $\mathbf{A}^{\sigma\flat}$ is given (up to renaming of variables) by one of the terms in T.

For each sort $\langle i, j\rangle$, let $N_{\langle i,j\rangle} \subset T$ be the set of all terms $t(v_0, v_1, \dots)$ such that v_0 has sort $\langle i, j\rangle$ and t is not left-invertible at v_0. Then the relations

$$a \propto_{\langle i,j\rangle} b \iff \bigwedge_{t \in N_{\langle i,j\rangle}} \forall \vec{u}\, t(a, \vec{u}) = t(b, \vec{u})$$

together comprise a definable equivalence relation on any $\mathbf{M} \in \mathrm{HSP}\left(\mathbf{A}^{\sigma\flat}\right)$. We will usually write $a \propto b$ instead of $a \propto_{\langle i,j\rangle} b$.

It is clear from the definition that $\mathbf{a} \propto \mathbf{b}$ in a product $\prod_{x \in X} \mathbf{B}_x$ if and only if $a^x \propto b^x$ in each stalk.

Proposition VII.8. *If $s(v_0, v_1, \dots, v_n)$ is a right-invertible term depending only on v_0, then for any $\mathbf{M} \in \mathrm{HSP}\left(\mathbf{A}^{\sigma\flat}\right)$, any $a \propto b \in M$, and any $x_1, \dots, x_n \in M$ of the appropriate sorts, $s(a, x_1, \dots x_n) \propto s(b, x_1, \dots x_n)$.*

Proof. Say the sort of v_0 is $\langle i, j\rangle$. Let

$$s(t_0(y, \vec{z}), v_1, \dots, v_n) = y$$

and let $t(v_0, \dots, v_\ell) \in N_{\langle i,j\rangle}$. It suffices to show that

$$t(s(v_0, \dots, v_n), v_{n+1}, \dots, v_{n+\ell}) \in N_{\langle i,j\rangle}$$

too.

Suppose otherwise: then for some essentially unary term $r(v_0, \dots)$ we have

$$r(t(s(v_0, \dots, v_n), v_{n+1}, \dots, v_{n+\ell})) = v_0$$

Then

$$\begin{aligned} t_0(y, \vec{z}) &= r(t(s(t_0(y, \vec{z}), \dots, v_n), v_{n+1}, \dots, v_{n+\ell})) \\ &= r(t(y, v_{n+1}, \dots, v_{n+\ell})) \\ y &= s(t_0(y, \vec{z}), v_1, \dots, v_n) \\ &= s(r(t(y, v_{n+1}, \dots, v_{n+\ell})), v_1, \dots, v_n) \end{aligned}$$

contradicting our assumption that t was not left-invertible at its first variable. □

VII–B. Proof of Theorem D

We are now ready to prove Theorem VII.1.

Proof. Let $\mathbf{A}$ be a finite algebra with strongly-solvable radical σ such that every strongly-solvable congruence in $\mathrm{HSP}(\mathbf{A})$ is strongly-abelian, and suppose that $\mathbf{A}^{\sigma\flat}$ is not essentially unary. By Lemma VII.4, we may fix a term $q(v_1, \dots, v_\ell)$ depending essentially on v_1, v_2 but not left-invertible at either. Let X be a sorted collection of free generators: one $x_{\langle i,j\rangle}$ for each sort $\langle i, j\rangle$, as well as two generators a_0, a_1 of the sort of v_1 and two more b_0, b_1 of the sort of v_2. Let

$$v_1 * v_2 = q(v_1, v_2, x_{\langle i_3,j_3\rangle}, \dots, x_{\langle i_\ell,j_\ell\rangle}) \in \mathsf{Pol}_2(\mathbf{F}(X))$$

and define elements

$$\begin{aligned} 0 &= a_0 * b_0 \\ 1 &= a_0 * b_1 \\ 2 &= a_1 * b_0 \\ 3 &= a_1 * b_1 \end{aligned}$$

(These elements are all distinct since $*$ depends on both variables.) Let $\langle i_0, j_0\rangle$ be the type of these four elements, and let $\mathbf{C} = \mathbf{F}'/\theta$, where $\mathbf{F}'$ and θ are built according to Construction VII.5, with 0 playing the role of f_0. As we remarked before, $\mathbf{F} \leq \mathbf{C}$.

We first observe that, by construction, for any $t(v_0, \dots, v_n) \in N_{\langle i_0,j_0\rangle}$ and any $\vec{u} \in F$,

$$\mathbf{C} \models t(0, \vec{u}) = t(z, \vec{u})$$

Since $\mathbf{C}$ is strongly-abelian, it follows that the polynomials $t(0, \bullet)$ and $t(z, \bullet)$ are equal: that is,

$$\mathbf{C} \models 0 \propto z.$$

Claim VII.1.1. $\{0, 1, 2, 3\}$ are pairwise $\propto$-inequivalent.

We will show that $0 \not\propto 1$; the remaining cases are similar.

Suppose for the sake of contradiction that $0 \propto 1$. Observe that $0 \propto 1$ in $\mathbf{F}$ also.

Subclaim VII.1.1a. Under the hypothesis that $0 \propto 1$, 3 is isolated modulo $\beta = \mathrm{Cg}_{\mathbf{F}}(\langle 0, 1\rangle)$.

To see this, let $3 \in \{g(0, \vec{u}), g(1, \vec{u})\}$ for some term g. Then we have

$$3 = a_1 * b_1 = g(a_0 * b, \vec{u})$$

for b either b_0 or b_1; since a_0 appears on the right but not the right and $\mathbf{F}$ is free,

$$g(a_0 * b, \vec{u}) = g(a_1 * b, \vec{u})$$

Thus the polynomial is not injective on $\langle i_0, j_0\rangle$, so $g(v_0, \vec{u})$ cannot be left-invertible, and hence belongs to $N_{\langle i_0, j_0\rangle}$.

Our assumption that $0 \propto 1$ now forces $g(0, \vec{u})$ to be equal to $g(1, \vec{u})$. $\dashv_{\text{VII.1.1a}}$

In particular, $2 \not\equiv_\beta 3$. But then

$$a_0 * b_0 \equiv_\beta a_0 * b_1$$
$$\text{but}$$
$$a_1 * b_0 \not\equiv_\beta a_1 * b_1$$

so β is not abelian. This is a contradiction; the remaining five cases are proved analogously. $\dashv_{\text{VII.1.1}}$

Our plan is to semantically interpret the class of bipartite graphs without isolated vertices into $\mathrm{HSP}\left(\mathbf{A}^{\sigma\flat}\right)$. (It is well known that the theory of bipartite graphs is computably inseparable from the set of sentences false in some finite bipartite graph.) Our strategy will be to define an algebra $\mathbf{D}(\mathbb{G})$ for each graph $\mathbb{G}$, and then to show that certain relations are uniformly first-order definable in these algebras. (Here "uniformly" means that the respective relations are defined via the same first-order formulas for all $\mathbf{D}(\mathbb{G})$; the subsets defined by these formulas in algebras in $\mathrm{HSP}\left(\mathbf{A}^{\sigma\flat}\right)$ but not of the form $\mathbf{D}(\mathbb{G})$ may be quite strange and bear no resemblance to the relations we intend.)

For us, a bipartite graph will be a two-sorted structure $\mathbb{G} = \left\langle R^{\mathbb{G}}, B^{\mathbb{G}} \,;\, E^{\mathbb{G}}\right\rangle$, where E has type signature $\langle R, B\rangle$.

Construction VII.9. Let $\mathbb{G}$ be a bipartite graph. We define a subpower $\mathbf{D} = \mathbf{D}(\mathbb{G}) \leq \mathbf{C}^{\Gamma}$ as follows: the index set $\Gamma = R^{\mathbb{G}} \sqcup B^{\mathbb{G}} \sqcup \{\clubsuit, \spadesuit\}$, and $\mathbf{D}$ is generated by all points

$$\begin{aligned}
\iota_x &= x_{|\Gamma} && (x \in X)\\
\chi_v &= a_{1|v} \oplus a_{0|\mathrm{else}} && (v \in R^{\mathbb{G}})\\
\chi_v &= b_{1|v} \oplus b_{0|\mathrm{else}} && (v \in B^{\mathbb{G}})\\
\chi_{e,\clubsuit} &= 2_{|v} \oplus 1_{|w} \oplus z_{|\clubsuit} \oplus 0_{|\mathrm{else}} && (e = \langle v, w\rangle \in E^{\mathbb{G}})\\
\chi_{e,\spadesuit} &= 2_{|v} \oplus 1_{|w} \oplus z_{|\spadesuit} \oplus 0_{|\mathrm{else}} && (e = \langle v, w\rangle \in E^{\mathbb{G}})
\end{aligned}$$

We let

$$\chi_R = \{\chi_v : v \in R^{\mathbb{G}}\}$$
$$\chi_B = \{\chi_v : v \in G^{\mathbb{G}}\}$$
$$\chi_E = \{\chi_{e,\clubsuit}, \chi_{e,\spadesuit} : e \in E^{\mathbb{G}}\}$$

By abuse of notation, X will still denote the set of diagonal generators ι_x. We will suppose that we have constant symbols for all the ι_x, so that X (and hence F, the subalgebra generated by X) is a uniformly definable subset of D.

Note that $\mathbf{D}$ is not quite a diagonal subpower; it contains all diagonal elements from $\mathbf{F}$, but none of those from $\mathbf{C} \setminus \mathbf{F}$.

Claim VII.1.2. If for some term t and elements $\vec{\mathbf{x}}$ of $\mathbf{D}$, $t(\mathbf{x}_1, \ldots, \mathbf{x}_n)$ is equal to one of the non-diagonal generators, then t is right-invertible (and hence essentially unary).

Suppose first that $t(\mathbf{x}_1, \ldots, \mathbf{x}_n) = \chi_v \in \chi_R$. Then

$$a_1 = \chi_v^v = t(x_1^v, \ldots, x_n^v)$$

and all the elements in this equality belong to F. Since $\mathbf{F}$ is free, this is precisely the statement that t is right-invertible.

The case where v is a blue vertex is the same.

Next let $t(\mathbf{x}_1, \ldots, \mathbf{x}_n) = \chi_{e,\clubsuit}$. Then

$$z = (\chi_{e,\clubsuit})^{\clubsuit} = t(x_1^{\clubsuit}, \ldots, x_n^{\clubsuit})$$

so that in $\mathbf{F}'$, $z \equiv_\theta t(y_1, \ldots, y_n)$ for $y_k/\theta = x_k^{\clubsuit}$. By Proposition VII.7,

$$t(y_1, \ldots, y_k) = z$$

once again showing that t is right-invertible. $\dashv_{\text{VII.1.2}}$

The set $\mathsf{NRINV} \subset D$ of all x such that

x is neither diagonal nor in the image of any term which is not right-invertible.

is uniformly first-order, and we have just shown that every off-diagonal generator lies in this set. While it would be nice if this were actually the set of off-diagonal generators, this might be too much to ask.

To get around this, define $x \le y$ in $\mathbf{D}$ if for some essentially unary term $t(v_0, \ldots)$ we have $x = t^{\mathbf{D}}(y, \ldots)$. Then $\le$ is a definable preorder, and its associated partial order $\sim$ is of course definable too, as is the property of being in a maximal $\sim$-equivalence class.

Claim VII.1.3. The map $\chi \mapsto \chi/\sim$ is a bijection of off-diagonal generators to $\le$-maximal $\sim$-classes containing a member of NRINV.

To prove this, we must first show that no two distinct off-diagonal generators are $\leq$-related. This is done by exhaustive case analysis; none of the cases are hard, but there are a lot of them. We show two, and leave the rest to the skeptic.

For our first model case, suppose v is a red vertex and $\chi_v \leq \chi_{e,\clubsuit}$ for some e. Then for some essentially unary term $t(v_0, \ldots)$,

$$\chi_v = t(\chi_{e,\clubsuit})$$
$$a_0 = (\chi_v)^{\spadesuit} = t\left((\chi_{e,\clubsuit})^{\spadesuit}, \ldots\right) = t(0, \ldots) = t(a_0 * b_0, \ldots)$$

Since a_0, b_0 were free generators, this would imply that the operation $v_0 * v_1$ is left-invertible at v_0, a contradiction.

Next suppose $\chi_{e,\clubsuit} \leq \chi_v$. Then

$$\chi_{e,\clubsuit} = t(\chi_v, \ldots)$$
$$z = (\chi_{e,\clubsuit})^{\clubsuit} = t\left((\chi_v)^{\clubsuit}, \ldots\right) \in F$$

a contradiction. The rest of the cases are handled similarly.

So we have that if we have generators $\mathbf{x}_1 \leq \mathbf{x}_2$ then $\mathbf{x}_1 = \mathbf{x}_2$. Now: suppose that $\mathbf{y} \in \mathsf{NRINV}$. We have $\mathbf{y} = t(\mathbf{x}_1, \ldots, \mathbf{x}_n)$ for some term t and some generators $\mathbf{x}_k$. But by assumption, t is right-invertible, hence depends only on one variable (say the first). In other words $\mathbf{y} \leq \mathbf{x}_1$. Hence every maximal $\sim$-class containing a member of NRINV contains a generator.

Lastly, if $\mathbf{x}_0$ is an off-diagonal generator and $\mathbf{x}_0\mathbf{y} \in \mathsf{NRINV}$, then $\mathbf{x}_0 \leq \mathbf{y} \leq \mathbf{x}_1$ for some generator $\mathbf{x}_1$. By the previous part, $\mathbf{x}_0 = \mathbf{x}_1$. This shows that the $\sim$-class of every off-diagonal generator is maximal. $\dashv_{\text{VII.1.3}}$

Let GEN be the set of all elements of D $\sim$-equivalent to an off-diagonal generator. As we have just seen, this set is uniformly definable: $\mathbf{y} \in \mathsf{GEN}$ if and only if

$\mathbf{y} \in \mathsf{NRINV}$ and for all $\mathbf{y}' \in \mathsf{NRINV}$, $\mathbf{y} \leq \mathbf{y}' \to \mathbf{y}' \leq \mathbf{y}$.

We want to be able to distinguish between edge-type and vertex-type generators. To do this, first observe that for any edge e, $\chi_{e,\clubsuit} \propto \chi_{e,\spadesuit}$ since the relation holds in every factor. This prompts us to set $\mathsf{EDGEGEN}$ to be the subset of GEN consisting of all $\mathbf{x}$ such that

There exist $\mathbf{x}', \mathbf{y} \in \mathsf{GEN}$ with $\mathbf{x} \sim \mathbf{x}'$, $\mathbf{x} \not\propto \mathbf{y}$, and $\mathbf{x}' \propto \mathbf{y}$.

This set is clearly definable.

Claim VII.1.4. For $\mathbf{y} \in \mathsf{GEN}$, $\mathbf{y} \in \mathsf{EDGEGEN}$ if and only if the (unique) generator in $\mathbf{y}/\sim$ has edge type.

By construction, each $\chi_{e,\clubsuit}$ and each $\chi_{e,\spadesuit}$ belong to $\mathsf{EDGEGEN}$. Also, $\mathsf{EDGEGEN}$ is clearly a union of $\sim$-classes.

Hence it suffices to show that $\chi_v \notin \mathsf{EDGEGEN}$ for any vertex v. Suppose this were false: then we would have elements $\mathbf{x} \sim \chi_v$ and $\mathbf{y} \not\sim \chi_v$ with $\mathbf{x} \propto \mathbf{y}$. Let γ be the generator $\sim$-equivalent to $\mathbf{y}$.

Since $\mathbf{x} \sim \chi_v$, they are connected by essentially unary terms

$$\mathbf{x} = f_1(\chi_v) \qquad \chi_v = f_2(\mathbf{x})$$

and likewise

$$\mathbf{y} = g_1(\gamma) \qquad \gamma = g_2(\mathbf{y})$$

Since all four of these elements are in GEN, the terms f_k, g_k must in fact be right-invertible. By Proposition VII.8,

$$\gamma = g_2(\mathbf{y}) \propto g_2(\mathbf{x}) = g_2 \circ f_1(\chi_v)$$

Note that $g_2 \circ f_1$ is right-invertible.

Case 1: $\gamma = \chi_w$ for some $w \neq v$.

Without loss of generality, w is a red vertex. We have $\chi_v^w = \chi_v^{\clubsuit}$, so

$$a_0 = \gamma^{\clubsuit} \propto g_2 \circ f_1(\chi_v^{\clubsuit}) = g_2 \circ f_1(\chi_v^w) \propto \gamma^w = a_1$$

which is impossible.

Case 2: $\gamma = \chi_{e,\clubsuit}$ for some edge. Then e contains an endpoint $w \neq v$, which we may suppose again to be red.

Since $\chi_v^w = \chi_v^{\clubsuit}$,

$$2 = \gamma^w \propto g_2 \circ f_1(\chi_v^w) = g_2 \circ f_1(\chi_v^{\clubsuit}) \propto \gamma^{\clubsuit} = z$$

But this is likewise impossible. $\dashv_{\text{VII.1.4}}$

With this in hand, we know that the set $\mathsf{VERTEXGEN}$ of all $\mathbf{x} \in \mathsf{GEN}$ which are not in $\mathsf{EDGEGEN}$ is (uniformly first-order) definable. This set is, of course, better known as the set of all $\mathbf{x}$ which are $\sim$-equivalent to one of the χ_v.

Lastly, let $\mathsf{EDGE}(x, y)$ be a formula asserting that

> $x \in \mathsf{VERTEXGEN}$ and $y \in \mathsf{VERTEXGEN}$ and there exist $x' \sim x$, $y' \sim y$ and $w \in \mathsf{EDGEGEN}$ such that $w \propto x' * y'$.

Claim VII.1.5. For $\mathbf{x}, \mathbf{y} \in \mathsf{VERTEXGEN}$, $\mathbf{D} \models \mathsf{EDGE}(\mathbf{x}, \mathbf{y})$ iff there exists an edge $e = \{v, w\}$ such that $\mathbf{x} \sim \chi_v$ and $\mathbf{y} \sim \chi_w$.

($\Leftarrow$): If the red vertex v has an edge to the blue vertex w, then

$$\begin{aligned}
\chi_v * \chi_w &= \left(a_{1|v} \oplus a_{0|\text{else}}\right) * \left(b_{1|w} \oplus b_{0|\text{else}}\right) \\
&= a_1 * b_{0|v} \oplus a_0 * b_{1|w} \oplus a_0 * b_{0|\text{else}} \\
&= 2_{|v} \oplus 1_{|w} \oplus 0_{|\text{else}} \\
&\propto 2_{|v} \oplus 1_{|w} \oplus z_{|\clubsuit} \oplus 0_{|\text{else}} \\
&= \chi_{e,\clubsuit}
\end{aligned}$$

($\Rightarrow$): Assume $\mathsf{EDGE}(\mathbf{x}, \mathbf{y})$. Fix

$$\begin{gathered}
\mathbf{x}' \sim \mathbf{x} \sim \chi_v \\
\mathbf{y}' \sim \mathbf{y} \sim \chi_w \\
\mathbf{x}' * \mathbf{y}' \propto \mathbf{w} \sim \chi_{e,\clubsuit}
\end{gathered}$$

(The proof is the same if $w \sim \chi_{e,\spadesuit}$.)

Since all these points are members of GEN, we may choose right-invertible terms so that

$$\chi_{e,\clubsuit} = f(\mathbf{w}) \qquad \mathbf{x}' = g(\chi_v) \qquad \mathbf{y}' = h(\chi_w)$$

Then f is right-invertible and

$$f(\mathbf{w}) \propto f(\mathbf{x}' * \mathbf{y}') = f(g(\chi_v) * h(\chi_w))$$

We will be done if we can show that $e = \langle v, w \rangle$.

If this were false, we could choose an endpoint $u \in e \setminus \{v, w\}$, which we may suppose is red. Then

$$\chi_v^u = \chi_v^\clubsuit = a_0 \qquad \chi_w^u = \chi_w^\clubsuit = b_0$$

so

$$\begin{aligned}
2 = \chi_{e,\clubsuit}^u = f(w^u) &\propto f(g(\chi_v^u) * h(\chi_w^u)) \\
&= f(g(\chi_v^\clubsuit) * h(\chi_w^\clubsuit)) \propto f(w^\clubsuit) = \chi_{e,\clubsuit}^\clubsuit = z
\end{aligned}$$

a contradiction. $\dashv_{\text{VII.1.5}}$

Observe that since $\mathbb{G}$ has no isolated vertices, the subsets $\mathsf{VERTEXRED}$ and $\mathsf{VERTEXBLUE}$ of $\mathsf{VERTEXGEN}$ consisting of those $\mathbf{x}$ which are $\sim$-equivalent to a red (resp. a blue) vertex are definable using the EDGE relation.

The foregoing shows that

$$\langle \mathsf{VERTEXRED}/\sim, \mathsf{VERTEXBLUE}/\sim\, ;\ \mathsf{EDGE} \rangle$$

is isomorphic to our original bipartite graph $\mathbb{G}$; since all the relations in this isomorphism are uniformly definable, we have effected a semantic embedding of bipartite graphs into HSP $\left(\mathbf{A}^{\sigma\flat}\right)$. $\square$

Proof of Theorem D. Theorem VII.1 shows that, if $\mathbf{A}$ has a σ-sorted term operation depending on too many variables, then HSP $\left(\mathbf{A}^{\sigma\flat}\right)$ is hereditarily finitely undecidable. But we have already seen in Lemma VI.12 that HSP $\left(\mathbf{A}^{\sigma\flat}\right)$ semantically embeds into HSP $(\mathbf{A})$. Since semantic interpretability is transitive, we are done. $\square$

CHAPTER VIII

NEXT STEPS

Pawel Idziak is fond of saying that if you want to prove something is undecidable, use local structure; if you want decidability, think global. We have followed the first part of this maxim in this paper, by compressing "bad behavior" down to minimal sets and finding small sets of parameters which ensure that the propagation of the behavior through subpower generation is well-controlled.

However, all the results in this paper are, in a sense, negative; we do not provide any sufficient conditions guaranteeing that $\mathrm{HSP}_{\mathrm{fin}}(\mathbf{A})$ has decidable theory. A very few such conditions exist in the literature; as mentioned, the most comprehensive is that given by Idziak in [Idz97], which not only provides sufficient conditions, but characterizes finitely-decidable congruence-modular varieties up to the (still wide-open) problem of determining for which finite rings $\mathbf{R}$ the variety of $\mathbf{R}$-modules is finitely-decidable. In particular, the construction provides an effective procedure for answering the finite decidability problem for $\mathrm{HSP}_{\mathrm{fin}}(\mathbf{A})$ so long as all congruence covers in this prevariety have the boolean type.

The other broad sufficient condition is provided by Valeriote and McKenzie; it asserts that a strongly-abelian locally-finite variety is either semantically bi-interpretable with a multi-sorted unary variety whose free algebra is chain-preordered by divisibility, or else (by a variant of the construction in Chapter VII) admits a semantic interpretation of graphs. The latter theorem was generalized in [HV91] to remove the hypothesis of local-finiteness; their proof allows the variety to be multi-sorted.

In the case of decidable varieties, [MV89, Chapter 13] shows that an ordinary single-sorted variety which is abelian and which does not decompose as the varietal product of affine and strongly-abelian subvarieties must be undecidable and finitely undecidable. Hart, Starchenko and Valeriote subsequently pushed the model-theoretic method further in [HSV94], which showed that for any variety $\mathcal{V}$, either $\mathcal{V}$ decomposes as the varietal product of an affine subvariety and a strongly-abelian subvariety, or $\mathcal{V}$ is large in the sense of stability theory: it must have continuously many countable models and fail to be superstable.

These two last-mentioned proofs have radically different character from each other; and my immediate project is to understand the details of each, and determine if either can be adapted (or an entirely new method developed) to show the following:

Suppose we consider finite algebras $\mathbf{A}$ such that the solvable radical $\sigma = \mathrm{Rad}(\mathbf{A})$ is comparable to every congruence of $\mathbf{A}$. (This is not as restrictive as it sounds; one can show that in every variety which we might hope to show is finitely-decidable, every algebra is residually in this class.) Let σ_1, σ_2 be respectively the least congruence below $\mathrm{Rad}(\mathbf{A})$ such that the included interval consists of only unary-type (resp. affine-type) covers.

Problem 1. (1) Show that $\mathbf{A}^\sigma$ is the direct product (in the sense of L^σ) of $(\mathbf{A}/\sigma_1)^\sigma$ and $(\mathbf{A}/\sigma_2)^\sigma$.

(2) Show that $\mathrm{HSP}(\mathbf{A}^\sigma)$ is the varietal product of $\mathrm{HSP}((\mathbf{A}/\sigma_1)^\sigma)$ and $\mathrm{HSP}((\mathbf{A}/\sigma_2)^\sigma)$.

If this is done, it should become manageable to establish necessary and sufficient conditions, as in [Idz97], for a computable reduction of $\mathrm{Th}(\mathcal{V})$ to the theory of some modules obtainable in a nice way from the affine part of the variety. (This is accomplished, Pawel would say, by thinking globally and very hard.)

This would do for finitely-generated varieties. The corresponding problems for locally-finite varieties which are not finitely-generated may not be much harder, though there are open questions here which sound like they ought to be easy, and are not:

Problem 2. It is known that every finitely-generated discriminator variety is decidable. Characterize the decidability of locally-finite discriminator varieties.

Lastly, in all cases that I am aware of, the theory of a locally-finite variety is either computable or equivalent to the Halting Problem.

Problem 3. (1) Does there exist a locally-finite variety whose first-order theory is undecidable but properly Turing-below $\emptyset'$?

(2) Does there exist a locally-finite variety whose set of finitely-refuted sentences is undecidable but properly Turing-below $\emptyset'$?

(3) Given an arbitrary c.e. Turing degree D, does there exist a variety $\mathcal{V}$ such that $\mathrm{Th}(\mathcal{V}) \equiv_T D$?

(4) Given an arbitrary Π_1 Turing degree D, does there exist a variety $\mathcal{V}$ such that $\mathrm{Th}_{\mathrm{fin}}(\mathcal{V}) \equiv_T D$?

It is, naturally, of interest to know if the latter two can be answered affirmatively if we require the language of $\mathcal{V}$ to be finite.

BIBLIOGRAPHY

[ARV12] Jiří Adámek, Jiří Rosickỳ, and Enrico M Vitale. "Birkhoff's variety theorem in many sorts". In: *Algebra universalis* 68.1–2 (2012), pp. 39–42.

[BS81] Stanley Burris and Hantamantagouda Sankappanavar. *A course in universal algebra.* Vol. 78. Graduate Texts in Mathematics. Springer-Verlag, New York, 1981.

[EF72] Paul C Eklof and Edward R Fischer. "The elementary theory of abelian groups". In: *Annals of Mathematical Logic* 4.2 (1972), pp. 115–171.

[Ers72] Yuri Ershov. "Elementary Theories of Groups". In: *Doklady Akademii Nauk SSSR* 203 (1972), pp. 1240–1243.

[FM87] Ralph Freese and Ralph McKenzie. *Commutator theory for congruence modular varieties.* Cambridge University Press, 1987.

[HSV94] Bradd Hart, Sergei Starchenko, and Matthew Valeriote. "Vaughts conjecture for varieties". In: *Transactions of the American Mathematical Society* 342.1 (1994), pp. 173–196.

[HV91] Bradd Hart and Matthew Valeriote. "A structure theorem for strongly abelian varieties with few models". In: *The Journal of Symbolic Logic* 56.3 (1991), pp. 832–852.

[HM88] David Hobby and Ralph McKenzie. *The Structure of Finite Algebras.* Vol. 76. Contemporary Mathematics. American Mathematical Society, 1988.

[Hod93] Wilfrid Hodges. *Model theory.* Vol. 42. Cambridge University Press, 1993.

[Idz86] Katarzyna Idziak. "Undecidability of Brouwerian Semilattices". In: *Algebra Universalis* 22 (1986), pp. 298–301.

[II88] Katarzyna Idziak and Pawel Idziak. "Decidability Problem for Finite Heyting Algebras". In: *Journal of Symbolic Logic* 53.3 (1988), pp. 729–735.

[Idz89a] Pawel Idziak. "Varieties with decidable finite algebras I: Linearity". In: *Algebra Universalis* 26 (1989), pp. 234–246.

[Idz89b] Pawel Idziak. "Varieties with decidable finite algebras II: Permutability". In: *Algebra Universalis* 26 (1989), pp. 247–256.

[Idz97] Pawel Idziak. "A characterization of finitely decidable congruence modular varieties". In: *Transactions of the American Mathematical Society* 349.3 (1997), pp. 903–934.

[IMV09] Pawel Idziak, Ralph McKenzie, and Matthew Valeriote. "The Structure of Locally Finite Varieties with Polynomially Many Models". In: *Journal of the American Mathematical Society* 22.1 (2009), pp. 119–165.

[IV01] Pawel Idziak and Matthew Valeriote. "A Property of the Solvable Radical in Finitely Decidable Varieties". In: *Fundamenta Mathematicae* 170 (2001), pp. 69–86.

[Jeo99] Joohee Jeong. “A Decidable Variety that is Finitely Undecidable”. In: *The Journal of Symbolic Logic* 64.2 (1999), pp. 651–677.

[Kea93] Keith Kearnes. “An Order-theoretic Property of the Commutator”. In: *International Journal of Algebra and Computation* 3.4 (1993), pp. 491–533.

[Kea91] Keith A Kearnes. “Residual bounds for varieties of modules”. In: *Algebra Universalis* 28.3 (1991), pp. 448–452.

[KS95] Olga Kharlampovich and Mark Sapir. “Algorithmic problems in varieties”. In: *International Journal of Algebra and Computation* 5 (1995), pp. 379–602.

[Mal65] A.I. Malcev. “On a correspondence between rings and groups”. In: *AMS Translations* 45 (1965). Russian original article in Mathematics: Sbornik 50 (1960), pp. 221–231.

[MMT87] Ralph McKenzie, George McNulty, and Walter Taylor. *Algebras, lattices, varieties*. Vol. I. Wadsworth & Brooks Cole, 1987.

[MS05] Ralph McKenzie and John Snow. “Congruence modular varieties: commutator theory and its uses”. In: *Structural Theory of Automata, Semigroups, and Universal Algebra*. NATO Science Series 207 (2005), pp. 273–329.

[MV89] Ralph McKenzie and Matthew Valeriote. *The Structure of Decidable Locally Finite Varieties*. Progress in Mathematics. Birkhäuser, 1989.

[Ols91] Alexander Yu. Olshanskii. *Geometry of defining relations in groups*. Vol. 70. Mathematics and its applications (Soviet series). Kluwer Academic Publishers, 1991.

[Szm55] Wanda Szmielew. “Elementary properties of Abelian groups”. In: *Fundamenta Mathematicae* 41 (1955), pp. 203–271.

[TMR53] Alfred Tarski, Andrzej Mostowski, and Raphael Robinson. *Undecidable Theories*. North Holland, 1953.

[Val94] Matthew Valeriote. “On Solvable Congruences in Finitely Decidable Varieties”. In: *Mathematical Logic Quarterly* 40 (1994), pp. 398–414.

[VW92] Matthew Valeriote and Ross Willard. “Some Properties of Finitely Decidable Varieties”. In: *International Journal of Algebra and Computation* 2.1 (1992), pp. 89–101.

[Zam76] A.P. Zamyatin. “Varieties of associative rings whose elementary theory is decidable”. In: *Doklady Akademii Nauk SSSR* 229.2 (1976), pp. 276–279.

[Zam78] A.P. Zamyatin. “A non-abelian variety of groups has an undecidable elementary theory”. In: *Algebra i Logika* 17.1 (1978), pp. 20–27.