GEORGE F. MCNULTY

Equational Logic

DRAWINGS BY THE AUTHOR

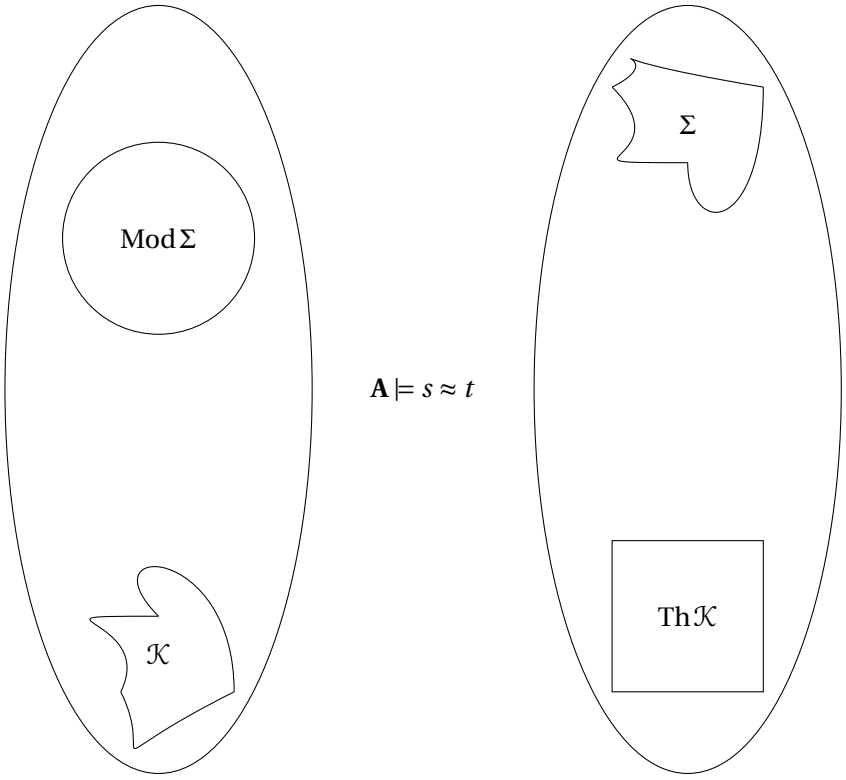

UNIVERSITY OF SOUTH CAROLINA

SPRING 2017

PREFACE

The concepts that can be expressed by means of equations and the kinds of proofs that may be devised using equations are central concerns of equational logic. The concept of a ring is ordinarily presented by saying that a ring is a system $\langle R, +, -, \cdot, 0, 1\rangle$ in which the following equations are true:

$$\begin{array}{lll} x+(y+z)\approx(x+y)+z & x\cdot(y\cdot z)\approx(x\cdot y)\cdot z & x\cdot(y+z)\approx x\cdot y+x\cdot z \\ x+y\approx y+x & x\cdot 1\approx x & (x+y)\cdot z\approx x\cdot z+y\cdot z \\ -x+x\approx 0 & 1\cdot x\approx x & \\ x+0\approx x & & \end{array}$$

A ring is an *algebra*—meaning here a nonempty set endowed with a system of finitary operations. *Equations*, on the other hand, are certain strings of formal symbols. The concept of truth establishes a binary relation between equations and algebras: an equation $s\approx t$ *is true* in an algebra $\mathbf{A}$. This relationship underlies virtually all work in equational logic. By way of this relation each *equational theory*—that is, each set of equations closed under logical consequence—is associated with a *variety of algebras*: the class of all algebras in which each equation of the theory is true. Through this connection syntactical and computational tools developed for dealing with equations can be brought to bear on algebraic questions about varieties. Conversely, algebraic techniques and concepts from the theory of varieties can be employed on the syntactical side.

It turns our that a variety is a class of similar algebras closed with respect to forming homomorphic images, subalgebras, and arbitrary direct products. The classification of algebras into varieties is compatible with most commonly encountered algebraic constructions. It allows us to gather algebras into classes that can be easily comprehended and manipulated. In many cases, it allows us to distinguish algebras that strike our intuitions as genuinely different, as well as to group together algebras which seem to belong together.

This exposition will focus on topics like

1. **Finite Axiomatizability:** Which varieties can be specified by finitely many equations?
2. **Decision Problems:** For which varieties is it possible to have a computer algorithm that determines which equations are true in the variety? Is there a computer algorithm which would determine whether a finite set of equations specifies exactly the class of all groups?
3. **The Lattice of Equational Theories:** Set-inclusion imposes a partial order on the set of equational theories. The structure of this ordered set reflects the comparative strength of the equational theories.

CONTENTS

LESSON 0

EQUATIONAL LOGIC—THE SET UP

Formal systems of mathematical logic are provided with a means of expression and a means of proof. Equational logic is perhaps the simplest example that is still able to comprehend a considerable portion of mathematics that arises in practice. From one perspective, equational logic is a fragment of elementary (or first-order) logic. In this fragment the only formulas are equations between terms—this logic is provided with neither connectives nor quantifiers (apart from implicit universal quantifiers). In comparison with elementary logic, equational logic has a meager means of expression. As a consequence, many of the powerful methods of model theory that have been developed for elementary logic seem to have limited applicability in equational logic. On the other hand, since the truth of equations is preserved under the formation of homomorphic images, subalgebras, and direction products, the methods of algebra can be brought into play.

0.1 THE SYNTAX OF EQUATIONAL LOGIC

A **signature** is a function which gives natural numbers as outputs. The inputs of a signature are called **operation symbols**. The outputs of a signature are called **ranks**. **Constant symbols** are those operation symbols of rank 0. In any particular investigation, we are free to choose a convenient signature. For example, for the equational logic of groups we might choose a signature that provides one two-place operation symbol (to denote the group product), a one-place operation symbol (to denote the formation of inverses), and one constant symbol (to denote the identity element).

Fix some signature. An **algebra** $\mathbf{A} = \langle A, F\rangle$ of the given signature is a system made up of a nonempty set A and a system F of operations of finite rank appropriate for the signature. That is, F is a function whose domain is the same as the domain of the signature and $F(Q)$ is an operation on A of the rank assigned by the signature to operation symbol Q. Ordinarily, we dispense with F and use $Q^{\mathbf{A}}$ to denote $F(Q)$. The set A is called the **universe of discourse** of the algebra $\mathbf{A}$. But we will be less formal and refer to it more simply as the **universe** of $\mathbf{A}$. The operation $Q^{\mathbf{A}}$ are the **basic** or **fundamental** operations of $\mathbf{A}$.

Example. Suppose our operation symbols are $+, \times, 0, 1$ and our algebra is $\mathbf{M} = \langle M_2(\mathbb{R}), F\rangle$ where $M_2(\mathbb{R})$ is the set of 2×2 matrices over the field of real numbers. Then $F(+)$ is

$$+^{\mathbf{M}} : M_2(\mathbb{R}) \times M_2(\mathbb{R}) \longrightarrow M_2(\mathbb{R})$$

where $+^{\mathbf{M}}$ is matrix addition. Likewise with take $F(\times) = \times^{\mathbf{M}}$ to be matrix multiplication, $F(1)$ to be the identity matrix, and $F(0)$ is our zero matrix.

We generally write algebras like $\mathbf{A} = \langle A, +, \cdot, \times, \ldots \rangle$ rather than $\langle A, F \rangle$ where F is appropriately defined.

Returning to the syntactical arrangements, we only other essential part of our syntax is a countably infinite list of distinct symbols for **variables**: $v_0, v_1, v_2, \ldots$. For convenience, we also supply ourselves with the symbol $\approx$ for equality.

Fix a signature. Terms are built up from the variables and the constant symbols with the help of the operation symbols of positive rank. We give a precise definition.

Definition. The set of **terms** is the smallest set T of finite sequences of operation symbols and variables satisfying the following constraints:

- Every variable belongs to T;
- If Q is an operation symbol of rank r and $t_0, \ldots, t_{r-1}$ belong to T, then $Q t_0 t_1 t_{r-1}$ also belongs to T.

More informally, every variable is a term and $Q t_0 \ldots t_{r-1}$ is a term, whenever $t_0, \ldots, t_{r-1}$ are terms and Q is an operation symbol with rank r.

According to this definition, operation symbols are to be followed by the terms they combine. This is sensible, since we have allowed operation symbols to have any finite rank. However, it is at variance with the customary practice of writing terms like this $x \cdot (y + z)$ rather than like this $\cdot x + yz$. This last string of symbols looks odd indeed, but it is nevertheless what is called for by our official definition. There are a couple of virtues that the official definition has. In the first place, we do not have to deal with parentheses—simplifying our syntax. In the second place, while the customary practice works well for two-place operation symbols, there seems to be no natural way to extend it to operation symbols of, let us say, rank seven. This system of notation was promoted by the Polish logician Jan Łukasiewicz in the 1920's and is sometimes called Łukasieicz or Polish notation. More simply it is called prefix notation.

Terms have a key property that we will employ without much reference: A string of symbols that is a term can be parsed in only one way. That is, if Q is an operation symbol of rank r and $t_0, \ldots, t_{r-1}$ and $s_0, \ldots, s_{r-1}$ are terms such that the string $Q t_0 t_1 \ldots t_{r-1}$ is the same as the string $Q s_0 s_1 \ldots s_{r-1}$, then it must be that $t_0 = s_0, \ldots, t_{r-1} = s_{r-1}$. This may seem so obvious as to require no proof. Nevertheless it is the subject of Problem Set 0 and the proof outlined there as some subtlety.

It is often helpful to depict terms as (ordered rooted) trees. Under such a scheme each term as a top node: the first symbol in the term. Beneath the top node, should in be an operation symbol of rank r, there will be r nodes arranged from left to right. These nodes will be the top nodes of trees depicting subterms of the original term. The following display gives the idea.

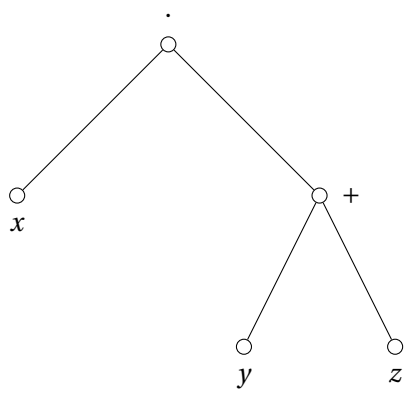

The tree depicting $x \cdot (y + z)$

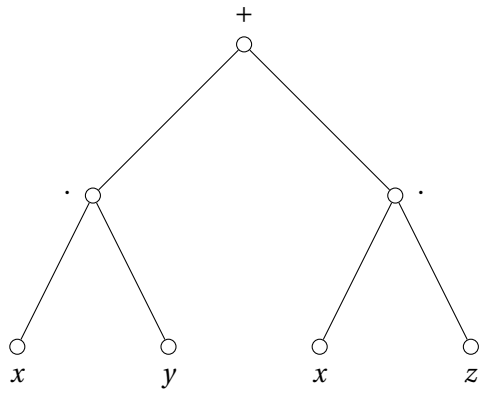

The tree depicting $(x \cdot y) + (x \cdot z)$

An **equation** is an ordered pair of terms. We denote these ordered pairs as $s \approx t$, rather than (s,t). So even though $\approx$ has been made officially part of our syntax, it just denotes ordered pair and could be omitted.

0.2 Problem Set 0

Problem Set About Unique Readability

In the problems below L is the set of operation and relation symbols of same signature and X is a set of variables.

Problem 0.
Define a function λ from the set of finite nonempty sequences of elements of $X \cup L$ into the integers as follows:

$$\lambda(w) = \begin{cases} -1 & \text{if } w \in X, \\ r-1 & \text{if } w \text{ is an operation symbol of rank } r, \\ \sum_{i<n} \lambda(u_i) & \text{if } w = u_0 u_1 \dots u_{n-1} \text{ where } u_i \in X \cup L \text{ and } n > 1. \end{cases}$$

Prove that w is a term if and only if $\lambda(w) = -1$ and $\lambda(v) \geq 0$ for every nonempty proper initial segment v of w.

Problem 1.
Let $w = u_0 u_1 \dots u_{n-1}$, where $u_i \in X \cup L$ for all $i < n$. Prove that if $\lambda(w) = -1$, then there is a unique cyclic variant $\hat{w} = u_i u_{i+1} \dots u_{n-1} u_0 \dots u_{i-1}$ of w that is a term.

Problem 2.
Prove that if w is a term and w' is a proper initial segment of w, then w' is not a term.

Problem 3.
Let $\mathbf{T}$ be the term algebra of L over X. Prove

> If Q and P are operation symbols, and $P^{\mathbf{T}}(p_0, p_1, \dots, p_{n-1}) = Q_1^{\mathbf{T}}(q_0, q_1, \dots, q_{m-1})$, then $P = Q$, $n = m$, and $p_i = q_i$ for all $i < n$.

0.3 The Semantics of Equational Logic

Let $\mathbf{A}$ be an algebra and let t be a term of the same signature. Then $t^{\mathbf{A}}$ will be a certain function from A^ω into A defined as follows:

- $v_i^{\mathbf{A}}(a_0, a_1, \dots) = a_i$ for all $a_0, a_1, \dots \in A$;
- Suppose $Q t_0 \cdots t_{r-1}$ is given. Then

$$(Q t_0 \cdots t_{r-1})^{\mathbf{A}}(a_0, a_1, a_2, \dots) = Q^{\mathbf{A}}\left(t_0^{\mathbf{A}}(a_0, a_1, a_2, \dots), t_1^{\mathbf{A}}(a_0, a_1, a_2, \dots), \dots, t_{r-1}^{\mathbf{A}}(a_0, a_1, a_2, \dots)\right)$$

The functions described above are called the **term functions** of **A**. Of course, even though we have made term functions to have rank ω, really each one depends only on finitely many of its denumerably many inputs. An alternatiive definition of term functions has certain functions of finite rank can be easily constructed, but it is more involved.

Example. Let's look at an example. Consider the familiar algebra (it is the ring of integers)

$$\mathbf{Z} = \langle \mathbb{Z}, +^{\mathbf{Z}}, \cdot^{\mathbf{Z}}, -^{\mathbb{Z}}, 0, 1 \rangle .$$

Then

$$((x_0 + x_1) + x_2)^{\mathbf{Z}}(3,7,5,4,4,4,\ldots)$$

is really

$$\left(x_0^{\mathbf{Z}}(3,7,5,4,4,4,\ldots) + x_1^{\mathbf{Z}}(3,7,5,4,4,4,\ldots)\right) +^{\mathbf{Z}} x_2^{\mathbf{Z}}(3,7,5,4,4,4,\ldots)$$

which is really

$$\left(3 +^{\mathbf{Z}} 7\right) +^{\mathbf{Z}} 5$$

and of course this is 15.

At last, here is the crucial element of our semantical arrangements—what it means for an equation $s \approx t$ to be **true** in the algebra **A**. Just as the operation symbols can be regarded as names for the basic operations of **A**, so the terms cqan be seen as names for the term functions of **A**. We say that $s \approx t$ is true in **A** if and only if s and t name the same term function of **A**—that is, if and only if $s^{\mathbf{A}} = t^{\mathbf{A}}$. Here are alternative ways to express this:

- $\mathbf{A} \models s \approx t$;
- **A** is a model of $s \approx t$;
- $s \approx t$ is true in **A**;
- $\mathbf{A} \models \forall \bar{y}[s \approx t]$, where $\bar{y}$ is any finite string of variables that includes all the variables occurring in the terms s and t.

The last alternative listed above reflects that equational logic is a fragment of elementary logic.

The truth of equations in algebras imposes a two-place relation between the class of all algebras of our fixed signature and the set of all equations of the same signature. $\models$ to denote this relation. Like any two-place relation, $\models$ gives rise to a Galois connection. This Galois connection is crucial to equational logic.

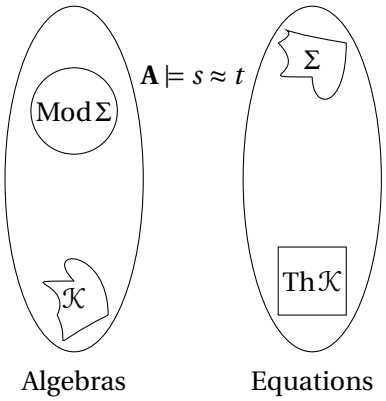

The polarities of the connection are

$$\operatorname{Mod}\Sigma = \{\mathbf{A} \mid A \models s \approx t \text{ for all } s \approx t \in \Sigma\}, \quad \text{where } \Sigma \text{ is any set of equations.}$$

and

$$\operatorname{Th}\mathcal{K} = \{s \approx t \mid \mathbf{A} \models s \approx t \text{ for all } \mathbf{A} \in \mathcal{K}\}, \quad \text{where } \mathcal{K} \text{ is any class of algebras.}$$

$\operatorname{Mod}\Sigma$ is called the **equational class** or the **variety based** on Σ. It is the class of all models of Σ. The set $\operatorname{Th}\mathcal{K}$ is called the **equational theory** of $\mathcal{K}$. It is the class of all equations that are true in every algebra belonging to $\mathcal{K}$.

As is true of every Galois connection, this one provides two closure operators

$$\operatorname{Mod}\operatorname{Th}\mathcal{K} \supseteq \mathcal{K} \text{ and } \operatorname{Th}\operatorname{Mod}\Sigma \supseteq \Sigma.$$

The closed classes on the algebra side are just the varieties of algebras, while the closed sets of t he equation side are exactly to the equational theories. Under the inclusion relation, these closed classes are ordered as complete lattices that are dually isomorhpic to each other. Problem Set 1 provides an overview of these general results about Galois connections.

The first task in the development of equational logic is to provide descriptions of these two closure operations.

0.4 Problem Set 1

Problem Set on Galois Connections

In Problem 4 to Problem 8 below, let A and B be two classes and let R be a binary relation with $R \subseteq A \times B$. For $X \subseteq A$ and $Y \subseteq B$ put

$$X^{\rightarrow} = \{b \mid x \, R \, b \text{ for all } x \in X\}$$
$$Y^{\leftarrow} = \{a \mid a \, R \, y \text{ for all } y \in Y\}$$

Problem 4.
Prove that if $W \subseteq X \subseteq A$, then $X^{\rightarrow} \subseteq W^{\rightarrow}$. (Likewise if $V \subseteq Y \subseteq B$, then $Y^{\leftarrow} \subseteq V^{\leftarrow}$.)

Problem 5.
Prove that if $X \subseteq A$, then $X \subseteq X^{\rightarrow\leftarrow}$. (Likewise if $Y \subseteq B$, then $Y \subseteq Y^{\leftarrow\rightarrow}$.)

Problem 6.
Prove that $X^{\rightarrow\leftarrow\rightarrow} = X^{\rightarrow}$ for all $X \subseteq A$ (and likewise $Y^{\leftarrow\rightarrow\leftarrow} = Y^{\leftarrow}$ for all $Y \subseteq B$).

Problem 7.
Prove that the collection of subclasses of A of the form $Y^{\leftarrow}$ is closed under the formation of arbitrary intersections. (As is the collection of subclasses of B of the form $X^{\rightarrow}$.) We call classes of the form $Y^{\leftarrow}$ and the form $X^{\rightarrow}$ closed.

Problem 8.
Let $A = B = \{q \mid 0 < q < 1 \text{ and } q \text{ is rational}\}$. Let R be the usual ordering on this set. Identify the system of closed sets. How are they ordered with respect to inclusion?

LESSON 1

THE DESCRIPTION OF $\mathrm{Mod}\,\mathrm{Th}\,\mathcal{K}$

The closure operator on the algebra side of the Galois connection established by truth, that is by $\models$, between algebras and equations is given by

$$\mathrm{Mod}\,\mathrm{Th}\,\mathcal{K},$$

for any class $\mathcal{K}$ of algebras, all of the same signature. This operator takes, as input, the class $\mathcal{K}$ of algebras, and returns as output a class, perhaps somewhat larger, of algebras. Obtaining this output, in the way described, requires a detour through our syntactical arrangements with the help of out semantical notion of truth. What we desire here is a description of this closure operator that is entirely algebraic and avoids this detour.

1.1 ALGEBRAIC PRELIMINARIES

Let $\mathbf{A}$ and $\mathbf{B}$ be algebras of the same signature.

Definition. A function $h : A \longrightarrow B$ is a **homomorphism** provided it preserves the operations of the signature; that is, provided for every operation symbol Q and every $a_0, \ldots, a_{r-1} \in A$, where r is the rank of Q, we have

$$h\left(Q^{\mathbf{A}}(a_0, a_1, \ldots, a_{r-1})\right) = Q^{\mathbf{B}}\left(h(a_0), h(a_1), \ldots, h(a_{r-1})\right)$$

A straightforward argument by induction on the complexity of a term t reveals that

$$h\left(t^{\mathbf{A}}(a_0, a_1, \ldots)\right) = t^{\mathbf{B}}\left(h(a_0), h(a_1), \ldots\right)$$

Here is some notation.

$$h : \mathbf{A} \twoheadrightarrow \mathbf{B}$$

denotes a homomorphism from $\mathbf{A}$ onto $\mathbf{B}$. In this case, we say that $\mathbf{B}$ is a **homomorphic image** of $\mathbf{A}$. We use

$$h : \mathbf{A} \rightarrowtail \mathbf{B}$$

to denote a one-to-one homomorphism—these are called **embeddings** . A homomorphism that is both one-to-one and maps A onto B, is called an **isomorphism**. Isomorphisms are invertible and the inverse of an isomorphism is easily shown to be an isomorphism. If there is an isomorphism between $\mathbf{A}$ and $\mathbf{B}$ we say that $\mathbf{A}$ and $\mathbf{B}$ are **isomorphic** and we denote this by $\mathbf{A} \cong \mathbf{B}$.

A homomorpism from **A** into **A** is called an **endomorphism** of **A**. The set of all endomorphisms of **A** is denoted by End**A**. An isomorphism from **A** to **A** is called an **automorphism** of **A**. The set of all automorphisms of **A** is denoted by Aut**A**.

Let $\mathcal{K}$ be a class of algebras of our signature. Then

$$\mathcal{HK} := \{\mathbf{B} : \mathbf{B} \text{ is a homomorphic image of some algebra in } \mathcal{K}\}$$

Fact. Suppose that **B** is a homomorphic image of the algebra **A**. Every equation true in **A** is also true in **B**.

Put another way, the failure of an equation in **B** can be pulled back to become a failure of the same equation in **A**.

Definition. Let **A** be an algebra. The set B is a **subuniverse** of **A** provided

- $B \subseteq A$;
- B is closed under all the operations of **A**.

If B is nonempty, we can make **B** (a **subalgebra** of **A**) by restricting the operations of **A** to the set B.

Let $\mathcal{K}$ be a class of algebras of our signature. Then

$$\mathcal{SK} := \{\mathbf{B} : \mathbf{B} \text{ is isomorphic to a subalgebra of an algebra in } \mathcal{K}\}$$

Fact. Let **A** be a subalgebra of **B**. Then every equation that holds in **B** must also hold in **A**.

Put another way, if an equation fails in **A**, then the variables can be assigned values from A in such a way that the term functions from the left and right sides of the equation will produce different values. Since every element of A is also an element of B and the basic operations of **A** evaluate in the same way as the corresponding basic operation of **B**, then the same assignment will witness the failure of the equation in **B** as well.

Let I be any set and, for each $i \in I$, let $\mathbf{A}_i$ be an algebra of our signature. Then we know $A_i \neq \varnothing$ for all i. We define

$$\prod_{i\in I} A_i := \left\{s | s : I \to \bigcup_{i\in I} A_i \text{ so that } s_i \in A_i \text{ for each } i \in I\right\}.$$

In the expression above $s = \langle s_i : i \in I\rangle$ is an I-tuple.

Definition. The **direct product** $\prod_{i\in I} \mathbf{A}_i$ is the algebra with universe $\prod_{i\in I} A_i$ so that for each operation symbol Q we put

$$Q^{\prod_{i\in I}\mathbf{A}_i}\left(s^0, s^1, \ldots, s^{r-1}\right) = \left\langle Q^{\mathbf{A}_i}(s_i^0, s_i^1, \ldots, s_i^{r-1}) | i \in I\right\rangle$$

where r is the rank of Q and $s^k \in \prod_{i\in I} A_i$ for each $k < r$.

Example. Suppose the rank of our operation Q is 2 and that we have three algebras $\mathbf{A}_0, \mathbf{A}_1$, and $\mathbf{A}_2$. Then

$$Q^{\mathbf{A}_0\times\mathbf{A}_1\times\mathbf{A}_2}\left((a_0, a_1, a_2), (b_0, b_1, b_2)\right) = \left(Q^{\mathbf{A}_0}(a_0, b_0), Q^{\mathbf{A}_1}(a_1, b_1), Q^{\mathbf{A}_2}(a_2, b_2)\right)$$

It might clarify matters if we display the members of the direct product as columns instead of rows:

$$Q^{\mathbf{A}_0\times\mathbf{A}_1\times\mathbf{A}_2}\left(\begin{pmatrix} a_0 \\ a_1 \\ a_2 \end{pmatrix} \quad \begin{pmatrix} b_0 \\ b_1 \\ b_2 \end{pmatrix}\right) = \begin{pmatrix} Q^{\mathbf{A}_0}(a_0, b_0) \\ Q^{\mathbf{A}_1}(a_1, b_1) \\ Q^{\mathbf{A}_2}(a_2, b_2) \end{pmatrix}$$

We say that the operations on the direct product have be defined coordinate-wise.

A straightforward argument by induction on the complexity of a term t reveals that

$$t^{\prod_{i\in I}\mathbf{A}_i}(s^0, s^1, s^2, \ldots) = \langle t^{\mathbf{A}_i}(s_i^0, s_i^1, s_i^2, \ldots) \mid i \in I\rangle.$$

What happens when $I = \varnothing$? Then we get a function from the empty set to the empty set. A function is a set of ordered pairs satisfying an additional constraint—remember the vertical line test?. The function in this case is the empty function, which is the empty set of orderer pairs. That is,

$$\prod_{i\in I} A_i = \{\varnothing\} = 1.$$

When we apply this to an empty system of algebras, the result is a one-element algebra.

Let $\mathcal{K}$ be a class of algebras of our signature. Then

$$\mathcal{P}\mathcal{K} := \{\mathbf{B} : \mathbf{B} \text{ is isomorphic to a product of a system of algebras from } \mathcal{K}\}$$

Fact. Let $\langle \mathbf{A}_i \mid i \in I\rangle$ be a system of algebras, all of the same signature and let $s \approx t$ be any equation of the signature. Then

$$\prod_{i\in I}\mathbf{A}_i \models s \approx t \quad \text{if and only if} \quad \mathbf{A}_i \models s \approx t \text{ for all } i \in I.$$

That is, an equation holds in a direct product exactly when it holds coordinate-wise.

1.2 An Algebraic Characterization of Mod Th $\mathcal{K}$: The HSP Theorem

The HSP Theorem (Tarski's version).
Let $\mathcal{K}$ be a class of algebras, all of the same signature. Mod Th $\mathcal{K} = \mathcal{HSP}\,\mathcal{K}$.

Proof. We see that $\mathcal{HSP}\,\mathcal{K} \subseteq$ Mod Th $\mathcal{K}$. Indeed, algebras in $\mathcal{HSP}\,\mathcal{K}$ must be models of every equation true in $\mathcal{K}$, so they must be in Mod Th $\mathcal{K}$.

For the reverse inclusion, let $\mathbf{C} \in$ Mod Th $\mathcal{K}$. We will prove that $\mathbf{C} \in \mathcal{HSP}\,\mathcal{K}$.

First, let I be the set of all equations of the given signature that fail in $\mathcal{K}$. So I is the set of all equations of the signature that do not belong the Th $\mathcal{K}$. For each $s \approx t \in I$ pick an algebra $\mathbf{A}_{s\approx t} \in \mathcal{K}$ so that $s \approx t$ fails to hold in $\mathbf{A}_{s\approx t}$. Let

$$\mathbf{A} = \prod_{s\approx t\in I} \mathbf{A}_{s\approx t}.$$

Then $\mathbf{A} \in \mathcal{P}\mathcal{K}$ and an equation is true in $\mathbf{A}$ if and only if it belongs to Th $\mathcal{K}$.

Let κ be a cardinal with $|C| \le \kappa$. Above, when we introduced terms and defined the notion of term functions, we restricted our attention to the set $\{v_0, v_1, v_2, \ldots\}$ of variables. This set is countably infinite. For our current purposes, we would like to replace this set with $\{v_\alpha \mid \alpha \in \kappa\}$. This introduces no essential changes in the notions of terms and term functions apart from changing the set of variables.

Given a term t, the term function of $\mathbf{A}$ denoted by t is a certain function

$$t^{\mathbf{A}} \colon A^\kappa \to A.$$

So

$$t^{\mathbf{A}} \in A^{A^\kappa}.$$

The term functions of $\mathbf{A}$ constitute, in a natural way, a subalgebra of $\mathbf{A}^{A^\kappa}$. Indeed, let Q be any operation symbol and let r be the rank of Q. Then for any terms $t_0, \ldots, t_{r-1}$ we have

$$Q^{\mathbf{A}^{A^\kappa}}(t_0^{\mathbf{A}}, \ldots, t_{r-1}^{\mathbf{A}}) = (Q t_0 \ldots t_{r-1})^{\mathbf{A}}.$$

Let $\mathbf{F}$ denote this algebra of term functions of $\mathbf{A}$. Then $\mathbf{F}$ is the subalgebra of $\mathbf{A}^{A^\kappa}$ generated by the set $\{\rho_\alpha \mid \alpha \in \kappa\}$ of projection functions. Here $\rho_\alpha(\bar{a}) = a_\alpha$ for every κ-tuple $\bar{a}$ of elements of A. So we see that $\mathbf{F} \in \mathcal{S}\mathcal{P}\mathcal{P}\,\mathcal{K}$.

Since $|C| \leq \kappa$ there is a function from κ onto C. Denote this function by $\bar{c} = \langle c_\alpha \mid \alpha \in \kappa \rangle$. We define a homomorphism $h\colon \mathbf{F} \twoheadrightarrow \mathbf{C}$ by

$$h(t^{\mathbf{A}}) = t^{\mathbf{C}}(\bar{c})$$

for all terms t. Now we have to see that this makes sense.

The first difficulty is that there might be quite different terms s and t so that $s^{\mathbf{A}} = t^{\mathbf{A}}$. In this event, we see that $s \approx t$ is true in $\mathbf{A}$. So $s \approx t$ is also true in $\mathbf{C}$, since $\mathbf{C} \models \mathrm{Th}\,\mathcal{K} = \mathrm{Th}\mathbf{A}$. So $s^{|mathbfC}(\bar{c}) = t^{\mathbf{C}}(\bar{c})$. So at least our definition of h is sound.

To see that h is a homomorphism, let Q be any operation symbol and let r be the rank of Q. Let $t_0, \ldots, t_{r-1}$ be any terms. Observe

$$\begin{aligned} h(Q^{\mathbf{F}}(t_0^{\mathbf{A}}, \ldots, t_{r-1}^{\mathbf{A}})) &= h((Qt_0 \ldots t_{r-1})^{\mathbf{A}}) \\ &= (Qt_0 \ldots t_{r-1})^{\mathbf{C}}(\bar{c}) \\ &= Q^{\mathbf{C}}\left(t_0^{\mathbf{C}}(\bar{c}), \ldots, t_{r-1}^{\mathbf{C}}(\bar{c})\right) \\ &= Q^{\mathbf{C}}\left(h(t_0^{\mathbf{A}}), \ldots, h(t_{r-1}^{\mathbf{A}})\right) \end{aligned}$$

So h is a homomorphism.

To see that h maps F onto C just observe that $h(v_\alpha) = c_\alpha$, for each $\alpha \in \kappa$, recalling that $\bar{c}$ maps κ onto C.

Our conclusion is that $\mathbf{C} \in \mathcal{H}\mathcal{S}\mathcal{P}\mathcal{P}\,\mathcal{K}$. This is not quite what we wanted. But it is easy to see that $\mathcal{P}\mathcal{P}\,\mathcal{K} = \mathcal{P}\,\mathcal{K}$ (and anyway, this is an exercise in the next Problem Set).

□

The first version of the HSP Theorem was proven by Garrett Birkhoff in 1935. Here it is.

The HSP Theorem (Birkhoff's version).
Let $\mathcal{V}$ be a class of algebras, all of the same signature. $\mathcal{V} = \mathrm{Mod}\,\Sigma$ for some set Σ of equations if and only if $\mathcal{V}$ is closed under $\mathcal{H}$, $\mathcal{S}$, and $\mathcal{P}$.

Proof.
($\Rightarrow$) This is easy, since the truth of equations is preserved uner the formation of direct products, the formation of subalgebras, and the formation of homomorphic images.
($\Leftarrow$) Let $\Sigma = \mathrm{Th}\,\mathcal{V}$. Then

$$\begin{aligned} \mathrm{Mod}\,\Sigma &= \mathrm{Mod}\,\mathrm{Th}\,\mathcal{V} \\ &= \mathcal{H}\mathcal{S}\mathcal{P}\,\mathcal{V} && \text{using Tarski's version of the HSP Theorem} \\ &= \mathcal{H}\mathcal{S}\,\mathcal{V} \\ &= \mathcal{H}\,\mathcal{V} \\ &= \mathcal{V} \end{aligned}$$

□

Tarski actually deduced his version of the HSP Theorem from Birkhoff's version. His line of reasoning is indicated in Problem Set 2. The proof of Tarski's version that we gave above follows the reasoning of A. I. Mal'cev in 1954.

1.3 Problem Set 2

Problem Set About $\mathcal{H}$, $\mathcal{S}$, and $\mathcal{P}$

In addition to $\mathcal{H}$, $\mathcal{S}$, and $\mathcal{P}$, there is one more class operator $\mathcal{I}$: $\mathbf{A} \in \mathcal{I}\,\mathcal{K}$ if and onoy if $\mathbf{A}$ is isomorphic to a member of the class $\mathcal{K}$. From these for operators on classes of algebras (these are functions that take a class $\mathcal{K}$ of algebras as inputs and produce classes of algebras as outputs), it is possible to compose them to obtain other class operators such as $\mathcal{S}\mathcal{P}\mathcal{H}$. We can order these compound operators by taking $\mathcal{O} \leq \mathcal{Q}$ to mean $\mathcal{O}\,\mathcal{K} \subseteq \mathcal{Q}\,\mathcal{K}$ for all classes $\mathcal{K}$ of algebras, all of the same signature, whenever $\mathcal{O}$ and $\mathcal{Q}$ are class operators.

Problem 9.
Prove $\mathcal{I} \leq H$, $\mathcal{I} \leq \mathcal{S}$ and $\mathcal{I} \leq \mathcal{P}$.

Problem 10.
Prove $\mathcal{H} = \mathcal{H}\mathcal{H}$, $\mathcal{S} = \mathcal{S}\mathcal{S}$ and $\mathcal{P} = \mathcal{P}\mathcal{P}$.

Problem 11.
Prove that if $\mathcal{K} \subseteq \mathcal{L}$, then $\mathcal{O}\,\mathcal{K} \subseteq \mathcal{O}\,\mathcal{L}$, for any classes $\mathcal{K}$ and $\mathcal{L}$ of algebras, all of the same signature and for any operator $\mathcal{O} \in \{\mathcal{H}, \mathcal{S}, \mathcal{P}\}$.

Problem 12.
Repeat the last three problems, but for the compound operators $\mathcal{H}\mathcal{S}$, $\mathcal{S}\mathcal{P}$, and $\mathcal{H}\mathcal{P}$.

Problem 13.
Prove $\mathcal{S}\mathcal{H} \leq \mathcal{H}\mathcal{S}$, $\mathcal{P}\mathcal{S} \leq \mathcal{S}\mathcal{P}$, and $\mathcal{P}\mathcal{H} \leq \mathcal{H}\mathcal{P}$.

Problem 14.
Prove $\mathcal{S}\mathcal{H} \neq \mathcal{H}\mathcal{S}$, $\mathcal{P}\mathcal{S} \neq \mathcal{S}\mathcal{P}$, and $\mathcal{P}\mathcal{H} \neq \mathcal{H}\mathcal{P}$.

Problem 15.
Let $\mathcal{K}$ be the class of algebras $\langle G, \cdot \rangle$ so that $\langle G, \cdot, ^{-1}, 1\rangle$ is a finite cyclic group. Prove that $\mathcal{S}\mathcal{P}\mathcal{H}\mathcal{S}\,\mathcal{K}$, $\mathcal{S}\mathcal{H}\mathcal{P}\mathcal{S}\,\mathcal{K}$, and $\mathcal{H}\mathcal{S}\mathcal{P}\,\mathcal{K}$ are different classes.

In 1972, Don Pigozzi proved that there are exactly 18 class operators that can be compounded from $\mathcal{I}, \mathcal{H}, \mathcal{S}$, and $\mathcal{P}$. Below is a Hasse diagram of the ordered set of these operators. The comparabilities in this diagram, as well as the fact that there are no more than the 18 operators indicated, follow from the first five problems in this problem set. The incomparabilities as well as the fact that these 18 class operators are distinct from each other are more difficult to establish. Establishing these requires the construction of cleverly devised classes $\mathcal{K}$ to separate the operators.

That the compound operator $\mathcal{H}\mathcal{S}\mathcal{P}$ is at the top of this ordered set amounts to a derivation of Tarski's version of the HSP Theorem from Birkhoff's version.

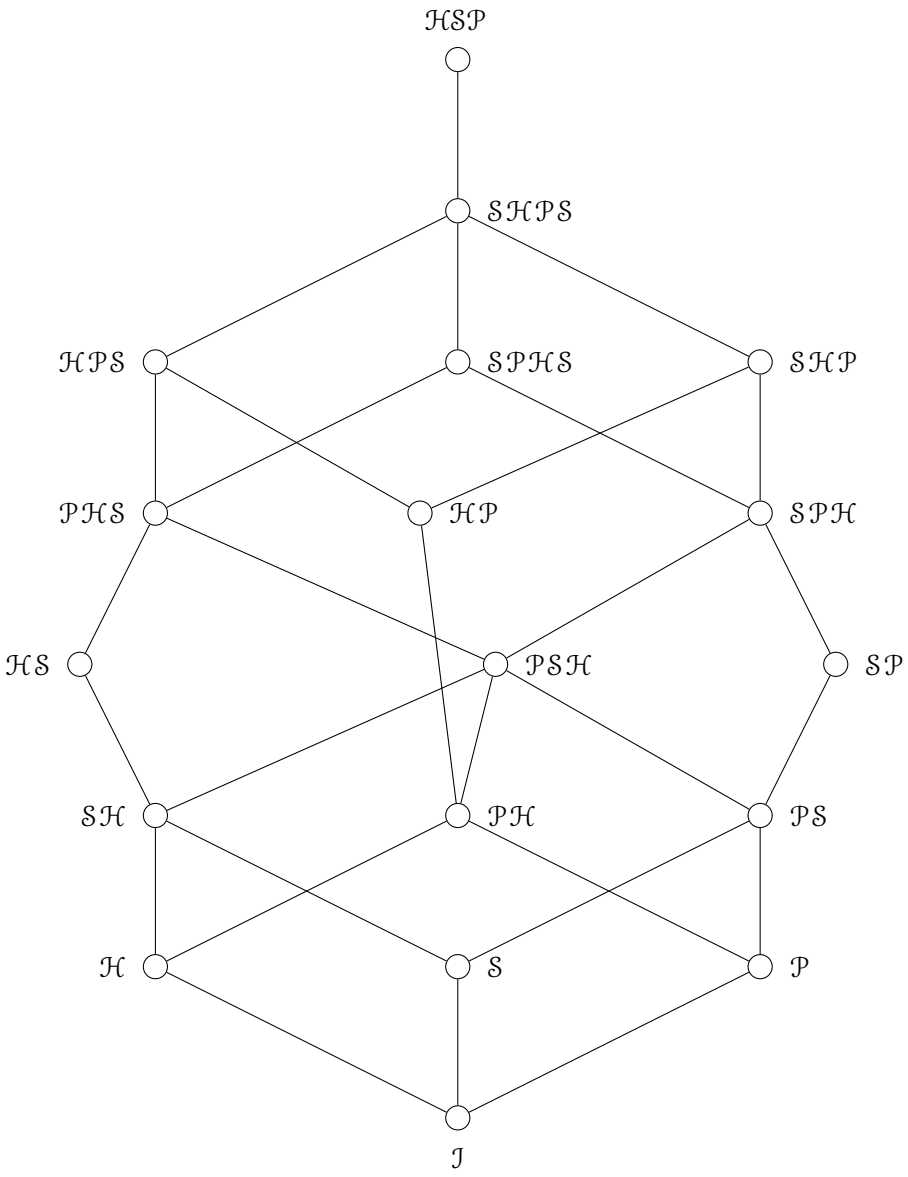

Don Pigozzi's Ordered Monoid of Class Operators

LESSON 2

THE DESCRIPTION OF $\mathrm{Th}\,\mathrm{Mod}\,\Sigma$

The HSP Theorem gives us an algebraic description of the closure operator $\mathrm{Mod}\,\mathrm{Th}$ on the algebraic side of the Galois connection established by truth between algebras and equations of a give signature. Our task now is to provide a description of the closure operator $\mathrm{Th}\,\mathrm{Mod}$ on the equational side of this Galois connection. Notice that $s \approx t \in \mathrm{Th}\,\mathrm{Mod}\,\Sigma$ relates the equation $s \approx t$ with the set Σ of equations. These are syntactical objects—made up of strings of symbols. But the closure operator is semantical, bringing in algebras and the truth of equations in algebras. What we seek here is a strictly syntactical description of the closure operator.

It helps to recast matters. Observe

$$s \approx t \in \mathrm{Th}\,\mathrm{Mod}\,\Sigma \quad \text{if and only if} \quad \mathbf{A} \models s \approx t \text{ for all } \mathbf{A} \text{ such that } \mathbf{A} \models \Sigma.$$

When the condition on the right above is satisfied we say that $s \approx t$ is a **logical consequence** of Σ.. We extend the meaning of $\models$ and denote the relation of logical conseuqence by $\Sigma \models s \approx t$. It is also convenient to contract $\mathrm{Th}\,\mathrm{Mod}$ to Th_{e}. We say that $\mathrm{Th}_{\mathrm{e}}\,\Sigma$ is the equational theory **based on** Σ and that Σ is a **basis** for this equational theory.

So what we desire is a syntactical description of logical consequence. That is, we would like to devise a means of proof for equational logic that leads from a set Σ of equational axioms by strictly syntactical means to each logical consequence $s \approx t$.

The chief attributes we require of our formal system of inference are:

Soundness: If the equation $s \approx t$ can be inferred from the set Σ of equations, then $s \approx t \in \mathrm{Th}\,\mathrm{Mod}\,\Sigma$.

Adequacy: If $s \approx t \in \mathrm{Th}\,\mathrm{Mod}\,\Sigma$, then $s \approx t$ can be inferred from Σ.

Effectiveness: In so far as it is reasonable (for computable signatures), there is an algorithm for the recognition of inferences.

In addition, we would like our system of inference to be as primitive as possible so that we can readily establish facts about inferences themselves. Since our system of inference is, in essence, a detailed description of the closure operator $\mathrm{Th}\,\mathrm{Mod}$, this will enable us to obtain far-reaching results about equational theories and varieties. While we will occasionally use our system of inference to actually deduce equations, this will not be its principal purpose. Deductions actually put forward in mathematics rarely have a completely formal, syntactical character.

Here is an informal example of an equational deduction. Any ring satisfying the equation $x^3 \approx x$ is commutative. The following derivation of this fact illustrates some of the rules of inference that are commonly used in reasoning about equations. Some steps invoking obvious uses of the axioms of ring theory are omitted. Expressions like $2yxy$ are shorthand for $yxy + yxy$—which itself abbreviates the more elaborate $(yx)y + (yx)y$.

(1)	$x^3 - x \approx 0$	by adding $-x$ to both sides of $x^3 \approx x$
(2)	$(x+y)^3 - (x+y) \approx 0$	by substituting $(x+y)$, a term, for the variable x in (1)
(3)	$(x-y)^3 - (x-y) \approx 0$	by substituting $(x-y)$ for x in (1)
(4)	$(x+y)^3 + (x-y)^3 - 2x \approx 0$	by adding (2) and (3)
(5)	$2(x^3 - x + y^2x + yxy + xy^2) \approx 0$	by standard ring theory from (4)
(6)	$2(y^2x + yxy + xy^2) \approx 0$	by replacing $x^3 - x$ by 0 in (5) using (1)
(7)	$2(y^3x - xy^3) \approx 0$	by substituting $(yx - xy)$ for x in (6) and standard ring theory
(8)	$2(yx - xy) \approx 0$	by replacing y^3 with y twice in (7)
(9)	$(x^2 - x)^3 - (x^2 - x) \approx 0$	by substituting $(x^2 - x)$ for x in (1)
(10)	$x^6 - 3x^5 + 3x^4 - x^3 - x^2 + x \approx 0$	by standard ring theory from (9)
(11)	$x^2 - 3x + 3x^2 - x - x^2 + x \approx 0$	by replacing x^3 by x six times in (10)
(12)	$3(x^2 - x) \approx 0$	by standard ring theory from (11)
(13)	$3[(x+y)^2 - (x+y)] \approx 0$	by substituting $(x+y)$ for x in (12)
(14)	$3(x^2 - x + xy + yx + y^2 - y) \approx 0$	by ring theory from (13)
(15)	$3(xy + yx) \approx 0$	by replacing $3(x^2 - x)$ and $3(y^2 - y)$ by 0 in (14)
(16)	$xy + 5yx \approx 0$	by adding (8) to (15)
(17)	$xy + 6yx \approx yx$	by adding yx to both sides of (16)
(18)	$6x^3 \approx 0$	by substituting x for y in (6)
(19)	$6x \approx 0$	by replacing x^3 by x in (18)
(20)	$xy \approx yx$	by replacing $6yx$ by 0 in (17) using (19)

In fact, any ring satisfying an equation of the form $x^n \approx x$, where $n > 1$, is a subdirect product of finite fields and hence commutative. This is a celebrated result of Nathan Jacobson. But every proof of this more general result we know is not purely syntactical, but rather uses a combination of syntactical and algebraic methods. Indeed, that mixture of the syntactic with the algebraic is typical of the derivations of equations in practice.

The kind of inference laid out above amounts to a list of equations, each of which is justified by certain axioms—here the axioms of ring theory and the equation $x^3 \approx x$—or by equations earlier in the list. Each justification is made according to some rule of inference: substituting terms for variables, replacing "equals by equals", adding "equals to equals", etc. The system of inference we are about to introduce is somewhat different. In fact, our inferences turn out to be sequences of terms rather than sequences of equations. Both the proof above and those appropriate to our system are syntactic in character—they are concerned with the formal manipulation of terms and equations considered as strings of symbols—but they embody the semantical notion of logical consequence.

2.1 Further Algebraic Preliminaries

Let $\mathbf{A}$ be an algebra. We say that θ is a **congruence relation** of $\mathbf{A}$ provided

- θ is an equivalence relation on A, and

- $Q^{\mathbf{A}}(a_0,\dots,a_{r-1})\ \theta\ Q^{\mathbf{A}}(b_0,\dots,b_{r-1})$ whenever Q is an operation symbol, where r is the rank of Q, and whenever $a_0, b_0 \dots, a_{r-1}, b_{r-1} \in A$ such that $a_i\ \theta\ b_i$ for all $i < r$.

The second item listed in this definition is sometimes called the substitution property. The notion of a congruence relation is familiar from the theories of groups and rings, where such equivalence relations are used to construct quotient groups and quotient rings. Gauss made congruences on the ring of integers central to number theory. It is easy to see that the congruence relations of $\mathbf{A}$ are precisely those sublagebras of $\mathbf{A} \times \mathbf{A}$ that happen to be equivalence relations.

When θ is a congruence relation on $\mathbf{A}$ and $a, a' \in A$ we use $a\ \theta\ a'$, $(a,a) \in \theta$, and $a \equiv a' \mod \theta$ interchangeably.

We use $\mathrm{Con}\mathbf{A}$ to denote the set of all congruence relations of $\mathbf{A}$. A routine argument shows that the intersection of any nonempty set of congruences of $\mathbf{A}$ is again a congruence of $\mathbf{A}$. Now the congruences of $\mathbf{A}$ are ordered by the inclusion relation $\subseteq$. Under this ordering, any set of congruences has a greatest lower bound, namely the intersection of all the congruences in the set, as well as a least upper bound, namely the intersection of all those congruences that contain each of the given congruences. There is a smallest conguence, namely the identity relation restricted to A. There is a largest congruence, namely $A \times A$. We use 0_A to denote the smallest congruence and 1_A to denote the largest. In this way, $\mathrm{Con}\mathbf{A}$ can be given the structure of a complete lattice.

Congruences are connected to homomorphisms in the same way that normal subgroups in groups and ideals in rings are connected to homomorphisms. At the center of this business is the notion of a **quotient algebra**. Let $\mathbf{A}$ be an algebra and let θ be a congruence of $\mathbf{A}$. For each $a \in A$ we use a/θ to denote the congruence class $\{a' \mid a' \in A \text{ and } a \equiv a' \mod \theta\}$. Moreover, we use A/θ to denote the partition $\{a/\theta \mid a \in A\}$ of A into congruence classes. We make the quotient algebra $\mathbf{A}/\theta$ by letting its universe be A/θ and, for each operation symbol Q of the signature of $\mathbf{A}$, and all $a_0, a_1, \dots, a_{r-1} \in A$, where r is the rank of Q, we define

$$Q^{\mathbf{A}/\theta}(a_0/\theta, a_1/\theta, \dots, a_{r-1}/\theta) := Q^{\mathbf{A}}(a_0, a_1, \dots, a_{r-1})/\theta.$$

Because the elements of A/θ are congruence classes, we see that the r inputs to $Q^{\mathbf{A}/\theta}$ must be congruence classes. On the left side of the equation above the particular elements a_i have no special standing—they could be replaced by any a_i' provided only that $a_i \equiv a_i' \mod \theta$. Loosely speaking, what this definition says is that to evaluate $Q^{\mathbf{A}/\theta}$ on an r-tuple of θ-classes, reach into each class, grab an element to represent the class, evaluate $Q^{\mathbf{A}}$ at the r-tuple of selected representatives to obtain say $b \in A$, and then output the class b/θ. A potential trouble is that each time such a process is executed on the same r-tuple of congruence classes, different representatives might be selected resulting in, say b', instead of b. But the substitution property, the property that distinguishes congruences from other equivalence relations, is just what is needed to see that there is really no trouble. To avoid a forest of subscripts, here is how the argument would go were Q to have rank 3. Suppose $a, a', b, b', c, c' \in A$ with

$$\begin{aligned} a/\theta &= a'/\theta \\ b/\theta &= b'/\theta \\ c/\theta &= c'/\theta. \end{aligned}$$

So a and a' can both represent the same congruence class—the same for b and b' and for c and c'. Another way to write this is

$$\begin{aligned} a &\equiv a' \mod \theta \\ b &\equiv b' \mod \theta \\ c &\equiv c' \mod \theta. \end{aligned}$$

What we need is $Q^{\mathbf{A}}(a,b,c)/\theta = Q^{\mathbf{A}}(a',b',c')/\theta$. Another way to write that is

$$Q^{\mathbf{A}}(a,b,c) \equiv Q^{\mathbf{A}}(a',b',c') \mod \theta.$$

But this is exactly what the substitution property provides. Hard-working graduate students will do the work to see that what works for rank 3 works for any rank.

Now suppose $h\colon \mathbf{A} \to \mathbf{B}$ is a homomorphism. By the **kernel** of the homomorphism h we mean

$$\ker h := \{(a,a') \mid a,a' \in A \text{ and } h(a) = h(a')\}.$$

The definition of kernel departs a bit from it use in the theories of groups and rings, but we see in the next theorem that this departure is not essential. The kernel of an homomorphism is easily seen to be a congruence. The theorem below, sometimes called the First Isomorphism Theorem, shows among other things, that congruence relations are exactly the kernels of homomorphisms.

The Homomorphism Theorem.
Let $\mathbf{A}$ *be an algebra, let* $f : \mathbf{A} \twoheadrightarrow \mathbf{B}$ *be a homomorphism from* $\mathbf{A}$ *onto* $\mathbf{B}$*, and let* θ *be a congruence relation of* $\mathbf{A}$*. All of the following hold.*

(a) *The kernel of* f *is a congruence relation of* A*.*

(b) $\mathbf{A}/\theta$ *is an algebra of the same signature as* $\mathbf{A}$*.*

(c) *The map* η *that assigns to each* $a \in A$ *the congruence class* a/θ *is a homomorphism from* $\mathbf{A}$ *onto* $\mathbf{A}/\theta$ *and its kernel is* θ*.*

(d) *If* θ *is the kernel of* f*, then there is an isomorphism* g *from* $\mathbf{A}/\theta$ *to* $\mathbf{B}$ *such that* $f = g \circ \eta$*, where* $\eta : \mathbf{A} \twoheadrightarrow \mathbf{A}/\theta$ *with* $\eta(a) = a/\theta$ *for all* $a \in A$*.*

While we provide no proof of this theorem here, we note that any proof of the cooresponding theorems for groups or rings can be easily modified to obtain such a proof.

A congruence θ of $\mathbf{A}$ is said to be **fully invariant** provided

$$a \equiv a' \mod \theta \text{ implies } f(a) \equiv f(a') \mod \theta$$

for all $a, a' \in A$ and all endomorphisms f of $\mathbf{A}$. We can say this another way. Expand the signature of $\mathbf{A}$ by adding a new one-place operation symbol to name each endomorphism of $\mathbf{A}$. Let $\mathbf{A}^* = \langle \mathbf{A}, f \rangle_{f \in \mathrm{End}\mathbf{A}}$ be the expansion of $\mathbf{A}$ by adjoining each endomorphism as a new basic operation. Then $Con\mathbf{A}^*$ is the set of all fully invariant congruences of $\mathbf{A}$.

The set T of all the terms of our given signature is the heart of our syntactical arrangements. This set becomes an algebra $\mathbf{T}$ of our signature in a natural way. Indeed let Q be any operation symbol. Let r be the rank of Q. Then we define the corresponding basic operation on $\mathbf{T}$ via

$$Q^{\mathbf{T}}(t_0,\dots,t_{r-1}) := Qt_0\dots t_{r-1},$$

for all $t_0\dots,t_{r-1} \in T$. We call the algebra $\mathbf{T}$ the **term algebra** of our signature.

The term algebra $\mathbf{T}$ has two crucial properties.

- $\mathbf{T}$ is generated by the set $\{v_i \mid i \in \omega\}$ of all variables, and
- For every algebra $\mathbf{A}$ of the signature, every function $f\colon \{v_i \mid i \in \omega\} \to A$ can be extended to a homomorphism from $\mathbf{T}$ into $\mathbf{A}$.

To see this last property, suppose that $f(v_i) = a_i$ for each $i \in \omega$. Define the extension $\widehat{g}$ of f by

$$\widehat{f}(t) \colon = t^{\mathbf{A}}(a_0, a_1, \ldots)$$

for all $t \in T$. That $\widehat{g}$ has been given a sound definition relies on the unique readiblity of terms—that is, there is no way to parse the string t as a different term. We see also that the function $\widehat{f}$ is the only way to extend f to a homomorphism. The function $\hat{f}$ is an evaluation map.

These properties of $\mathbf{T}$ are a instance of a useful notion. Let $\mathcal{K}$ be a class of algebras, all of the same signature, let $\mathbf{F}$ be an algebra of the same signature and $X \subseteq F$. Then we say that $\mathbf{F}$ is $\mathcal{K}$-**freely generated** by X provided

- The algebra $\mathbf{F}$ is generated by the set X, and
- For every algebra $\mathbf{A} \in \mathcal{K}$, every function $f \colon X \to A$ can be extended to a homomorphism $\widehat{f} \colon \mathbf{F} \to \mathbf{A}$.

This is a familiar property of vector spaces over a field $\mathbf{K}$. Let $\mathbf{F}$ be any vector space and X be a basis of $\mathbf{F}$. The $\mathbf{F}$ is $\mathcal{K}$-freely generated by X, where $\mathcal{K}$ is the class of all vector spaces over the field $\mathbf{K}$.

An examination of the proof of the HSP Theorem reveals that the algebra $\mathbf{F}$ in that proof is $\mathrm{Mod}\,\mathrm{Th}\,\mathcal{K}$-freely generated by the projections.

Let $\mathbf{A}$ be an algebra and let $\mathbf{T}$ be the term algebra of the same signature. Recalling that $s \approx t$ is just another way to denote (s, t), we contend that

$$\mathrm{Th}\,\mathbf{A} = \bigcap \{\varphi \mid \varphi \text{ is the kernel of a homomorphism from } \mathbf{T} \text{ into } \mathbf{A}\}.$$

To see this, first suppose that $s \approx t \in \mathrm{Th}\,\mathbf{A}$. This means that for every ω-tuple $\langle a_0, a_1, \ldots \rangle$ of elements of A, we have $s^{\mathbf{A}}(a_0, a_1, \ldots) = t^{\mathbf{A}}(a_0, a_1, \ldots)$. Since homomorphisms from $\mathbf{T}$ into $\mathbf{A}$ are completely determined by where they send the variables, we see that (s, t) belongs to the kernel of every such homomoprhism. Therefore $\mathrm{Th}\,\mathbf{A} \subseteq \bigcap \{\varphi \mid \varphi \text{ is the kernel of a homomorphism from } \mathbf{T} \text{ into } \mathbf{A}\}$. For the reverse inclusion, suppose $s \approx t \notin \mathrm{Th}\,\mathbf{A}$. Then there must be an ω-tuple $\langle a_0, a_1, \ldots \rangle$ of elements of A so that $s^{\mathbf{A}}(a_0, a_1, \ldots) \neq t^{\mathbf{A}}(a_0, a_1, \ldots)$. Let f be the homomorphism from $\mathbf{T}$ into $\mathbf{A}$ so that $f(v_i) = a_i$ for all $i \in \omega$. Then $(s, t) \notin \ker f$ and so $(s, t) \notin \bigcap \{\varphi \mid \varphi \text{ is the kernel of a homomorphism from } \mathbf{T} \text{ into } \mathbf{A}\}$. This establishes the reverse inclusion.

Now in the proof of the HSP Theorem we saw that, for every class $\mathcal{K}$ of algebras, all of the same signature, it was possible to devise an algebra $\mathbf{A} \in \mathcal{P}\,\mathcal{K}$ so that $\mathrm{Th}\,\mathcal{K} = \mathrm{Th}\,\mathbf{A}$. So we see that every equational theory has the form $\mathrm{Th}\,\mathbf{A}$ for some single algebra $\mathbf{A}$. This means that every equational theory is a congruence relation on the term algebra. More is true.

Theorem on Fully Invariant Congruences of Term Algebras.
Fix a signature. The equational theories of the signature are exactly the fully invariant congruence relations on the term algebra.

Proof. We have already seen that equational theories are intersections of congruences of the term algebra. So they are congruence relations themselve. First, we see that they are fully invariant.

So consider the equational theory $\mathrm{Th}\,\mathbf{A}$. Suppose that $s \approx t \in \mathrm{Th}\,\mathbf{A}$ and that f is an endomorphism of the term algebra $\mathbf{T}$. We want to establish that $f(s) \approx f(t) \in \mathrm{Th}\,\mathbf{A}$ or, what is the same, that $((f(s)), (f(t))) \in \ker h$ for every homomorphism $h \colon \mathbf{T} \to \mathbf{A}$. Now $h \circ f$ is a homomorphism from $\mathbf{T}$ into $\mathbf{A}$. Since $s \approx t \in \mathrm{Th}\,\mathbf{A}$, we know that $(s, t) \in \ker h \circ f$. But this is the same as $(f(s), f(t)) \in \ker h$, our desired conclusion.

So each equational theory is, indeed, a fully invariant congruence relation on the term algebra.

We also desire the converse. So let θ be a fully invariant congruence on $\mathbf{T}$. Let $\mathbf{A}$ be $\mathbf{T}/\theta$. We contend that $\mathrm{Th}\,\mathbf{A}$ and θ are identical. First, suppose that $(s, t) \in \theta$. Let $a_0, a_1, \ldots$ be any elements of A. Pick terms

$p_0, p_1, \ldots$ so that $a_i = p_i/\theta$ for each $i \in \omega$. Let f be the endomorphism of $\mathbf{T}$ such that $f(v_i) = p_i$ for all $i \in \omega$. Now observe

$$\begin{aligned}(s,t) \in \theta &\implies (f(s), f(t)) \in \theta \\ &\implies s^{\mathbf{A}}(a_0, a_1, \ldots) = t^{\mathbf{A}}(a_0, a_1, \ldots)\end{aligned}$$

But the ω-tuple $\langle a_0, a_1, \ldots\rangle$ was arbitrary. So $s^{\mathbf{A}} = t^{\mathbf{A}}$. That is, we have $s \approx t \in \mathrm{Th}\mathbf{A}$. Our conclusion:

$$(s,t) \in \theta \implies s \approx t \in \mathrm{Th}\mathbf{A}.$$

We also need the reverse implication. To this end, suppose $s \approx t \in \mathrm{Th}\mathbf{A}$. For each $i \in \omega$ let a_i be x_i/θ. Observe

$$\begin{aligned}s \approx t \in \mathrm{Th}\mathbf{A} &\implies s^{\mathbf{A}}(a_0, a_1, \ldots) = t^{\mathbf{A}}(a_0, a_1, \ldots) \\ &\implies s^{\mathbf{T}/\theta}(x_0/\theta, x_1/\theta, \ldots) = t^{\mathbf{T}/\theta}(x_0/\theta, x_1/\theta, \ldots) \\ &\implies s^{\mathbf{T}}(x_0, x_1, \ldots)/\theta = t^{\mathbf{T}}(x_0, x_1, \ldots)/\theta \\ &\implies s/\theta = t/\theta \\ &\implies (s,t) \in \theta\end{aligned}$$

In this way, we see that each fully invariant congruence of the term algebra $\mathbf{T}$ is indeed an equational theory. □

2.2 A Syntactic Characterization of Th Mod Σ: The Completeness Theorem for Equational Logic

Now we can reframe our task. Given a set Σ of equations, we see that Th Mod Σ is the smallest equational theory that includes the set Σ. Regarding Σ as a set of ordered pairs of terms, we see that Th Mod Σ is the smallest fully invariant congruence relation on the term algebra that includes the set Σ. That is, Th Mod Σ is the fully invariant congruence relation on $\mathbf{T}$ that is generated by the set Σ. So our task is to give a description of how the fully invariant congruence relations on the term algebra are generated from a given set Σ.

Given two terms w and r we say that w is a **subterm** of r provided there are strings u and v, possibly empty, of symbols so that $r = uwv$. The term w might occur as a subterm of r is several different ways. Given an equation $p \approx q$ and terms r and r' we will say that r and r' are **equivalent in one step using** $p \approx q$ provide there is an endomorphism f of the term algebra and strings u and v, possibly empty, so that

$$\{r, r'\} = \{uf(p)v, uf(q)v\}.$$

We denote this relation by $r \overset{p\approx q}{\longleftrightarrow} r'$. To say this another way, r' is obtained from r by replacing a substitution instance of one side of $p \approx q$ by the same subsitution applied to the other side. There is a directional variant of this notion that is useful. We say that r **rewrites in one step using** $p \approx q$ provide there is an endomorphism f of the term algebra and strings u and v, possibly empty, so that

$$r \text{ is } uf(p)v \text{ and } r' \text{ is } uf(q)v.$$

We denote this by $r \overset{p\approx q}{\longrightarrow} r'$.

We will say that the equation $s \approx t$ is **deducible** from the set Σ of equations provided there is some finite sequence of terms $r_0, r_1, \ldots, r_n$ so that

- s is r_0 and r_n is t, and

- r_i and r_{i+1} are equivalent in one step using some equation from Σ for each $i < n$.

We call the sequence $r_0, \ldots, r_n$ a **deduction** of $s \approx t$ from Σ. Notice that deductions with $n = 0$ are permitted. This means $s \approx s$, where s is any term, is deducible from every set of equations, even the empty set. We can display such deductions as

$$s \overset{e_0}{\longleftrightarrow} r_1 \overset{e_1}{\longleftrightarrow} r_2 \overset{e_2}{\longleftrightarrow} \ldots \overset{e_{n-1}}{\longleftrightarrow} t$$

where each e_i is an equation belonging to Σ.

We use $\Sigma \vdash s \approx t$ to denote that there is a deduction of $s \approx t$ from Σ.

The Completeness Theorem for Equational Logic.
Let Σ be a set of equations and let $s \approx t$ be an equation. $\Sigma \models s \approx t$ if and only if $\Sigma \vdash s \approx t$.

Proof. Let $\theta = \{(s,t) \mid \Sigma \vdash s \approx t\}$. We need to show that θ is smallest fully invariant congruence on the term algebra $\mathbf{T}$ that includes the set Σ. It is easy to see that θ is an equivalence relation on the set T of terms and that $\Sigma \subseteq \theta$.

Our first goal is to show that $\theta \subseteq \varphi$ for every fully invariant congruence relation φ on $\mathbf{T}$ that includes Σ. We do this by induction on the length of deductions.

Base Step
In this case, $\Sigma \vdash s \approx t$ is witnessed by a deduction with $n = 0$. This means that s and t are the same. Evidently, (s,s) belongs to every congruence on $\mathbf{T}$.

Inductive Step
Here we assume that $(r, r') \in \varphi$ for all fully invariant φ as long as $\Sigma \vdash r \approx r'$ is witnessed by a deduction of length no more than n. Let $\Sigma \vdash s \approx t$ be witnessed by the deduction of length $n+1$ below:

$$s = r_0 \overset{e_0}{\longleftrightarrow} r_1 \overset{e_1}{\longleftrightarrow} r_2 \ldots r_n \overset{e_n}{\longleftrightarrow} r_{n+1} = t.$$

So we see that (s, r_n) belongs to every fully invariant congruence that includes Σ and that $r_n \overset{p \approx q}{\longleftrightarrow} t$, where $p \approx q \in \Sigma$. Let φ be a fully invariant congruence that includes Σ. Let f be an endomorphism of $\mathbf{T}$ and u and v be strings of symbols so that $\{r_n, t\} = \{uf(p)v, uf(q)v\}$. It is harmless to suppose that $r_n = uf(p)v$ and $t = uf(q)v$. We know that $(f(p), f(q)) \in \varphi$ since φ is a fully invariant congruence that includes Σ. Now let y be a variable that does occurs in neither r_n nor in t. Let $t^* = uyv$. Observe that t^* is the term obtained from t by replacing the designated occurrence of term $f(q)$ by the new variable y. We contend that $f(p) \equiv f(q) \mod \varphi$ entails that $t^{*\mathbf{T}}(\ldots, f(p), \ldots) \equiv t^{*\mathbf{T}}(\ldots, f(q), \ldots) \mod \varphi$. In fact, this is true about arbitrary algebras and arbitrary congruences on them. A routine induction of the complexity of the term t^* does the job.

So our induction is complete and we know that $\theta \subseteq \varphi$ whenver φ is a fully invarient congruence of $\mathbf{T}$ that includes Σ.

So it only remains for us to show that θ itself is a fully invariant congruence of $\mathbf{T}$. We have already observe that θ is an equivalence relation on the set of terms. Let Q be an operation symbol. To avoid a morass of indices, we show the case when the rank of Q is 2. Suppose that $\Sigma \vdash s_0 \approx t_0$ and $\Sigma \vdash s_1 \approx t_1$. To see that θ is a congruence, we must have $\Sigma \vdash Qs_0s_1 \approx Qt_0t_1$. So let

$$s_0 \overset{e_0}{\longleftrightarrow} \ldots \overset{e_{n-1}}{\longleftrightarrow} t_0$$

and

$$s_1 \overset{g_0}{\longleftrightarrow} \ldots \overset{g_{m-1}}{\longleftrightarrow} t_1$$

be deductions from Σ of $s_0 \approx t_0$ and $s_1 \approx t_1$ respectively. We can piece these two deductions together:

$$Qs_0s_1 \overset{e_0}{\longleftrightarrow} \ldots \overset{e_{n-1}}{\longleftrightarrow} Qt_0s_1 \overset{g_0}{\longleftrightarrow} \ldots \overset{g_{m-1}}{\longleftrightarrow} Qt_0t_1.$$

So we find $\Sigma \vdash Qs_0s_1 \approx Qt_0t_1$, as desired. The case of operation symbols of arbitrary rank holds no mysteries. Our conclusion so far is that θ is a congruence relation of $\mathbf{T}$.

Finally, to see that θ is fully invariant, suppose that $\Sigma \vdash s \approx t$ and the f is an endomorphism of $\mathbf{T}$. We need to see that $\Sigma \vdash f(s) \approx f(t)$. The idea is straightforawrd. Let

$$s \overset{e_0}{\longleftrightarrow} r_1 \overset{e_1}{\longleftrightarrow} r_2 \overset{e_2}{\longleftrightarrow} \ldots \overset{e_{n-1}}{\longleftrightarrow} t$$

be a deduction of $s \approx t$ from Σ. We claim that

$$f(s) \overset{e_0}{\longleftrightarrow} f(r_1) \overset{e_1}{\longleftrightarrow} f(r_2) \overset{e_2}{\longleftrightarrow} \ldots \overset{e_{n-1}}{\longleftrightarrow} f(t)$$

is also a deduction of $f(s) \approx f(t)$ from Σ. To establish this, we need only consider an arbitrary step in the deduction. So let us suppose that $r \overset{p\approx q}{\longleftrightarrow} r'$. It is harmless to suppose that $r = ug(p)v$ and $r' = ug(q)v$ where g is an endomorphism of $\mathbf{T}$ and u and v are certain strings of symbols. The $f(r) = \widehat{u}f(g(p))\widehat{v}$ and $f(r') = \widehat{u}f(g(q))\widehat{v}$. Here $\widehat{u}$ is obtained from u by replacing each variable v_i by the term $f(v_i)$. The string $\widehat{v}$ is obtained from v in the same way. But notice that $f \circ g$ is itself an endomorphism of $\mathbf{T}$. This means $f(r) \overset{p\approx q}{\longleftrightarrow} f(r')$, as desired. It follows that θ is a fully invariant congruence of $\mathbf{T}$. Since $\Sigma \subseteq \theta \subseteq \varphi$ for each fully invariant congruence φ that includes Σ, we find that θ is the least fully invariant congruence of $\mathbf{T}$ that includes Σ. In this way, our theorem is established. □

Birkhoff proved the Theorem on Fully Invariant Congruences of Term Algebras in 1935 and drew from it a completeness theorem for equational logic. Loosely speaking, in Birkhoff's framework a deduction is a sequence of equations, rather than a sequence of terms. His rules of inference reflect the definition of fully invariant congruence relation on the term algebra. Birkhoff system of equational inference can be found in Problem Set 3. There you can also find a somewhat different system of equational inference put forward by Tarski. The system we have given, part of the folklore, was inspired by Mal'cev's description of how to generated congruence relations in arbitrary algebras. It is convenient for giving proofs on the lenghth of deductions.

It is frequently possible to prove theorems concerning deducibilty by induction on the length of derivations. We say that a set Γ of terms is **closed with respect to deductions based on** Σ iff $t \in \Gamma$ whenever $s \in \Gamma$ and $\Sigma \vdash s \approx t$.

The Principle of Induction on Deductions.
Let Γ be any set of terms and Σ be any set of equations. If

$$s \overset{e}{\longleftrightarrow} t \quad \textit{implies} \quad t \in \Gamma \quad \textit{whenever} \quad s \in \Gamma \quad \textit{and } e \in \Sigma$$

then Γ is closed with respect to deductions based on Σ. □

Now that we have a formal system of inference in hand, we invite the reader to write out a derivation of $xy \approx yx$ from $x^3 \approx x$ and the axioms of ring theory. The deduction provided at the beginning of this section should be of help. (We should also warn the reader that a fully detailed derivation of this within our formal system is fairly long.)

One of the advantages of our system of inference is that it gives us a simple test for detecting equational theories.

Corollary 2.2.1. *Let T be a set of equations. The following statements are equivalent:*

i. *T is an equational theory.*

ii. (a) *$s \approx s \in T$ for all terms s.*

(b) $s \approx t \in T$ *whenever* $s \approx r \in T$ *and* $r \xrightarrow{e} t$ *for some term* r *and some equation* $e \in T$.

iii. (a) $s \approx s \in T$ *for all terms* s.

(b) T *is closed under substitution.*

(c) $s \approx t \in T$ *whenever* t *is obtained by replacement for* s *on the basis of some equation in* T.

(d) $s \approx t \in T$ *whenever,* $s \approx r$ *and* $r \approx t$ *belong to* T, *for some term* r.

Proof. The conditions in (ii) assert little more than that T is closed under deductions of lengths zero, one, and two. We argue by induction on the length of deducttions that T is closed under all derivations. Suppose inductively that $s \xrightarrow{p}_1 \xrightarrow{p}_2 \xrightarrow{\cdots} \xrightarrow{p}_{n-1} \xrightarrow{t}$. The inductive hypothesis gives $s \approx p_{n-1} \in T$. Set $r = p_{n-1}$ and let $e \in T$ be an equation such that $r \xrightarrow{e} t$. By (b) $s \approx t \in T$. So (ii) implies (i). That (i) implies (ii) is immediate from the Completeness Theorem for Equational Logic.

(iii) is an easy consequence of (i) and the Completeness Theorem for Equational Logic. Suppose that (iii) holds and that $s \approx r \in T$ and that $r \xrightarrow{e} t$, where $e \in T$. By (iii–b) and (iii–c), $r \approx t \in T$. By (iii-d), $s \approx t \in T$. This means that (iii) implies (ii). □

Just as knowing that $\mathcal{K}$ is a variety iff $\mathcal{K} = \mathcal{HSP}\,\mathcal{K}$ gives us a way to check, in some instances, whether a given class of algebras is a variety, the corollary above often allows us to determine whether a given set of equations is an equational theory.

2.3 Problem Set 3

Problem Set About Equational Inference

Problem 16.
Let **A** be any algebra and Γ be any collection of fully invariant congruence relations on **A**. Prove that the join $\bigvee\Gamma$ of the set Γ in Con**A** is again a fully invariant congruence relation on **A**.

Problem 17.
Write down a detailed definition of the notion "the tree T **depicts** the term t."

Problem 18.
Prove that a sequence s of symbols is a subterm of the term t iff s is a term and $t = AsB$ for some possibly empty strings A and B of symbols.

Problem 19.
Are there substitutions f and g such that $f(x+(x+x)) = g((x+x)+x)$?

Problem 20.
Let F and G be unary operation symbols and let $s = F^2G^2FGx$ and $t = F^2G^3FGx$. Let u be a subterm of s that is different from x. Is there a substitution instance of u that is also a substitution instance of t—possibly by means of a different substitution?

PROBLEM 21.
(Birkhoff's System of Inference, Birkhoff **1935**) Fix a signature. In this system there are three kinds of rules of inference:

The Substitution Rule: From $\{e\}$, it is permitted to infer any substitution instance of e.

Equivalence Rules: From the any set of equations, it is permitted to infer $s \approx s$, for any term s.

From $\{s \approx t\}$ it is permitted to infer $t \approx s$.

From $\{s \approx t, t \approx u\}$ it is permitted to infer $s \approx u$.

Rules for Operating on Equations: For each operation symbol Q, it is permitted to infer the equation

$$Qp_0p_1\dots p_{r-1} \approx Qq_0q_1\dots q_{r-1}$$

from $\{p_i \approx q_i : i < r\}$, where r is the rank of Q.

Take $\Sigma \vdash_B e$ to mean that there is a finite sequence $e_0, e_1, \dots, e_n$ of equations such that e is just e_n and each member of the sequence either belongs to Σ or is obtainable from some subset of its predecessors by one of the rules of inference above. Prove that

$$\Sigma \vdash_B e \text{ iff } \Sigma \vdash e$$

for all sets Σ of equations and all equations e.

PROBLEM 22.
(Tarski's System of Inference, Tarski **1968**) In this system there are three rules of inference:

The Substitution Rule: From $\{e\}$ it is permitted to infer any substitution instance of e.

The Tautology Rule: From the any set of equations it is permitted to infer any tautology, i.e. any equation of the form $s \approx s$.

The Replacement Rule: From $\{s \approx t, e\}$ it is permitted to infer $t \approx u$ where u is obtained from s by replacement on the basis of e.

Take $\Sigma \vdash_T e$ to mean that there is a finite sequence of equations that ends with e such that each member of the sequence either belongs to Σ or can be inferred from some subset of its predecessors in the sequence by means of the rules given above. Prove that

$$\Sigma \vdash_T e \text{ iff } \Sigma \vdash e$$

for all sets Σ of equations and every equation e.

PROBLEM 23.
Devise formal inferences of $xy \approx yx$ from $x^3 \approx x$ and the usual axioms of ring theory, using the formal system of inference presented in this section and each of the systems described in the two preceding exercises.

PROBLEM 24 (harder)**.**
From the usual axioms of the theory of commutative rings with unit, supplemented by $x^6 \approx x$, find a derivation of $x + x \approx 0$.

PROBLEM 25 (harder)**.**
(Hand **1976**) From the usual axioms of the theory of commutative rings with unit, supplemented by $x^{48} \approx x$, find a derivation of $x^2 \approx x$.

PROBLEM 26 (harder)**.**
(Levi **1944** and Stormquist **1974**) Let Γ be the usual set of axioms for group theory. Prove that

$$\Gamma \cup \{x^p y^p \approx (xy)^p, x^q y^q \approx (xy)^q\} \vdash xy \approx yx \text{ iff } 2 = \gcd(p^2 - p, q^2 - q).$$

PROBLEM 27.
Fix a similarity type σ with no constant symbols. Denote by $\mathbf{2}_\sigma$ the algebra of type σ with universe $2 = \{0, 1\}$ such that

$$F(\bar{a}) = \begin{cases} 1, & \text{if every entry in } \bar{a} \text{ is } 1 \\ 0. & \text{otherwise} \end{cases}$$

for each fundamental operation F. Prove that $s \approx t$ is regular iff $\mathbf{2}_\sigma \models s \approx t$, for all equations $s \approx t$ of type σ.

PROBLEM 28 (harder)**.**
(Graczyńska **1983**) Let T be any equational theory, containing nonregular equations, in a similarity type without constant symbols. Prove that Reg T $\cup \{e\} \vdash$ T, for any nonregular equation $e \in$T.

PROBLEM 29 (harder)**.**
(Płonka **1967**) Let σ be a similarity type with no constant symbols. Let $\mathbf{S} = \langle S, \curlyvee \rangle$ be a semilattice and denote by $\leq$ the join semilattice order on $\mathbf{S}$. Let $\langle \mathbf{A}_i : i \in S \rangle$ be a system of algebras such that A_i and A_j are disjoint whenever i and j are distinct elements of S. Finally, let $H = \langle h_{ij} : i, j \in S \text{ and } i \leq j \rangle$ be a system of homomorphisms such that

$$h_{ij} : \mathbf{A}_i \to \mathbf{A}_j \text{ for } i \leq j \text{ in } \mathbf{S}$$

$$h_{ii} \text{ is the identity map on } \mathbf{A}_i \text{ for all } i \in S, \text{ and}$$

$$h_{ik} = h_{jk} \circ h_{ij} \text{ for } i \leq j \leq k \text{ in } \mathbf{S}$$

The **Płonka sum** of the system $\langle \mathbf{A}_i : i \in S \rangle$ with respect to the semilattice $\mathbf{S}$ and the system H of homomorphisms is the algebra with universe $A = \bigcup_{i \in S} A_i$ such that for any operation symbol Q and any $a_0, \ldots, a_{r-1} \in A$, where r is the rank of Q, we have

$$Q^{\mathbf{A}}(a_0, \ldots, a_{r-1}) = Q^{\mathbf{A}_k}(h_{i_0 k}(a_0), \ldots, h_{i_{r-1} k}(a_{r-1})$$

where $a_0 \in A_{i_0}, \ldots, a_{r-1} \in A_{i_{r-1}}$ and $k = i_0 \curlyvee \cdots \curlyvee i_{r-1}$. For any class $\mathcal{K}$ of algebras, we say that $\mathcal{K}$ is **closed with respect to Płonka sums** provided every algebra isomorphic to a Płonka sum of a system of algebras from $\mathcal{K}$ belongs, itself, to $\mathcal{K}$. Prove that $\mathcal{V}$ is closed with respect to Płonka sums iff $\mathcal{V}$ can be axiomatized by regular equations, for any variety $\mathcal{V}$.

PROBLEM 30.
An equation $s \approx t$ is **balanced** iff whenever u is a variable or an operation symbol with rank less than two, then $|s|_u = |t|_u$. Prove that in any similarity type the set of balanced equations is an equational theory.

PROBLEM 31.
Let u be any term and set $\Delta^u = \{s \approx t : s = t$ or both $u \lhd s$ and $u \lhd t\}$. Prove that Δ^u is an equational theory for every term u.

PROBLEM 32.
Let $\Sigma = \{x(yz) \approx (xy)z\}$. Since it is irrelevant with respect to Σ how a term is associated, in this exercise we suppress parentheses. For any natural number n, let e_n denote the following equation:

$$v_0 v_1 \dots v_{n-1} v_n v_n v_{n-1} \dots v_1 v_0 \approx v_n v_{n-1} \dots v_1 v_0 v_0 v_1 \dots v_{n-1} v_n$$

Prove the $\Sigma \cup \{e_i : i < n\} \nvdash e_n$, for every natural number n.

PROBLEM 33 (harder)**.**
Call a similarity type **bold** provided it has an operation symbol of rank at least two or at least two unary operation symbols. Prove that there are 2^ω equational theories for any countable bold similarity type.

LESSON 3

FIRST INTERLUDE: THE RUDIMENTS OF LATTICE THEORY

Lattice theory is a branch of algebra, just as group theory in a branch of algebra. Lattices arise naturally in the course of investigating equational logic, so you will find gathered here key definitions and examples, as well as the basic facts about lattices that will be useful later in this exposition. On the other hand, no proofs will be provided—apart from some sketched in the problem set. Lattice theory will also provide us examples of equational theories.

3.1 BASIC DEFINITIONS AND EXAMPLES

A lattices can be construed as an algebra $\mathbf{L} = \langle L, \vee, \wedge \rangle$ with two two-place basic operations called join and meet. A lattice can also be contrued as an ordered set $\mathbf{L} = \langle L, \leq \rangle$. This works like two sides of the same coin. We reserve the word **lattice** for the algebraic version and use **lattice ordered set** for the other version. On the ordered set side, $\leq$ is a partial ordering of L such that for all $x, y \in L$, there is a least upper bound of x and y, as well as a greated lower bound. On the algebraic side, evaluating the join $\vee$ produces the least upper bound and evaluating the meet $\wedge$ produces the greatest lower bound. Given a lattice ordered set, the join and meet can be defined by elementary formulas. Given a lattice, the ordering can be defined via

$$x \leq y \iff x \vee y \approx y \iff x \wedge y \approx x.$$

The class of lattices is a variety, based on a small handful of easily understood equations. The class of lattice ordered sets, on the other hand, is axiomatized by small set of easily understood elementary sentences, which, however, have more involved syntactical structure requiring several alternations of quantifiers. Here are both systems of axioms:

An Equational Base for the Class of all Lattices

$$x \vee (y \vee z) \approx (x \vee y) \vee z \qquad x \wedge (y \wedge z) \approx (x \wedge y) \wedge z$$
$$x \vee y \approx y \vee x \qquad x \wedge y \approx y \wedge x$$
$$x \vee x \approx x \qquad x \wedge x \approx x$$
$$x \vee (x \wedge y) \approx x \qquad x \wedge (x \vee y) \approx x$$

An Axiomatization of the Class of all Lattice-Ordered Sets

$$\forall x[x \leq x]$$
$$\forall x \forall y[(x \leq y \between y \leq x) \implies x \approx y]$$
$$\forall x \forall y \forall z[(x \leq y \between y \leq x) \implies x \leq z]$$
$$\forall x \forall y \exists z[z \leq x \between z \leq y \between (\forall u[u \leq x \between u \leq y \implies u \leq z])]$$
$$\forall x \forall y \exists z[x \leq z \between y \leq z \between (\forall u[x \leq u \between y \leq u \implies z \leq u])]$$

The equational base, the first two lines assert associativity and commutativity for both join and meet. The property reflected in the third line is called idempotence, while the last line constains the two absorption laws. The equational base gives us the curious fact that if $\langle L, \vee, \wedge \rangle$ is a lattice, then so is $\langle L, \wedge, \vee \rangle$. You cannot interchange plus and times in a ring and expect the result to still be a ring!

The first three sentences in the axiomatization of lattice ordered sets just assert the reflexive, antisymmetric, and transitive properties of the orderig, while the last two sentences assert the existence of greatest lower bounds and least upper bounds. It only takes a little work to see that if $\langle L <\leq \rangle$ is a lattice, then so in $\langle L, \geq \rangle$. This is the analog of interchanging join and meet on the algebraic side. It means that turning a lattice upside down results again in a lattice.

Lattices arise naturally in several ways. We have already seen that the closed sets on either side of a Galois connection comprise lattices—in fact, they are just the upside down versions of each other. Indeed, more general closure systems give rise to lattices of closed set. Here is familiar example: the set of natural numbers is lattice ordered by divisibility. This ordering puts 1 at the bottom of the lattice (since $1 \mid n$ for every natural number n) and 0 at the top (since $n \mid 0$ for every natural number n). The meet in this lattice is just the greatest common divisor and the join is the least common multiple.

On of the attractive features of lattice theory is that lattices can be displayed in Hasse diagrams. In these diagrams, the vertices are the elements of the lattice, the edges give the covering relation (the is nothing in between), and getting higher in the diagram reflects getting larger in the ordering. Of course, this works best for finite lattices.... Here are two important lattices.

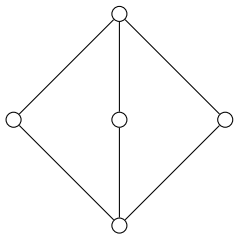

The Lattice $\mathbf{M}_3$

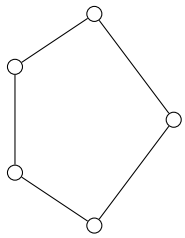

The Lattice $\mathbf{N}_5$

3.2 The First Facts of Lattice Theory

A lattice in which the equation $x \wedge (y \vee z) \approx (x \wedge y) \vee (x \wedge z)$ holds is said to be **distributive**. Of course, there is another distributive law: $x \vee (y \wedge z) \approx (x \vee y) \wedge (x \vee z)$. Conveniently, it turns out that a lattice in which one of these distributive laws holds, the other must also hold. While the analogy is limited, the variety of distributive lattice plays a role in lattice theory akin the to role that the variety of Abelian groups plays in group theory. Distributive lattice, particularly finite distributive lattices, have a much nicer structure than lattices in general, just as Ablian groups,particularly finite Abelian groups, have a much nicer structure than groups is general.

Fact 1. Let $\mathbf{L}$ be a lattice. The following are equivalent:

(a) $\mathbf{L}$ is a distributive lattce.

(b) $\mathbf{L} \models x \vee (y \wedge z) \approx (x \vee y) \wedge (x \vee z)$.

(c) $\mathbf{L} \models x \wedge (y \vee z) \leq (x \wedge y) \vee (x \wedge z)$.

(d) $\mathbf{L} \models (x \vee y) \wedge (x \vee z) \leq x \vee (y \wedge z)$.

(e) The lattice $\mathbf{N}_5$ is not isomorphic to any sublattice of $\mathbf{L}$ and neither is the lattice $\mathbf{M}_3$.

Fact 2. Let $\mathbf{L}$ be any lattice. The lattice $\mathrm{Con}\,\mathbf{L}$ of congruences of $\mathbf{L}$ is a distributive lattice.

There is a way to weaken the distributive law that results in a very important class of lattices. The modular law, discovered by Richard Dedekind, is

$$\forall x \forall y \forall z [x \leq z \implies x \vee (y \wedge z) \approx (x \vee y) \wedge (x \vee z)].$$

So every distributive lattice is modular. The converse is false. $\mathbf{M}_3$ is a distributive lattice that fails to be modular. The modular law, as formulated above, is not an equation. How it can be replaced by an equation, as asserted in the next Fact. So the class of modular lattices is itself a variety.

Fact 3. Let $\mathbf{L}$ be any lattice. The following are equivalent:

(a) $\mathbf{L}$ is a modular lattice.

(b) $\mathbf{L} \models \big((x \wedge z) \vee y\big) \wedge z \approx (x \wedge z) \vee (y \wedge z)$.

(c) $\mathbf{L} \models \big((x \vee z) \wedge y\big) \vee z \approx (x \vee z) \wedge (y \vee z)$.

(d) $\mathbf{L} \models \big((x \wedge z) \vee y\big) \wedge z \leq (x \wedge z) \vee (y \wedge z)$.

(e) $\mathbf{L} \models (x \vee z) \wedge (y \vee z) \leq \big((x \vee z) \wedge y\big) \vee z$.

(f) The lattice $\mathbf{N}_5$ is not isomorphic to any sublattice of $\mathbf{L}$.

Fact 4. Let $\mathbf{A}$ be any group or any ring or any module. The lattice $\mathrm{Con}\,\mathbf{A}$ of congruences of $\mathbf{A}$ is a modular lattice.

One consequence of the characterizations above of distributive and of modular lattices is that turning them upside down (i.e. interchange join and meet) results in a lattice that is distributive or modular as the case may be.

There is another way to weaken the distributive law. A lattice $\mathbf{L}$ is said to be **meet-semidistributive** provided the sentence below in true in $\mathbf{L}$.

$$\forall x \forall y \forall z [x \wedge y \approx x \wedge z \implies x \wedge (y \vee z) \approx (x \wedge y) \vee (x \wedge z)]. \qquad (\mathrm{SD}_\wedge)$$

Meet-semidistributive lattices arise as congruence lattices of any algebra that has an associative, commutative, idempotent two-place basic operation. The modular lattice $\mathbf{M}_3$ fails to be meet-semidistributive, as does the lattice depicted on the left below:

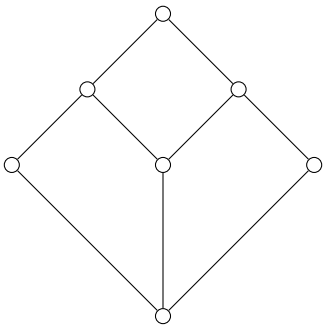

Not meet-semidistributive

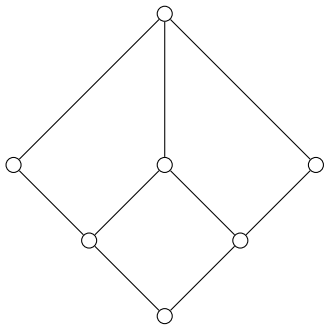

meet-semidistributive

However, the upside-down version on the right is meet-semidistributive. Neither of these lattices is modular (can you spot the $\mathbf{N}_5$'s?).

Now let A be any set. We use Eqv A to denote the set of all equivalence relations on the set A. While it is easy to verify that the intersection of any nonempty collection of equivalence relations is again an equivalence relation, the union of even two equivalence relations will usually fail to be transitive. Nevertheless, $eqvA$ is lattice ordered by $\subseteq$. While we take $\varphi \wedge psi = \varphi \cap \psi$ for any two equivalence relations φ and ψ, we must take $\varpi \vee \psi$ to be the transitive closure of $\varphi \cup \psi$. Actually,

$$\varphi \vee \psi = \varphi \cup \varphi \circ \psi \cup \varphi \circ \psi \circ \varphi \cup \varphi \circ \psi \circ \varphi \circ \psi \cup \ldots,$$

where $R \circ S = \{(a,c) \mid a,c, \in A$ and there is $b \in A$ such that $aRbSc\}$ for any relations R and S on A. So the equivalence relations on A form a lattice **Eqv** A. In lattice theory, the lattice **Eqv** A as a role similar to the role group of permutations **Sym** A plays in group theory.

Fact. Let **L** be any lattice. There is a set A so that **L** is isomorphic to a sublattice of **Eqv** A.

Now suppose **A** is an algebra with universe A. How do **Con A** and **Eqv** A compare? Since every congruence relation is an equvialence relation we might expect **Con A** to be a sublattice of **Eqv** A. For this to be true, the meet and join in **Con A** must be the restrictions to congruence relations of the meet and join in **Eqv** A. With the help of the epression displayed several lines above, you can easily work out the details for $\vee$. Since it both lattice the meet is just intersection, we have the next fact.

Fact. Let **A** be an algebra with universe A. Then **Con A** is a sublattice of **Eqv** A.

3.3 Problem Set 4

Problem Set About Lattices

Problem 34.
Prove that if $\langle L, \leq \rangle$ is a lattice ordered set, then $\langle L, \vee, \wedge \rangle$ is a lattice, $a \vee b$ is the least upper bound of $\{a,b\}$ and $a \wedge b$ is greatest lower bound of $\{a,b\}$ for all $a,b \in L$.

Problem 35.
Prove that if $\langle L, \vee, \wedge \rangle$ is a lattice, then $\langle L, \leq \rangle$ is a lattice ordered set, where $a \leq b$ means that $a \vee b = b$ for all $a,b \in L$.

Problem 36.
Prove the Fact that characterizes distributive lattices.

Problem 37.
Prove the Fact that characterizes modular lattices.

Problem 38.
Let **A** be any algebra with universe A. Prove that **Con A** is a sublattice of **Eqv** A.

Problem 39.
Prove the the lattice of normal subgroups of any group is a modular latice.

LESSON 4

EQUATIONAL THEORIES THAT ARE NOT FINITELY AXIOMATIZABLE

One question that immediately presents itself about any give euational theory is whether that theory is finitely based. As the subjects are usually presented, rings, groups, lattices, and Boolean algebras, among others, are defined via finite sets of equations. This can be a bit misleading. For example, groups first arose in a concrete setting: they were sets of permutations endowed with the operations of composition of permutations, formation of inverse permutations, and in identity permutation. Abstract groups are those algebras that are isomorphic to concrete groups. In this light, the familiar theorem usually called the Cayley Representation Theorem should really be called the Cayley Finite Basis Theorem. It gives a finite list of equations to axiomatize the class of all (abstract) groups. On the other hand, the equational axioms for rings emerged as a finite list of properties common to a diverse assortment of algebras.

One might ask whether a given variety or even a given algebra is finitely based. Even restricted to finite algebras of finite signature, this question as turned out to be subtle. Our current concern will be to give examples of finite algebras that are not finitely based and devise general means to construct such examples.

It has turned out that almost all the finite algebras that emerged in the 19^{th} century are finitely based. This applies to each finite group, each finite ring, and each finite lattice—although this is by no means obvious. Indeed, in an asymptotic sense, almost every finite algebra is finitely based. As a consequence, most nonfinitely based finite algebras seem pathological.

4.1 THE BIRKHOFF BASIS

An algebra $\mathbf{A}$ is **locally finite** provided every finitely generated subalgebra of $\mathbf{A}$ is finite. A class $\mathcal{K}$ of algebras is locally finite when each algebra belonging to $\mathcal{K}$ is locally finite.

Fact. Every variety generated by a finite algebra is locally finite.

Proof. Let $\mathbf{A}$ be a finite algebra and let $\mathcal{V} = \mathcal{HSP}\,\mathbf{A}$. Let $\mathbf{C}$ be an algebra in $\mathcal{V}$ that is generated by the finite set X. Recall from the proof of the HSP Theorem that we made a subalgebra $\mathbf{B}$ of $\mathbf{A}^{A^C}$ that was generated by the projection function ρ_c for each $c \in C$. Then we formed the homomorphism $h : \mathbf{B} \twoheadrightarrow \mathbf{C}$. We make here a small change. Instead, let $\mathbf{B}$ be the subalgebra of $\mathbf{A}^{A^X}$ generated by the projections ρ_x for each $x \in X$. Then $\mathbf{B}$ will be finite since A^{A^X} is finite. So $\mathbf{C}$ must be finite as well. □

A small change in the argument above yields the next Fact.

Fact. Let $\mathcal{V}$ be a locally finite variety and let n be a natural number. Then $\operatorname{Th}\mathcal{V}$ induces an equivalence relation of the set of terms on $\{v_0, \ldots, v_{n-1}\}$ that has only finitely many equivalence classes.

Proof. Suppose that $s \approx t$ is an equation with variables all drawn from $\{v_0, \ldots, v_{n-1}\}$ that fails in $\mathcal{V}$. Then there must be an algebra $\mathbf{C} \in \mathcal{V}$ in which $s \approx t$ fails. Moreover, we can insist that $\mathbf{C}$ is generated by a set X of cardinality n. Now in the proof of the HSP Theorem we produced an algebra $\mathbf{A}$ so that $\mathcal{V} = \mathcal{HSP}\,\mathbf{A}$. As in the proof of the Fact above, we see $\mathbf{C}$ is a homomorphic image of a finitely generated subalgebra $\mathbf{B}$ of $\mathbf{A}^{A^X}$. Since $\mathcal{V}$ is locally finite, we have that $\mathbf{B}$ is finite. But then so it $\mathbf{C}$. □

Let $\mathcal{V}$ be a locally finite variety and let n be a natural number. By $\mathcal{V}^{(n)}$ we mean the class of algebras, of the same signature as $\mathcal{V}$, that are models of all the equations true in $\mathcal{V}$ in no more than n distinct variables occur. So $\mathcal{V}^{(n)}$ is a variety and $\mathcal{V} \subseteq \mathcal{V}^{(n)}$. It is easy to check that

$$\mathcal{V} \subseteq \cdots \subseteq \mathcal{V}^{(n+1)} \subseteq \mathcal{V}^{(n)} \subseteq \cdots \subseteq \mathcal{V}^{(1)} \subseteq \mathcal{V}^{(0)} \qquad \text{and} \qquad \mathcal{V} = \bigcap_{n \in \omega} \mathcal{V}^{(n)}.$$

The sequence $\langle \mathcal{V}^{(n)} \mid n \in \omega \rangle$ is called the **descending varietal chain of** $\mathcal{V}$.

Fact. Let $\mathcal{V}$ be a variety and n be a natural number. Then for all algebras $\mathbf{B}$, we have $\mathbf{B} \in \mathcal{V}^{(n)}$ if and only if every subalgebra of $\mathbf{B}$ with n or fewer generators belongs to $\mathcal{V}$.

Proof. First suppose that $\mathbf{B} \in \mathcal{V}^{(n)}$ and $\mathbf{C}$ is a subalgebra of $\mathbf{B}$ generated by $\{b_0, \ldots, b_{n-1}\}$. To see that $\mathbf{C} \in \mathcal{V}$ let $s \approx t$ be any equation true in $\mathcal{V}$. Pick $c_0, c_1, c_2, \cdots \in C$. For each natural number k pick a term $p_k(x_0, \ldots, x_n - 1)$ so that $c_k = p^{\mathbf{B}}(b_0, \ldots, b_{n-1})$. Then the equation

$$s(p_0, p_1, p_2, \ldots) \approx t(p_0, p_1, p_2, \ldots)$$

is a logical consequence of $s \approx t$ and so is true in $\mathcal{V}$. But the displayed equation has only variables from the set $\{x_0, x_1, \ldots, x_{n-1}\}$. So the displayed equation is true in $\mathbf{B}$, since $\mathbf{B} \in \mathcal{V}^{(n)}$. So

$$s^{\mathbf{C}}(p_0^{\mathbf{C}}(b_0, b_1, \ldots, b_{n-1}), \ldots) = t^{\mathbf{C}}(p_0^{\mathbf{C}}(b_0, b_1, \ldots, b_{n-1}), \ldots)$$
$$s^{\mathbf{C}}(c_0, c_1, \ldots) = t^{\mathbf{C}}(c_0, c_1, \ldots).$$

Since $c_0, c_1, \ldots$ were arbitrary elements of C, we see that $s \approx t$ holds true in $\mathbf{C}$. This entails that $\mathbf{C} \in \mathcal{V}$.

For the converse, suppose $\mathbf{B} \notin \mathcal{V}^{(n)}$. Then there is some equation $s \approx t$, in which at most n distinct varialbles occur, that is true in $\mathcal{V}$ but fails in $\mathbf{B}$. So pick $b_0, b_1, \ldots, b_{n-1} \in B$ so that $s^{\mathbf{B}}(b_0, b_1, \ldots, b_{n-1} \neq t^{\mathbf{B}}(b_0, b_1, \ldots, b_{n-1})$. Let $\mathbf{C}$ be the subalgebra of $\mathbf{B}$ generated by $\{b_0, b_1, \ldots, b_{n-1}\}$. Evidently, the equation $s \approx t$ fails in $\mathbf{C}$. This means that $\mathbf{B}$ has a subalgebra, generated by at most n elements, that does not belong to $\mathcal{V}$. □

Birkhoff's Finite Basis Theorem.
Let $\mathcal{V}$ be a locally finite variety of finite signature and let n be a natural number. Then $\mathcal{V}^{(n)}$ is finitely based.

Proof. We can assume that $\mathcal{V}$ is not the trivial variety, since otherwise $\mathcal{V}^{(n)}$ is also trivial (and hence finitely based) for all $n \geq 2$. We leave it in the hands of the eager graduate students to devise finite bases for the varieties $\mathcal{V}^{(1)}$ and $\mathcal{V}^{(0)}$ in case $\mathcal{V}$ is the trivial variety. (It might help to read through the rest of this proof. . . .)

Let $\mathbf{T}_n$ be the subalgebra of the term algebra consisting of those terms in which only variables from $\{x_0, x_1, \ldots, x_{n-1}\}$ occur. Let $\theta = \{(p, q) \mid p, q \in T_n \text{ and } \mathcal{V} \models p \approx q\}$. It is easy to check that θ is a congruence of $\mathbf{T}_n$ and that $\mathbf{T}_n/\theta$ is the $\mathcal{V}$-freely generated algebra on n free generators. (It is even easier to see that $\mathbf{T}_n/\theta \in \mathcal{V}$, and this is all we need.)

Since $\mathcal{V}$ is locally finite this means that $\mathbf{T}_n/\theta$ is finite. In particular, θ partitions T_n into finitely many blocks. From each block pick a representative element that is as short as possible (and is a variable if the block contains a variable). Let Σ be the set of all equations of the form

$$Q\, t_0\, t_1 \ldots t_{r_1} \approx s$$

where Q is an operation symbol, r is the rank of Q, the terms $t_0, t_1, \ldots,$ and t_{r_1} are representative terms, and the term s is the representative of the θ-class to which $Qt_0t_1\ldots t_{r-1}$ belongs. Because there are only finitely many operations symbols and only finitely many representative terms, the set Σ of equations is finite. It is easy to check that each equation in Σ is true in $\mathcal{V}$. We can conclude that Σ is a base for $\mathcal{V}^{(n)}$ if we can prove that every equation using only variables from $\{x_0, \ldots, x_{n-1}\}$ that is true in $\mathcal{V}$ is derivable from Σ. This will follow from the next claim.

Claim. Let $p \in T_n$. Then $\Sigma \vdash p \approx q$, where q is the representative of the θ-class to which p belongs.

Proof. We induct on the complexity of p.

The base step of the induction splits into to cases: when p is a variable and when p is a constant symbol. In the first alternative p already a representative term. In the second alternative, $p \approx q$ has been included in Σ.

For the inductive step, we have $p = Qp_0p_1\ldots p_{r-1}$ for some operation symbol Q of postive rank r. By the induction hypothesis we know that

$$\Sigma \vdash p_k \approx q_k$$

for each $k < r$, where q_k is the θ-representative of p_k. So

$$\Sigma \vdash Qp_0p_1\ldots p_{r_1} \approx Qq_0q_1\ldots q_{r-1}.$$

But $Qq_0q_1\ldots q_{r-1} \approx q$ is an equation in Σ. So putting things together, we get

$$\Sigma \vdash p \approx q.$$

This finishes the claim. □

To conclude the proof of the theorem, notice that if $s \approx t$ is an equation with variables only from $\{x_0, \ldots, x_{n-1}\}$ that is true in $\mathcal{V}$, then the θ-representative of s is the same as the θ-representative of t. Say it is q. Then $\Sigma \vdash s \approx q, t \approx q$. So $\Sigma \vdash s \approx t$, as desired. □

Corollary 4.1.1. *Let $\mathcal{V}$ be a locally finite variety of finite signature. Then $\mathcal{V}$ is finitely based if and only if $\mathcal{V} = \mathcal{V}^{(n)}$ for some natural number n.*

As an easy consequence, if $\mathcal{V}$ be a locally finite variety of finite signature and the only basic operations of $\mathcal{V}$ are either of rank 0 or rank 1, then every equation true in $\mathcal{V}$ can have at most two variables. In other words, $\mathcal{V} = \mathcal{V}^{(2)}$ and so $\mathcal{V}$ is finitely based.

While the variety generated by a finite algebra is always locally finite, the variety generated by a locally finite algebra need not be locally finite. You are asked to construct an example in the next problem set. We call an algebra **A uniformly locally finite** provide there is a function $b\colon \omega \to \omega$ on the natural numbers so that every n-generated subalgebra of **A** has cardinality less than $b(n)$ for every natural number n.

Fact. The variety generated by a uniformly locally finite algebra must be locally finite.

You are also asked to prove this fact in the next problem set.

4.2 Lyndon's Example of a Nonfinitely Based Finite Algebra

In 1954 Roger Lyndon published an example of an algebra with seven elements that is not finitely based. His example has two basic operations: a two-place operation and a distinquished element. That he used a constant symbol to denote one element of his algebra was convenient but not essential to his argument.

So Lyndon's seven-element algebra really needed only one basic operation and that one operation is a two-place operation. Some forty years after the publication of Lyndon's result, Zoltan Szekely observed that Lyndon's algebra (as well as several others that arose in the intervening decades) is associated with a finite automaton. Szekely's observation allows us to present Lyndon's example via a diagram.

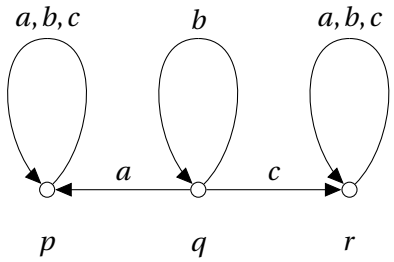

0

Lyndon's Algebra **L**

The elements of the algebra **L** fall into three disjoint sets: the 3-element set $\{p,q,r\}$, which we call the set of states, the 3-element set $\{a,b,c\}$, which we call the set of letters, and the 1-element set $\{0\}$. We call 0 the default element. The binary operation, which we denote by juxtaposition, is defined by following the arrows in the diagram. That is

$$\begin{array}{lll} ap = p & bp = p & cp = p \\ aq = p & bq = q & cq = r \\ ar = r & br = r & cr = r \end{array}$$

with all other products resulting in the default element 0. In this way each letter can be thought to act on the states to produce a state. For instance, the letter a acts on the state q to produce the state p.

The equivalence relation θ on L that makes 0 and r equivalent, but isolates all the other elements is easily seen to be a congruence of **L**. The quotient algebra $\mathbf{L}/\theta$ is isomorphic to $\mathbf{L}^*$, which is displayed below. Evidently, $\mathbf{L}^*$ is in the variety generated by **L**. In 2008, Edmond Lee found $\mathbf{L}^*$ and proved that **L** and $\mathbf{L}^*$ generate the same variety. He also pointed out that $\mathbf{L}^*$ is subdirectly irreducible (the diligent graduate students will check that $(p,0)$ is a critical pair).

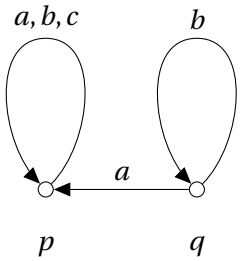

0

Lee's Algebra $\mathbf{L}^*$

Lyndon's Nonfinite Basis Theorem.
The algebra **L** *is not finitely based.*

Proof. Let $\mathcal{V}$ denote the variety generated by **L**. Our plan is to invoke Corollary 4.1.1. To do this, for each (large enough) natural number n we have to devise an algebra $\mathbf{C}_n$ and an equation ϵ_n satisfying the following constraints:

(a) $\mathbf{C}_n \in \mathcal{V}^{(n)}$,

(b) ϵ_n holds in $\mathcal{V}$, and

(c) ϵ_n fails in $\mathbf{C}_n$.

The last two constraints are equivalent to the constraint that $\mathbf{C}_n \notin \mathcal{V}$. Taken together, all three constraints are equivalent to the contention that $\mathcal{V} \neq \mathcal{V}^{(n)}$ for all n.

So fix a natural number n. We first build the algebra $\mathbf{A}_n$, which will be a subalgebra of the n-fold direct power of $\mathbf{L}^*$. So $\mathbf{A}_n \in \mathcal{V}$, since we know that $\mathbf{L}^*$ is a homomorphic image of **L**. So the elements of A_n will be n-tuples of elements of L^*. To avoid the accumulation of notation, we drop the commas and parentheses typical of n-tuples and instead write things like $b \quad b \quad b \quad a \quad c \quad c$, which would be a 6-tuple.

Here are some n-tuples that will belong to A_n:

$$\begin{aligned}
\beta_n &= q \quad q \quad q\cdots q \quad q \quad q \\
\alpha_{n-1} &= b \quad b \quad b\cdots b \quad b \quad a \\
\alpha_{n-2} &= b \quad b \quad b\cdots b \quad a \quad c \\
&\vdots \quad \vdots \\
\alpha_1 &= b \quad a \quad c\cdots c \quad c \quad c \\
\alpha_0 &= a \quad c \quad c\cdots c \quad c \quad c
\end{aligned}$$

We take $\mathbf{A}_n$ to be the subalgebra of $(\mathbf{L}^*)^n$ generated by the $n+1$ elements that were just listed. Of course, A_n has more than just these $n+1$ elements. For example, we see that

$$\alpha_{n-1}\beta_n = q \quad q \quad q\cdots q \quad q \quad p.$$

We will call this new element β_{n-1}. More generally,

$$\alpha_k\beta_{k+1} = \beta_k.$$

This makes, for example, $\beta_1 = q \quad p \quad p \quad \cdots p$ and $\beta_0 = p \quad p \quad p \quad p\cdots p \quad p \quad p$.

Here is another product:

$$\alpha_3\beta_2 = \beta_2$$

and in general

$$\alpha_k\beta_j = \beta_j$$

whenever $j \leq k$. But every other kind of product of these n-tuple produces an n-tuple with a 0 in at least one position. Let θ_n be the equivalence relation that collapses all the n-tuple that have at least one entry 0, but that isolates all the other elements of A_n. It is routine to check that θ_n is a congruence on $\mathbf{A}_n$. Let $\mathbf{B}_n = \mathbf{A}_n/\theta_n$. It is convenient to let $b_k = \beta_k/\theta_n = \{\beta_k\}$ for $k \leq n$ and let $a_k = \{\alpha_k\}$ for $k < n$. We accept the ambiguity of letting 0 denote the one large congruence class. So $\mathbf{B}_n$ belongs to $\mathcal{V}$, it has $2n+2$ elements, and just like **L** and $\mathbf{L}^*$ it can be displayed in a diagram, as follows.

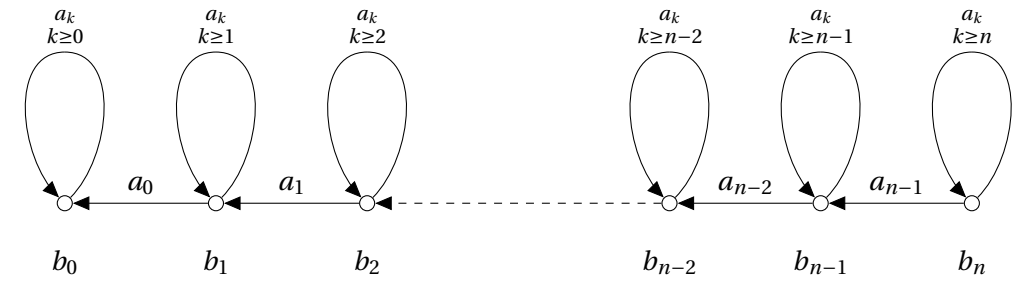

0

The Algebra $\mathbf{B}_n$

Notice that the rightmost loop actually has no label since there are no a_k's when $k \geq n$. The loop was included in the picture for the sake of uniformity.

In $\mathbf{B}_n$ we have $b_n b_0 = 0$. We obtain the desired algebra $\mathbf{C}_n$ by changing this one product—make $b_n b_0 = b_0$ in $\mathbf{C}_n$, but leave everything else unchanged. Here is the diagram for $\mathbf{C}_n$.

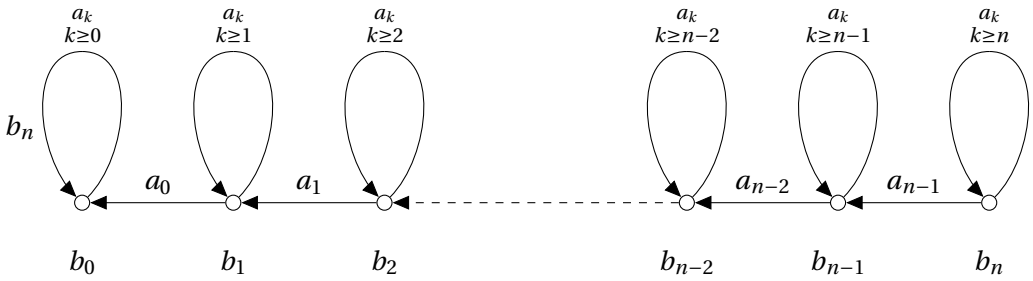

0

The Algebra $\mathbf{C}_n$

The algebras $\mathbf{B}_n$ and $\mathbf{C}_n$ have the same elements and the operations work almost always the same way. If we were displaying these algebras via multiplication tables, our two tables would differ at only one square of the table: in $\mathbf{B}_n$ the entry would be 0 while in $\mathbf{C}_n$ the entry would be b_0.

This small change is enough to keep $\mathbf{C}_n$ out of $\mathcal{V}$. Indeed, it is easy to see that the equation

$$y(x_0(x_1(x_2 \cdots (x_{n-1}y)\cdots))) \approx yy$$

holds in $\mathbf{L}$ but fails in $\mathbf{C}_n$. (Notice yy always evaluates to 0 in both $\mathbf{L}$ and $\mathbf{C}_n$.) So we take the equation displayed above to be ϵ_n.

It only remains to show that $\mathbf{C}_n \in \mathcal{V}^{(n)}$. To achieve this we have to show that every subalgebra of $\mathbf{C}_n$ with n or fewer generators belongs to $\mathcal{V}$. Let $\mathbf{D}$ be such a subalgebra. Notice that the $n+1$ elements

$$a_0, a_1, \ldots, a_{n-1}, \text{ and } b_n$$

are not outputs of the basic operation of $\mathbf{C}_n$ and. moreover, removing any one of these elements results in a subalgebra of $\mathbf{C}_n$. At least one of these elements cannot be in $\mathbf{D}$. It follows that $\mathbf{D}$ is itself a subalgebra of

the subalgebra of $\mathbf{C}_n$ formed by removing that one element. So it suffices to suppose that $\mathbf{D}$ itself arises by removing from $\mathbf{C}_n$ one of the elements listed above.

Suppose first that $\mathbf{D}$ results from removing b_n. Then the diagram below displays $\mathbf{D}$.

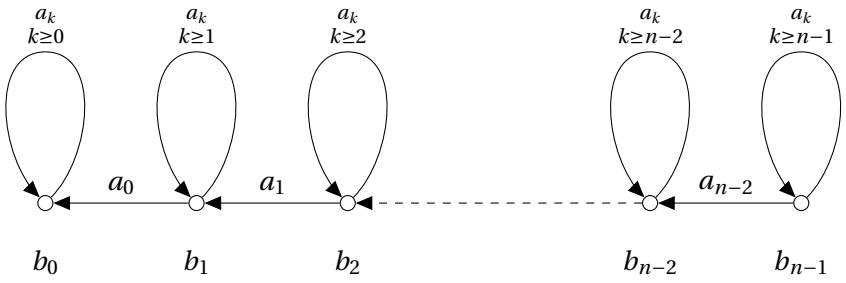

0

The Algebra $\mathbf{D} = \mathbf{C}_n$ with b_n removed.

We see, in this case, that $\mathbf{D}$ is a subalgebra of $\mathbf{B}_n$, so $\mathbf{D} \in \mathcal{V}$ as desired.

The last case we have to consider is that $\mathbf{D}$ arises by removing a_ℓ from $\mathbf{C}_n$. Then the diagram below deplays $\mathbf{D}$.

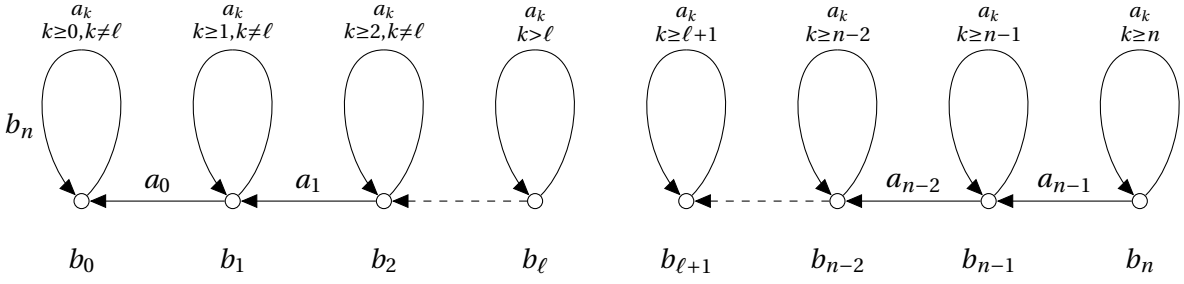

0

The Algebra $\mathbf{D} = \mathbf{C}_n$ with a_ℓ removed.

This diagram allows us to see two subalgebras of $\mathbf{D}$, namely $\mathbf{H}$, which has universe

$$\{0\} \cup \{b_{\ell+1}, \ldots, b_n\} \cup \{a_k \mid \ell < k < n\},$$

and $\mathbf{K}$, which has universe

$$\{0\} \cup \{b_0, \ldots, b_\ell\} \cup \{a_k \mid \ell \neq k < n\}.$$

Evidently, $\mathbf{H}$ is isomorphic to $\mathbf{B}_{n-(\ell+1)}$. This means $\mathbf{H} \in \mathcal{V}$. Now consider the algebra $\mathbf{K}$. Notice that b_n labels an edge but it does not label a node in our diagram. As we want the conclusion that $\mathbf{K} \in \mathcal{V}$, we need to modify our construction of $\mathbf{B}_n$ to add an extra label. Here is how. Let $\mathbf{A}'_n$ be the subalgebra of $(\mathbf{L}^*)^n$ generated by $A_n \cup \{\gamma\}$ where γ is the n-tuple $c \quad c \quad c \cdots c \quad c$. Then let $\mathbf{B}'_n$ be the quotient algebra formed from $\mathbf{A}'_n$ by the congruence that lumps together all n-tuples having at least one entry 0 and isolating all other elements. Let $g = \{\gamma\}$ be the congruence class of γ. The algebra $\mathbf{B}'_n$ has the following diagram.

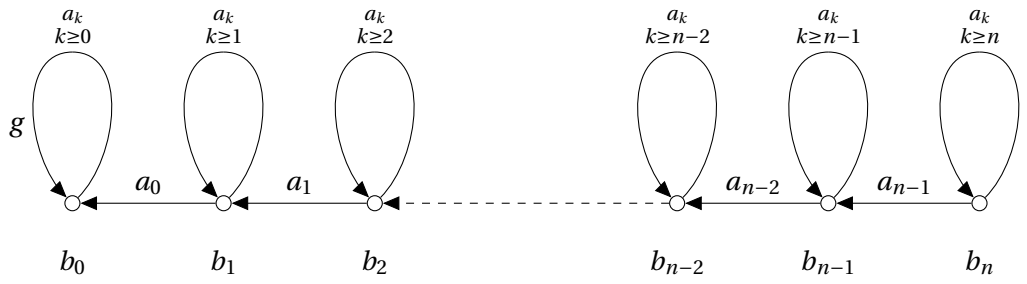

○

0

The Algebra $\mathbf{B}'_n$

So $\mathbf{B}'_n \in \mathcal{V}$ and evidently $\mathbf{K}$ is isomorphic to a subalgebra of $\mathbf{B}'_n$. This means that $\mathbf{K} \in \mathcal{V}$. Further, $\mathbf{K} \times \mathbf{H} \in \mathcal{V}$. But we wanted $\mathbf{D} \in \mathcal{V}$. So the last part of our proof will demonstrate that $\mathbf{D}$ is embeddable into $\mathbf{K} \times \mathbf{H}$. So what we need is two homomorphisms $\varphi : \mathbf{D} \to \mathbf{K}$ and $\psi : \mathbf{D} \to \mathbf{H}$ so that the system $\langle \varphi, \psi \rangle$ separates the points of D. Here are the maps:

$$\varphi(b_k) := \begin{cases} b_k & \text{if } k \leq \ell \\ b_n & \text{if } k = n \\ 0 & \text{Otherwise.} \end{cases}$$

$$\varphi(a_k) := a_k$$

$$\varphi(0) := 0$$

$$\psi(b_k) := \begin{cases} b_k & \text{if } k > \ell \\ 0 & \text{Otherwise.} \end{cases}$$

$$\psi(a_k) := \begin{cases} a_k & \text{If } k > \ell \\ 0 & \text{Otherwise.} \end{cases}$$

$$\psi(0) := 0$$

The work of showing that these two maps are homomorphisms is left for the entertainment of the graduate students. To see the separation of points, observe that that the kernel of φ places all the a_k's, all the b_k's for $k \leq \ell$, and b_n into singleton congruence classes. On the other hand, $\{0\} \cup \{b_k \mid \ell < k < n\}$ is a single congruence class of the kernel of φ. But the kernel of ψ separates all the elements this single big congruence class of kernel of φ. So the system $\langle \varphi, \psi \rangle$ separates the points of D. Consequently, $\mathbf{D}$ is embeddable into $\mathbf{K} \times \mathbf{H}$ and so $\mathbf{D} \in \mathcal{V}$, as desired.

This completes the proof of Lyndon's Nonfinite Basis Theorem.

□

4.3 More Algebraic Preliminaries

Let $\mathbf{A}$ be an algebra and let $\langle \mathbf{B}_i \mid i \in I \rangle$ be a system of algebras of the same signature. Further, let $\langle h_i \mid i \in I \rangle$ be a system of homomorphisms such that

$$h_i : \mathbf{A} \twoheadrightarrow \mathbf{B}_i \quad \text{for each } i \in I.$$

The system $\langle h_i \mid i \in I \rangle$ is a **subdirect representation** of $\mathbf{A}$ provided it **separates points**, that is for all $a, a' \in A$ with $a \neq a'$, there is $i \in I$ such that $h_i(a) \neq h_i(a')$. With such a representation of $\mathbf{A}$ in hand, observe that

$$h : \mathbf{A} \rightarrowtail \prod_{i \in I} \mathbf{B}_i,$$

where the embedding h is defined by

$$h(a) = \langle h_i(a) \mid i \in I \rangle \quad \text{for all } a \in A.$$

Moreover, $h_i = \rho_i \circ h$ for each $i \in I$, where ρ_i is the projection function.

A subdirect representation $\langle h_i \mid i \in I\rangle$ is **trivial** when h_j is one-to-one, for some $j \in I$. After all, that one h_j does all the work of separating points. We say **A** is **subdirectly irreducible** provided every subdirect representation of **A** is trivial.

Let $\theta_i = \ker h_i$. What does this mean? It means $a\,\theta_i\,a'$ if and only if $h_i(a) = h_i(a')$. That is, the homomorphism h_i separates a from a' if and only if a and a' belong to different congruence class modulo $\ker h_i$.

A system of congruences $\langle \theta_i : i \in I\rangle$ **separates points** means that $\bigcap \theta_i$ is the smallest congruence relation 0_A, the congruence that just identifies a with itself for each $a \in A$. That is, $\langle \theta_i : i \in I\rangle$ separates points means that

$$\bigcap_{i\in I} \theta_i = 0_A = \{(a,a) : a \in A\}.$$

So we could reframe the notion of subdirect representation in terms of congruence relations, rather than homomorphisms. A system of congruence relations $\langle \theta_i : i \in I\rangle$ is a **subdirect representation** of **A** provided $\bigcap_{i\in I} \theta_i = 0_A$.

In this setting, an algebra **A** is subdirectly irreducible means that for any subdirect representation $\langle \theta_i \mid i \in I\rangle$ of **A** by congruences, for at least one $i \in I$ we have $\theta_i = 0_A$.

Suppose **A** is a subdirectly irreducible algebra. Consider all the congruences of **A** that aren't trivial. Intersect them:

$$\bigcap\{\theta : \theta \in \mathrm{Con}\mathbf{A} \text{ and } \theta \neq 0_A\}.$$

Notice that this cannot be the smallest congruence 0_A. If it were, then the congruence in our intersection would have to be trivial. But **A** is subdirectly irreducible, so this is impossible. Hence

$$0_A \neq \bigcap\{\theta : \theta \in \mathrm{Con}\mathbf{A} \text{ and } \theta \neq 0_A\}.$$

This this intersection is the unique minimal least nontrivial congruence. We call it the **monolith** of **A**. The congruence lattice of **A** has 1_A at the top, the congruence that relates everything to everything. It has 0_A at the bottom, the identity congruence that just relates a to a and nothing else (for each $a \in A$). Directly above 0_A we find μ, the monolith of **A**. It is the only upper-cover of 0_A. So when **A** is subdirectly irreducible, the congruence lattice **Con A** looks like

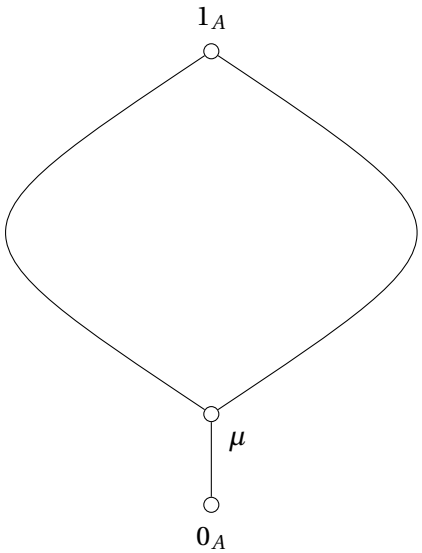

The Congruence Lattice of a Subdirectly Irreducible Algebra **A**

The proof of next theorem is left in the hands of the eager graduate students.

Theorem 4.3.1. *Let* **A** *be an algebra. Then the following are equivalent:*

a. **A** *is subdirectly irreducible;*

b. **A** *has a critical pair;*

c. *The congruence lattice of* **A** *has a monolith;*

d. *Every subdirect representation of* **A** *is trivial.*

Theorem 4.3.2 (Birkhoff's Subdirect Representation Theorem, 1944)**.** *Every algebra is a subdirect product of subdirectly irreducible factors.*

Proof. Let **A** be an algebra. Consider its congruence lattice Con**A**. It has a top element (namely 1_A) and a bottom element (namely 0_A). We want to find a flock of congruences $\langle \theta_i : i \in I \rangle$ that intersect to 0_A and so that $\mathbf{A}/\theta_i$ is subdirectly irreducible.

Let $I = \{(a,b) : a,b \in A \text{ and } a \neq b\}$. Pick $\theta_{(a,b)}$ to be a congruence of **A** maximal with respect to separating a and b. Evidently,

$$\bigcap_{(a,b)\in I} \theta_{(a,b)} = 0_A$$

and

$$\left(a/\theta_{(a,b)}, b/\theta_{(a,b)}\right)$$

is a critical pair of $\mathbf{A}/\theta_{(a,b)}$.

Why we can find a maximal congruence with respect to separating a and b? Ask Zorn.

Why it is that

$$\left(a/\theta_{(a,b)}, b/\theta_{(a,b)}\right)$$

is a critical pair? First, note that

$$a/\theta_{(a,b)} \neq b/\theta_{(a,b)}$$

since $(a,b) \notin \theta_{(a,b)}$. Second, suppose that $\varphi \in \mathrm{Con}\mathbf{A}/\theta_{(a,b)}$ which is nontrivial. Let $\psi \in \mathrm{Con}\mathbf{A}$ with

$$\theta_{(a,b)} \subsetneq \psi \text{ and } \varphi = \psi/\theta_{(a,b)} := \left\{(c/\theta_{(a,b)}, d/\theta_{(a,b)}) : (c,d) \in \psi\right\}$$

Notice that

$$a/\theta_{(a,b)} \varphi b/\theta_{(a,b)}$$

implies $(a,b) \in \psi$. But we know this by the maximality of $\theta_{(a,b)}$. □

Note that **Theorem 4.3.2** is just an existence theorem. It does not tell us how to get a subdirect representation, but just that there is one. Also, it does not say anything about uniqueness. Nor does it give us a description of the subdirectly irreducible algebras. So Birkhoff's Subdirect Representation Theorem falls short of the standard set by the Fundamental Structure Theorem for Finite Abelian Groups.

Corollary 4.3.3. *Every variety of algebras is determined by its subdirectly irreducible members.*

4.4 Inherently Nonfinitely Based Equational Theories

An algebra $\mathbf{A}$ is **inherently nonfinitely based** provided:

(i) $\mathbf{A}$ has only finitely many basic operations;

(ii) $\mathbf{A}$ belongs to some locally finite variety; and

(iii) $\mathbf{A}$ belongs to no locally finite variety that is finitely based.

Similarly, we say that a locally finite variety $\mathcal{V}$ of finite signature is **inherently nonfinitely based** provided it is not included in any finitely based locally finite variety.

Varieties that are inherently nonfinitely based have the nonfinite basis pathology in a contagious way: if $\mathcal{V}$ is inherently nonfinitely based and $\mathcal{W}$ is a locally finite variety with $\mathcal{V} \subseteq \mathcal{W}$, then $\mathcal{W}$ is also inherently nonfinitely based. Similarly if $\mathbf{A}$ is inherently nonfinitely based and $\mathcal{W}$ is a locally finite variety with $\mathbf{A} \in \mathcal{W}$, then $\mathcal{W}$ is inherently nonfinitely based. Finally, if $\mathbf{A}$ is inherently nonfinitely based and $\mathcal{HSP}\,\mathbf{B}$ is locally finite (e.g. $\mathbf{B}$ is finite) and $\mathbf{A} \in \mathcal{HSP}\,\mathbf{B}$, then $\mathbf{B}$ is also inherently nonfinitely based.

From Birkhoff's Finite Basis Theorem we deduce

Theorem 4.4.1. *Let $\mathcal{V}$ be a locally finite variety with a finite signature. Then the following conditions are equivalent:*

(a) *$\mathcal{V}$ is inherently nonfinitely based;*

(b) *The variety $\mathcal{V}^{(n)}$ is not locally finite for any natural number n;*

(c) *For arbitrarily large natural numbers n, there exists a non-locally finite algebra $\mathbf{B}_n$ whose n-generated subalgebras belong to $\mathcal{V}$.*

Thus to show that a locally finite variety $\mathcal{V}$ of finite signature is inherently nonfinitely based, it is enough to construct a family of algebras $\mathbf{B}_n$ (for each $n \in \omega$) so that each $\mathbf{B}_n$ fails to be locally finite, but $\mathbf{B}_n \in \mathcal{V}^{(n)}$.

We say that a variety $\mathcal{W}$ fails to be locally finite **in the finite sense** if there is a natural number p so that $\mathcal{W}$ contains arbitrarily large finite p-generated algebras. For $\mathcal{W}$ to fail to be locally finite, all it needs is an infinite, finitely generated algebra. However, for $\mathcal{W}$ to fail to be locally in the finite sense it must have arbitrarily large (but finite) p-generated algebras. A locally finite variety $\mathcal{V}$ is **inherently nonfinitely based in the finite sense** if whenever $\mathcal{V} \subseteq \mathcal{U}$ where $\mathcal{U}$ is a finitely based variety, then there is a natural number p so that $\mathcal{U}$ has arbitrarily large finite p-generated algebras. Equivalently: A locally finite variety $\mathcal{V}$ of finite signature is **inherently nonfinitely based in the finite sense** if and only if $\mathcal{V}^{(n)}$ fails to be locally finite in the finite sense for all natural numbers n. Combining our ideas thus far, for finite algebras, it is easy to see that

Inherently nonfinitely based in the finite sense
$\Downarrow$
Inherently nonfinitely based
$\Downarrow$
Nonfinitely based

In the opposite direction, it is known that there exist finite algebras that are nonfinitely based but that fail to be inherently nonfinitely based. Lyndon's seven-element algebra is one such example. Whether or not a variety being inherently nonfinitely based is enough to conclude that it is inherently nonfinitely based in the finite sense is not yet known. This question is related to a problem posed in 1976 by Eilenberg and Schützenberger:

The Problem of Eilenberg and Schützenberger:
Let $\mathcal{V}$ be a variety of finite signature generated by a finite algebra. Suppose that $\mathcal{W}$ is a finitely based variety such that $\mathcal{V}$ and $\mathcal{W}$ have exactly the same finite members. Must $\mathcal{V}$ be finitely based?

The same problem can be asked if $\mathcal{V}$ is locally finite, and and that problem is still open as well. It turns out that if there exists a finite, inherently nonfinitely based algebra that fails to be inherently nonfinitely based in the finite sense this algebra would resolve the Problem of Eilenberg and Schützenberger in the negative.

Consider an algebra $\mathbf{A}$ with a finite signature. We call an element $0 \in \mathbf{A}$ an **absorbing element** if any operation evaluated at any tuple containing 0 gives the value 0. The other nonabsorbing elements in $\mathbf{A}$ are called **proper elements**. Note that if it exists, the absorbing element 0 of $\mathbf{A}$ is unique as long as $\mathbf{A}$ has a basic operation of rank at least two. Moreover, if σ is any automorphism of $\mathbf{A}$, the absorbing element will be fixed by σ.

We call any tuple that does not contain the absorbing 0 a **proper tuple** or simply **proper**. If F is a basic operation of $\mathbf{A}$ of rank r, we can regard F as a set of $r+1$-tuples, with the first r entries as the inputs and the last entry as the output. The set of all proper tuples belonging to F is called the **proper part** of F.

Two proper elements of $\mathbf{A}$ are said to be **operationally related** provided they are entries in a tuple of proper elements that belongs to some basic operation of $\mathbf{A}$. An entry in such a tuple will be called an **essential element**.

We define a **shift automorphism algebra** as an infinite, locally finite algebra with only finitely many basic operations, with an absorbing element 0, and with an automorphism σ such that

a. The only finite σ-orbit of $\mathbf{A}$ is $\{0\}$;

b. The proper part of F is partitioned by σ into only finitely many orbits, for each fundamental operation F of $\mathbf{A}$;

c. There is a proper element a of A and a nonconstant unary polynomial function f of $\mathbf{A}$ such that $f(a) = \sigma(a)$.

A **shift automorphism variety** is any variety generated by a shift automorphism algebra. Notice that we have not stipulated that a shift automorphism algebra be uniformly locally finite. Every shift automorphism algebra has a countably infinite subalgebra that is also a shift automorphism algebra.

Almost all of the finite algebras known to be nonfinitely based come out of arguments involving the following Theorem:

The Shift Automorphism Theorem.
Every shift automorphism algebra is inherently nonfinitely based in the finite sense. Furthermore, every shift automorphism variety has a countably infinite subdirectly irreducible member.

Proof. Let $\mathbf{A}$ be a shift automorphism algebra and let $\mathcal{V}$ denote the shift automorphism variety generated by $\mathbf{A}$. Let n be any natural number larger than the rank of any basic operation of $\mathbf{A}$. Let E denote the set of essential elements of $\mathbf{A}$. Let $\mathbf{E}^\sharp$ be the algebra whose universe is $E \cup \{0\}$. We now argue that $\mathbf{E}^\sharp$ is itself a shift automorphism algebra. We begin by showing that the restriction of σ to $E^\sharp$ is an automorphism of $\mathbf{E}^\sharp$.

A set $S \subseteq A$ is **φ-invariant** provided $c \in S$ implies $\varphi(c) \in S$, where φ is an automorphism of $\mathbf{A}$. It is straightforward to check (why not do it?) that $E \cup \{0\}$ is φ-invariant, for every automorphism of $\mathbf{A}$. This means that the restrictions of both σ and σ^{-1} to $E \cup \{0\}$ are automorphisms of $\mathbf{E}^\sharp$. Without fear of ambiguity, we use σ to denote the restriction of σ to $E \cup \{0\}$. It is easy to check that $\mathbf{E}^\sharp$ is a shift-automorphism algebra. But we also note that $\mathbf{E}^\sharp$ is countable and has only finitely many σ-orbits.

We would like to have that $\mathbf{A}$ and $\mathbf{E}^\sharp$ generate the same variety. But this might not be quite true.

Certainly, every equation true in $\mathbf{A}$ must be true in its subalgebra $\mathbf{E}^\sharp$. Consider the converse. Suppose $s \approx t$ fails in $\mathbf{A}$. Then we can pick $c_0, c_1, \cdots \in A$ so that

$$s^{\mathbf{A}}(c_0, c_1, \ldots) \neq t^{\mathbf{A}}(c_0, c_1, \ldots).$$

It is harmless to suppose that $s^{\mathbf{A}}(c_0, c_1, \ldots) \neq 0$. This means that for all i if v_i occurs in s, then $c_i \in E$, provided s is not a variable. By reassigning $c_j = 0$ for all j so that v_j does not occur in s, we obtain a failure $s \approx t$ in $\mathbf{E}^\sharp$. So it remains to consider the case when s is a variable—in fact, $s = v_0$ harmlessly. In this case c_0 might be any proper element of A, even one that is not essential. Again, we could reassign $c_j = 0$ for all $j > 0$ and arrive at a failure of $s \approx t$ in $\mathbf{A}$. But this means that $\mathbf{E}^\sharp$ might not serve our purposes. So let $E^{\sharp\sharp} = E \cup \{0\} \cup \{\sigma^k(d) \mid d \in \mathbb{Z}\}$, where d is some fixed proper element of A that is not essential, if there is one. In this way, reasoning as above, we arrive at the subalgebra $\mathbf{E}^{\sharp\sharp}$ of $\mathbf{A}$ that is a countable shift-automorphism algebra with finitely many σ-orbits, with at most one of these orbits consisting of proper nonessential elements, and that generates the same variety has $\mathbf{A}$.

To simplify notation, we now assume that $\mathbf{A}$ itself is a countable shift-automorphism algebra with finitely many σ-orbits, with at most one of these orbits consisting of proper nonessential elements. We can replace $\mathbf{A}$ by $\mathbf{E}^{\sharp\sharp}$, if we need to.

Let m denote the number of σ-orbits into which A is partitioned. Pick representatives $a_0, a_1, \ldots, a_{m-1}$ from these orbits so that $a_0 = a$ where a is the element mentioned in condition (c). We also take a_{m-1} to be a proper element that is not essential, in case $\mathbf{A}$ has any such elements. Arrange the elements of A in a sequence

$$\ldots, a_{-2}, a_{-1}, a_0, a_1, \ldots, a_{m-1}, a_m, \ldots$$

with the order type of the integers so that $\sigma(a_j) = a_{j+m}$. Given this ordering of A, it is natural to talk about elements being "to the left of" or "to the right of" other elements in A, and the distance between elements in A is defined by the absolute value of the difference of their indices.

By condition (b), up to translation by σ, there can only be finitely many pairs of operationally related elements. Let d be the maximum distance between any two operationally related elements of A. We require that the choices for representatives for the σ-orbits and their ordering has been chosen in such a way that the parameter d is as small as possible. Notice that the parameters d and m depend only on the algebra $\mathbf{A}$ and its automorphism σ.

We now introduce a new parameter. Pick an element a_i from A. Let $X = \{a_j : i \leq j\}$. Notice that X is the set of elements to the right of a_i together with the element a_i itself. Now the subalgebra of $\mathbf{A}$ generated by X may include some elements to the left of a_i. However, any such elements must be generated by the finite set $\{a_j : i \leq j < i + d\}$. Since $\mathbf{A}$ is locally finite by assumption, we can conclude that X can generate only finitely many elements to the left of a_i. Let w_i denote the distance between a_i and the element farthest to the left of a_i which is generated by X. Since σ is an automorphism of $\mathbf{A}$, the only numbers that can arise in this way are precisely $w_0, w_1, \ldots, w_{m-1}$. Let w_L be the largest of these numbers. It is meant to denote the farthest distance to the left that we can get from any a_j using generating elements only to the right of a_j. In the same way, let w_R be defined in an analogous manner: it will be the largest member of an m-sequence of numbers.

Let τ be an automorphism of $\mathbf{A}$. We say that a subalgebra $\mathbf{S}$ of $\mathbf{A}$ is τ-**decomposable** provided there is a subalgebra $\mathbf{S}_0$ (called **core** of the decomposition) such that

$$S = \bigcup_{i \in \mathbb{Z}} \tau^i(S_0)$$

and no element of S_0 is operationally related to any element in any nontrivial τ-translate of S_0. Each τ-translate of $\mathbf{S}_0$ is also called a **core** of the decomposition. Note that the τ-decomposition is constructed so

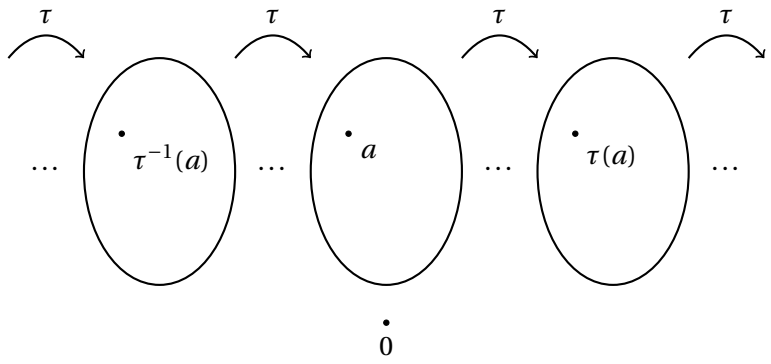

Figure 4.1: The τ-Decomposable Subalgebra

that elements from different blocks cannot be operationally related; that is, if F is any basic operation and $s_0, \ldots, s_{u-1} \in S$ such that

$$F(s_0, \ldots, s_{u-1}) \neq 0$$

we are guaranteed that $\{s_0, \ldots, s_{u-1}\} \subset \tau^j(S_0)$ for some unique integer j. Lastly, notice that every τ-decomposable subalgebra of $\mathbf{A}$ is τ-invariant and it is partitioned into τ-orbits.

Some of the subalgebras of $\mathbf{A}$ might turn out to be τ-decomposable while other might not. We need more control of the decomposable subalgebras. Let ℓ be the smallest natural number such that $\ell m > n(d + w_L + w_R)$. From this point, we reserve τ to denote σ^ℓ. Let $m' = \ell m$. This is the number of τ-orbits into which A is partitioned.

The τ-decomposability Lemma.
The union of any n τ-orbits of A generates a τ-decomposable subalgebra of $\mathbf{A}$ whose cores are n-generated.

Proof. Suppose that $Y \subseteq E$ is the union of n τ-orbits. Let $y_0 \in Y$. Scanning the elements of Y to the right we read the sequence

$$y_0, y_1, \ldots, y_{n-1}, y_n = \tau(y_0)$$

The distance between y_0 and y_n is $m' > n(d + w_L + w_R)$. Between adjacent pairs on this sequence there are n gaps. By the Pigeonhole Principle, at least one of these gaps must have length at least $d + w_L + w_R + 1$. So the set Y can be broken up into pieces of size n so that each is the τ translate of one to the left and the gaps that separate these pieces are of length at least $d + w_L + w_R + 1$. Let Y_0 be the piece that contains y_0. Let $\mathbf{S}_0$ denote the subalgebra generated by Y_0. By construction of w_L and w_R, the subalgebra $\mathbf{S}_0$ can extend no farther than w_L points to the left of Y_0 and no farther than w_R points to the right of Y_0. Define $\mathbf{S}_j = \tau^j(\mathbf{S}_0)$ for each integer j. Between any two such translates of $\mathbf{S}_0$, there must be a gap of length at least $d + 1$. It follows that no element of one translate can be operationally related to any other element of another translate. Furthermore, the subalgebra generated by Y is

$$\bigcup_{j \in \mathbb{Z}} \tau^j(\mathbf{S}_0)$$

and this subalgebra is τ-decomposable. □

The τ-decomposable subalgebra has the structure shown in Figure 4.1. Each core of the decomposition is represented by an ellipse and the element a is the particular element mentioned in condition (c).

For a variety to be inherently nonfinitely based in the finite sense, it first must be locally finite. Since the variety $\mathcal{V}$ is not assumed to be generated by a finite algebra, we must show directly that it is locally finite.

Claim. The variety $\mathcal{V}$ is locally finite.

Proof. It is enough to find a function $b(n)$ on the positive integers so that every n-generated subalgebra of $\mathbf{A}$ has no more than $b(n)$ elements. Let n be as before. Let X be any subset of A so that $|X| = n$. Apart from perhaps the default element 0, and some proper but nonessential elements, the only elements that can be generated from X must be essential elements. Since neither 0 nor any nonessential element can help us generate additional elements, by insisting that $b(n) > n+1$, we can suppose that X contains only essential elements. We now pick n τ-orbits in such a way that each element of X lies in one of the selected orbits. Thus the subalgebra generated by X is included in the subalgebra generated by the orbits of the elements in X. Since we have chosen n-many τ-orbits, the algebra generated by the union of the orbits of elements of X will be τ-decomposable. Let $\mathbf{S}_0$ be a core of this τ-decomposition. Then the subalgebra generated by X will be contained in the union of the τ-images of this core; that is,

$$\mathrm{Sg}^{\mathbf{A}} X \subseteq \bigcup_{i \in \mathbb{Z}} \tau^i(S_0).$$

Since the blocks of the decomposition are operationally unrelated, each of them form a subalgebra. That means that the size of the subalgebra generated by X can be no larger than n times the size of the core S_0. This number still depends on the cardinality of S_0, which is itself an n-generated subalgebra of $\mathbf{A}$. To bound this number, notice that there are only m' many τ-orbits in total. There are $\binom{m'}{n}$ ways to pick n of these orbits. For each such subalgebra we can pick a core algebra for the corresponding τ-decomposable subalgebra. Then n times the size of the largest of these will suffice for $b(n)$. □

Given that $\mathcal{V}$ is locally finite, then $\mathcal{V}$ is inherently nonfinitely based in the finite sense provided $\mathcal{V}^{(n)}$ has arbitrarily large finite p-generated algebras, for some p. Hence our goal is to find a natural number p and finite algebras $\mathbf{B}_{n,k}$ for infinitely many k so that

- Each $\mathbf{B}_{n,k}$ has at least k elements;
- Each $\mathbf{B}_{n,k}$ is generated by a set of p elements; and
- Each $\mathbf{B}_{n,k}$ is in $\mathcal{V}^{(n)}$.

We will now construct the algebras $\mathbf{B}_{n,k}$.

Recall that, apart from the elements in at most one of the orbits, all nonzero elements of A are assumed to be essential. We consider a partial subalgebra of the $2k$-fold direct power of this algebra. Let e be an essential element and let U be a subset of $\{0,1,\ldots,2k-1\}$. By $e|U$ we mean the $2k$-tuple that has e at the j^{th} position for all $j \in U$ and 0 at all other positions.

Let μ be the map from the set $\mathbb{Z}$ of integers to $\{0,1,\ldots,2k-1\}$ defined by

$$\mu(z) \equiv z \pmod{2k} \text{ for all } z \in \mathbb{Z}.$$

We say that the subset $U \subseteq \{0,1,\ldots,2k-1\}$ is **contiguous modulo** $2k$ provided it is the image under the map μ of an interval in $\mathbb{Z}$. Define G as the set of $2k$-tuples $e|U$ for which

- e is an essential element;
- U is a nonempty contiguous set of integers modulo $2k$; and
- $|U| < k$.

The set G is the universe of a partial subalgebra $\mathbf{G}$. We use $\mathbf{G}^\sharp$ to denote the subalgebra of $\mathbf{A}^{2k}$ obtained by adjoining the constantly 0 tuple to $\mathbf{G}$.

We now provide an example of such an $e|U$. Rather than considering $e|U$ as a usual $2k$-tuple, we will consider it as having a circular form where the positions are as they appear on an analog clock. Let $k = 4$, let $U = \{0, 1\}$, and let e be an essential element of $\mathbf{A}$. Then $e|U$ has the following form:

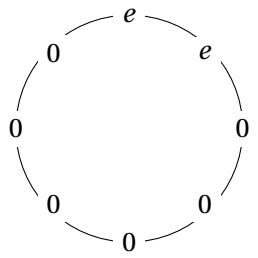

Notice that this depiction gives rise to two automorphisms in a natural way. The automorphism τ of $\mathbf{A}$ extends naturally to a coordinatewise defined automorphism on $\mathbf{G}^\sharp$; namely, we can define τ on $\mathbf{G}^\sharp$ as follows:

$$\tau(e|U) := \tau(e)|U$$

Another automorphism becomes clear using the circular depiction of $e|U$. Let ρ denote the cyclic shift on the coordinate set U. We note that the automorphisms τ and ρ commute. Furthermore, by applying $\tau \circ \rho$ to the particular $e|U$ that appears above, we find

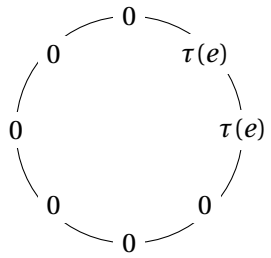

and so $\tau \circ \rho\,(e|U) = \tau(e)|\rho(U)$. Notice that the collection of U's that play a role in G is closed with respect to the cyclic shift.

Claim. $\tau \circ \rho$ partitions the partial algebra $\mathbf{G}$ into $2k^2m'$ orbits and these orbits are congruence classes of some congruence relation of $\mathbf{G}$.

Proof. As we noted before, there are m' many τ-orbits of A. The number of $\tau \circ \rho$-orbits of $\mathbf{G}$ will be the product of m' and the number of possible choices for the cyclically contiguous nonempty sets U with cardinality no more than k. To define a particular U, it is enough to describe the cardinality of U and the starting point of the contiguous interval (modulo $2k$). There are $2k$ choices for the starting point of the contiguous interval (modulo $2k$) of U. Given a starting point of this interval, there are k choices for its length. Thus there are $2k \cdot k = 2k^2$ possible sets U. Hence the total number of $\tau \circ \rho$-orbits in $\mathbf{G}$ is $2k^2 \cdot m'$, as claimed.

Since $\tau \circ \rho$ partitions $\mathbf{G}$, it forms an equivalence relation of $\tau \circ \rho$-equivalence classes. We just need to show that the underlying equivalence relation is a congruence. Suppose Q is a basic operation symbol of rank r

and that

$$\begin{aligned} Q^{\mathbf{G}}\left(g_0|U_0,\ldots,g_{r-1}|U_{r-1}\right) &= g_r|U_r \\ Q^{\mathbf{G}}\left(h_0|V_0,\ldots,h_{r-1}|V_{r-1}\right) &= h_r|V_r \\ \left(\tau\circ\rho\right)^{e_j} g_j|U_j &= h_j|V_j \text{ for } j<r. \end{aligned}$$

We need to verify that $g_r|U_r$ and $h_r|V_r$ lie in the same $\tau\circ\rho$-orbit.

Given that 0 is an absorbing element, the result of $Q^{\mathbf{G}}\left(g_0|U_0,\ldots,g_{r-1}|U_{r-1}\right)$ must contain 0 at each coordinate in the complement of $U_0\cap\cdots\cap U_{r-1}$. That is,

$$U_0\cap U_1\cap\cdots\cap U_{r-1}=U_r.$$

Furthermore, $Q^{\mathbf{A}}\left(g_0,\ldots,g_{r-1}\right)=g_r$. Likewise,

$$\begin{aligned} V_0\cap V_1\cap\cdots\cap V_{r-1}&=V_r \text{ and} \\ Q^{\mathbf{A}}\left(h_0,\ldots,h_{r-1}\right)&=h_r. \end{aligned}$$

Since $\tau\circ\rho$ acts in a coordinatewise manner, we have

$$\begin{aligned} \tau^{e_j}\left(g_j\right)&=h_j \text{ for all } j<r, \text{ and} \\ \rho^{e_j}\left(U_j\right)&=V_j \text{ for all } j<r \end{aligned}$$

Provided that $e_0=e_1=\cdots=e_{r-1}$, we will have our desired conclusion.

Having chosen n so that it is greater than the rank of any basic operation of $\mathbf{A}$, the elements $g_0,\ldots,g_{r-1}$ cannot be in more than n τ-orbits. Since each g_j is τ-related to each h_j (where $j<r$), the elements $h_0,\ldots,h_{r-1}$ are in those same τ-orbits. All these essential elements must belong to a single block of the τ-decomposable subalgebra. Since $g_0,\ldots,g_{r-1}$ are operationally related, they must belong to a single τ-translate of the core. Likewise, since $h_0,\ldots,h_{r-1}$ are operationally related, they must belong to a single τ-translate of the core. Therefore, there is a single e so that τ^e carries the block containing the g_j's to the block containing the h_j's. Hence $e=e_0=e_1=\cdots=e_{r-1}$. □

We now know that the $\tau\circ\rho$-orbits of $\mathbf{G}$ are really congruence classes of some congruence relation of $\mathbf{G}$. Let γ denote that congruence of $\mathbf{G}$. Define $\mathbf{G}_{n,k}$ as the algebra made by adding the default absorbing element to $\mathbf{G}/\gamma$. Since $\mathbf{G}/\gamma$ has $2k^2m'$ elements by the claim above, the algebra $\mathbf{G}_{n,k}$ has cardinality $2k^2m'+1$.

Claim. $\mathbf{G}_{n,k} \in \mathcal{V}^{(n)}$.

Proof. Consider n proper elements of $\mathbf{G}/\gamma$. These correspond to n many $\tau \circ \rho$-orbits of $\mathbf{G}$. From these, we can obtain n many τ-orbits of A since each $\tau \circ \rho$-orbit has an associated τ-orbit. Let $\mathbf{D}$ be the subalgebra generated by the union of these orbits. Since $\mathbf{D}$ is n-generated, it is τ-decomposable by τ-decomposability Lemma. Let $\mathbf{S}$ be a core of this decomposition. Now we can pick elements $g_0, \ldots, g_{n-1}$ of S and appropriate nonempty subsets $U_0, \ldots, U_{n-1}$ of $\{0, 1, \ldots, 2k-1\}$ so that

$$g_0|U_0, \ldots, g_{n-1}|U_{n-1}$$

is a system of representatives of our original $\tau \circ \rho$ orbits. Let $\mathbf{H}$ be the subalgebra of $\mathbf{G}^\sharp$ generated by

$$\{g_0|U_0, \ldots, g_{n-1}|U_{n-1}\}.$$

Let π be the map from H to $G_{n,k}$ defined so that each proper element of H is assigned the $\tau \circ \rho$-orbit to which it belongs and the absorbing element of $\mathbf{H}$ is assigned the absorbing element of $\mathbf{G}_{n,k}$. We aim to show that π is a homomorphism from $\mathbf{H}$ into $\mathbf{G}_{n,k}$.

Let Q be a basic operation symbol of rank r and let $h_0|V_0, \ldots, h_{r-1}|V_{r-1}$ be proper elements of H. Since the operation $Q^{\mathbf{H}}$ is defined coordinatewise, we have

$$\pi\left(Q^{\mathbf{H}}(h_0|V_0, \ldots, h_{r-1}|V_{r-1})\right) = \pi\left(Q^{\mathbf{A}}(h_0, \ldots, h_{r-1})\,|\,V_0 \cap \cdots \cap V_{r-1}\right). \tag{4.1}$$

Since each $h_j|V_j$ is a proper element (where $j < r$), the map π sends each $h_j|V_j$ to $h_j|V_j/\gamma$. Thus

$$Q^{\mathbf{G}_{n,k}}(\pi(h_0|V_0), \ldots, \pi(h_{r-1}|V_{r-1})) = Q^{\mathbf{G}_{n,k}}\left(h_0|V_0/\gamma, \ldots, h_{r-1}|V_{r-1}/\gamma\right). \tag{4.2}$$

We aim to show that equation (4.1) and equation (4.2) are equal. If the right side of equation (4.2) is the default element, the equivalence is clear. If it is not the default element, then there must be representatives $w_0|W_0, \ldots, w_{r-1}|W_{r-1}$ of the γ-classes such that

$$Q^{\mathbf{G}}(w_0|W_0, \ldots, w_{r-1}|W_{r-1})$$

is defined in $\mathbf{G}$. In this case, the τ-decomposability of $\mathbf{D}$ entails that

$$Q^{\mathbf{G}}(h_0|V_0, \ldots, h_{r-1}|V_{r-1})$$

is defined as well. Given that this is defined, all four sides of the above equations are equivalent, and hence π is a homomorphism.

The original n elements of $\mathbf{G}_{n,k}$ all belong to the image of the homomorphism π, and so the subalgebra they generate belongs to $\mathcal{V}$ as $\mathbf{H}$ belongs to $\mathcal{V}$. □

The algebra $\mathbf{G}_{n,k}$ has more than k elements and is found in $\mathcal{V}^{(n)}$; however, its generating set is not of the right cardinality. We will now define p and the algebra $\mathbf{B}_{n,k}$.

Condition (c) above provides a unary polynomial f and an essential element a so that $f(a) = \sigma(a)$. We need such a polynomial that will do the same with τ. Let $t(x, y_1, \ldots, y_r)$ be a term in which x occurs and let $c_1, \ldots, c_r$ be essential elements so that

$$f(x) = t^{\mathbf{A}}(x, c_1, \ldots, c_r)$$

which we will denote by $t^{\mathbf{A}}(x, \bar{c})$. The polynomial that works with τ and a is $g(x)$ displayed below.

$$g(x) = t^{\mathbf{A}}\left(t^{\mathbf{A}}\left(\cdots t^{\mathbf{A}}(x, \bar{c}), \sigma^{\ell-2}(c_1), \ldots, \sigma^{\ell-2}(c_r)\right), \sigma^{\ell-1}(c_1), \ldots, \sigma^{\ell-1}(c_r)\right)$$

Notice $g(a) = \tau(a)$. Let $s(x, y_1, \ldots, y_k)$ be a term in which x occurs and let $d_1, \ldots, d_j$ be essential elements so that

$$g(x) = s^{\mathbf{A}}\left(x, d_1, \ldots, d_j\right) \tag{4.3}$$

which we will denote by $s^{\mathbf{A}}(x, \bar{d})$.

First, observe that x cannot be the only variable to occur in s, for otherwise, $\tau(a) = s^{\mathbf{A}}(a)$ entails $\tau^2(a) = s^{\mathbf{A}}(\tau(a)) = s^{\mathbf{A}}(s^{\mathbf{A}}(a))$. More generally, we would have

$$t^k(a) = s^{\mathbf{A}}(s^{\mathbf{A}}(\ldots s^{\mathbf{A}}(a)\ldots)),$$

for every natural number k. This would entail that each $\tau^k(a)$ is in the subalgebra of $\mathbf{A}$ generated by $\{a\}$. But $\mathbf{A}$ is locally finite, so this is this subalgebra is finite. This is impossible since the τ-orbit of a is infinite.

Now $s(x, y_0, \ldots)$ must contain a shortest subtern w such that both x and at least one of the variables $y_0, \ldots$ occur in w. This means that $w = Qu_0u_1\ldots u_{r-1}$ where some the u_i's are terms in which no variable other than x occurs (and in at least one of these x does occur) and in the remaining u_j's the variable x does not occur (and in at least one of these some variable from $y_0, \ldots$ does occur). Here Q must be an operation symbol of rank at least two.

In equation 4.3, we constructed a polynomial $g(x)$ so that $g(a) = \tau(a)$ where a is the element mentioned in condition (c). We also found a term $s(x, y_1, \ldots, y_j)$ and essential elements $d_1, \ldots, d_j$ of $\mathbf{A}$ so that

$$g(x) = s^{\mathbf{A}}\left(x, d_1, \ldots, d_j\right).$$

Let $p = j + 1$. Notice p is one greater than the number of "coefficients" of $g(x)$. Furthermore, notice that $p = \ell r + 1$, where r is the number of "coefficients" in the original polynomial $f(x)$. Let $Y_j = \{0, \ldots, i-1\}$ for each $i < k$. Let Z_i be the preimage of Y_i under the cyclic shift ρ. That is, Z_i is Y_i turned one notch counter-clockwise; explicitly,

$$Z_i = \{2k-1, 0, 1, \ldots, i-2\}.$$

Let $\mathbf{B}_{n,k}$ be the subalgebra of $\mathbf{G}_{n,k}$ generated by the following elements:

$$a|Y_k/\gamma, d_1|Y_k/\gamma, \ldots, d_{p-1}|Y_k/\gamma$$

where each d_i (for $i < p$) are those found in equation (4.3).

As $\mathbf{B}_{n,k}$ is generated by p elements and is found in the variety $\mathcal{V}^{(n)}$, the only thing left to show is that $\mathbf{B}_{n,k}$ has at least k elements. We claim that the following k elements

$$a|Z_1/\gamma, a|Z_2/\gamma, \ldots, a|Z_k/\gamma$$

are distinct and all belong to $\mathbf{B}_{n,k}$.

Notice

$$\begin{aligned} s^{\mathbf{G}}\left(a|Y_k, d_1|Y_k, \ldots, d_{p-1}|Y_k\right) &= s^{\mathbf{A}}\left(a, d_1, \ldots, d_{p-1}\right)|Y_k \cap \cdots \cap Y_k \\ &= s^{\mathbf{A}}\left(a, d_1, \ldots, d_{p-1}\right)|Y_k \\ &= \tau(a)|Y_k. \end{aligned}$$

Now, $\tau(a)|Y_k$ is γ-related to $a|Z_k$ since the preimage of $\tau(a)$ under τ is a and the preimage of Z_k under ρ is Y_k. Using a similar argument, we see

$$\begin{aligned} s^{\mathbf{G}}\left(a|Z_k, d_1|Y_k, \ldots, d_{p-1}|Y_k\right) &= s^{\mathbf{A}}\left(a, d_1, \ldots, d_{p-1}\right)|Z_k \cap Y_k \cap \cdots \cap Y_k \\ &= s^{\mathbf{A}}\left(a, d_1, \ldots, d_{p-1}\right)|Y_{k-1} \\ &= \tau(a)|Y_{k-1} \\ &\gamma\ a|Z_{k-1} \end{aligned}$$

and

$$\begin{aligned} s^{\mathbf{G}}\left(a|Z_{k-1}, d_1|Y_k, \ldots, d_{p-1}|Y_k\right) &= s^{\mathbf{A}}\left(a, d_1, \ldots, d_{p-1}\right)|Z_{k-1} \cap Y_k \cap \cdots \cap Y_k \\ &= s^{\mathbf{A}}\left(a, d_1, \ldots, d_{p-1}\right)|Y_{k-2} \\ &= \tau(a)|Y_{k-2} \\ &\gamma\ a|Z_{k-2}. \end{aligned}$$

We can use this same argument to verify that

$$a|Z_k/\gamma, a|Z_{k-1}/\gamma, \ldots, a|Z_1/\gamma$$

are all distinct elements. Since they all belong to $\mathbf{B}_{n,k}$, it has more than k elements, as desired.

At this stage, we know that $\mathcal{V}$ is inherently nonfinitely based in the finite sense.

Now we turn to the construction of an infinite subdirectly irreducible algebra in $\mathcal{V}$. To show that $\mathcal{V}$ has an infinite subdirectly irreducible algebra, we let θ be a maximal congruence of $\mathbf{A}$ separating a and 0. Then $\mathbf{A}/\theta$ will be subdirectly irreducible. To see that it is infinite we argue that θ separates $\tau^{-j}(a)$ and $\tau^{-q}(a)$, whenever j and q are distinct positive integers.

Notice

$$\begin{array}{c} a = s^{\mathbf{A}}\left(\tau^{-1}(a), \tau^{-1}(\bar{d})\right) \\ \Downarrow \\ \tau^{-1}(a) = s^{\mathbf{A}}\left(\tau^{-2}(a), \tau^{-2}(\bar{d})\right) \end{array}$$

Combining these, we see

$$a = s^{\mathbf{A}}\left(s^{\mathbf{A}}\left(\tau^{-2}(a), \tau^{-2}(\bar{d})\right), \tau^{-1}(\bar{d})\right)$$

In general, we obtain

$$a = s^{\mathbf{A}}\left(\ldots(s^{\mathbf{A}}(\tau^{-j}(a), \tau^{-j}(\bar{d}))\ldots\tau^{-1}(\bar{d})\right).$$

Claim. Let j and q be natural numbers, with $j < q$ and let φ be any congruence of $\mathbf{A}$. Then

$$\tau^{-j}(a)\ \varphi\ \tau^{-q}(a) \implies a\ \varphi\ 0.$$

Proof. As we saw above,

$$a = s^{\mathbf{A}}\left(\ldots(s^{\mathbf{A}}(\tau^{-j}(a), \tau^{-j}(\bar{d}))\ldots\tau^{-1}(\bar{d})\right).$$

Consider the deepest piece: $s^{\mathbf{A}}(\tau^{-j}(a), \tau^{-j}(\bar{d}))$. Recall the subterm $w = Qu_0\ldots u_{r-1}$ of the term $s(x, y_1, \ldots, y_k)$. We make the harmless supposition that $u_0, \ldots u_{c-1}$ contain no variables other than x and $u_c, \ldots, u_{r-1}$ have variables drawn only from $y_0, \ldots$. Let

$$e_c = u_c^{\mathbf{A}}(\tau^{-j}(\bar{d})), \ldots, e_{r-1} = u_{r-1}^{\mathbf{A}}(\tau^{-j}(\bar{d})).$$

Likewise, let

$$e_0 = u_0^{\mathbf{A}}(\tau^{-j}(a)), \ldots, e_{c-1} = u_{c-1}^{\mathbf{A}}(\tau^{-j}(a)).$$

Because n is at least as large as the rank of any fundamental operation, we see that $e_0, \ldots, e_c, \ldots, e_{r-1}$ account for no more than n elements. These elements are all operationally related by Q. The subalgebra of $\mathbf{A}$ generated by the union of the τ-orbits of these elements is τ-decomposable. Since the $e_0\ldots, e_{r-1}$ are operationally related, they must belong to the same core of the τ-decomposition. But recall that $\tau^{-j}(a)$ belongs to the core S_{-j}. Since S_{-j} is a subuniverse. So $e_0, \ldots, e_{c-1}$ must belong to S_{-j}. It follows that $e_0, \ldots, e_{r-1}$ all belong to S_{-j}.

But our hypothesis is that $\tau^{-j}(a)\ \varphi\ \tau^{-q}(a)$. So we see that

$$Q^{\mathbf{A}}(u_0^{\mathbf{A}}(\tau^{-j}(a)),\ldots,u_{c-1}^{\mathbf{A}}(\tau^{-j}(a)),e_c,\ldots,e_{r-1})\ \varphi\ Q^{\mathbf{A}}(u_0^{\mathbf{A}}(\tau^{-q}(a)),\ldots,u_{c-1}^{\mathbf{A}}(\tau^{-q}(a)),e_c,\ldots,e_{r-1})$$

But on the right side of this congruence we see that the first c entries belong to S_{-q} and so are operationally unrelated to the last $r-c$ entries. This means that right side evaluates to 0. Taken altogether, we have

$$a = s^{\mathbf{A}}\Big(\ldots(s^{\mathbf{A}}(\tau^{-j}(a)),\tau^{-j}(\bar{d}))\ldots\tau^{-1}(\bar{d})\Big)\ \varphi\ s^{\mathbf{A}}\Big(\ldots(s^{\mathbf{A}}(\tau^{-q}(a)),\tau^{-j}(\bar{d}))\ldots\tau^{-1}(\bar{d})\Big) = s^{\mathbf{A}}\big(\ldots(0)\ldots\tau^{-1}(\bar{d})\big) = 0.$$

This is the desired result. □

We could reframe the last claim to assert that any congruence φ of $\mathbf{A}$ that separates a and 0 must also separate all the elements of the form $\tau^{-j}(a)$ where j is any natural number. Invoking Zorn as needed, take θ to be a maximal congruence separating a and 0. Then $\mathbf{A}/\theta$ is a subdirectly irreducible algebra is $\mathcal{V}$ and it is infinite since there must be infinitely many congruence classes to accomodate all those $\tau^{-j}(a)$'s.

As a and 0 are not related by θ, we have shown that

$$\Big(\tau^{-j}(a),\tau^{-q}(a)\Big)\notin\theta$$

for any distinct natural numbers j and q. The quotient algebra $\mathbf{A}/\theta$ will have countably infinitely many elements and it is subdirectly irreducible by the maximality of θ with respect to separating a and 0. Furthermore, this algebra is in $\mathcal{V}$.

Taking everything into account, we have shown that $\mathcal{V}$ is a locally finite variety, that $\mathcal{V}^{(n)}$ fails to be locally finite in the finite sense, for all large enough values of n (so $\mathcal{V}$ is inherently nonfinitely based in the finite sense), and that $\mathcal{V}$ contains an infinite subdirectly irreducible algebra. This finishes the proof of the Shift Automorphism Theorem. □

The Shift Automorphism Theorem asserts two conclusions beyond the central conclusion that shift automorphism algebras are inherently nonfinitely based. The first conclusion, that they are inherently nonfinitely based in the finite sense, means that the Shift Automorphism Theorem cannot be used to directly construct a counterexample to the Eilenberg-Schützenberger Conjecture. The second conclusion, that shift automorphism varieties must always have infinite subdirectly irreducible members, means that this theorem cannot be used directly not construct a counterexample to

> **Jónsson's Speculation**
> Is every finite algebra of finite signature, which generates a variety with the finite residual bound, actually finitely based?

A variety is said to have residual bound κ provided every subdirectly irreducible algebra in the variety has cardinality less than κ. Bjarni Jónsson offered this speculation in the mid-1970's, noting finite algebras that generate varieties so that the congruence lattice of each algebra in the variety is distributive have finite residual bounds. In 1971, Kirby Baker had proven that such finite algebras of finite signature are finitely based.

4.5 Examples of Inherently Nonfinitely Based Finite Algebras

Our first example is like Lyndon's algebra. It is the automatic algebra $\mathbf{R}$ associated with the automaton displayed below:

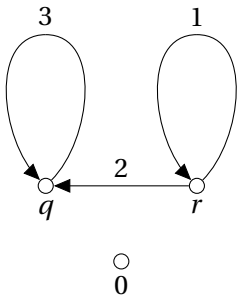

McKenzie's algebra **R**

Of course, **R** is finite, so we cannot apply the Shift Automorphism Theorem to it directly. Instead, we will construct an algebra $\mathbf{A} \in \mathcal{HSP}\,\mathbf{R}$ to which the Shift Automorphism Theorem does apply. Here it is.

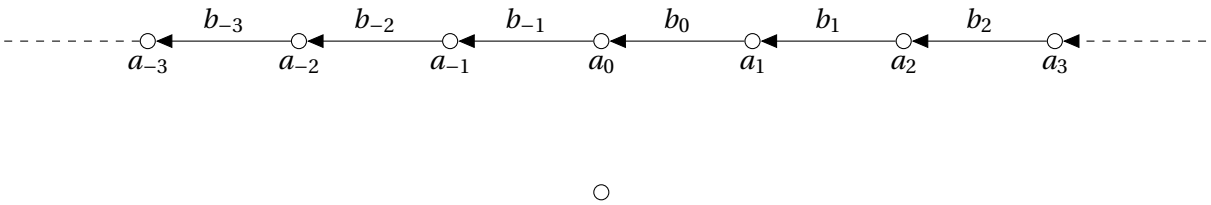

The algebra **A**

Here, we define the product $\cdot$ as follows:

$$b_k \cdot a_{k+1} = a_k$$

and all other products are 0. The automorphism σ that we need works as follows

$$\sigma(a_{k+1}) = a_k \qquad \sigma(b_{k+1}) = b_k \qquad \sigma(0) = 0.$$

So there are two infinite orbits and one finite orbit, namely $\{0\}$. The proper triples have the form (b_k, a_{k+1}, a_k) and σ induces just one orbit on these triples. Finally, $\sigma(a_1) = b_0 \cdot a_1 = f(a_1)$ where $f(x) = b_0 \cdot x$. So **A** is a shift automorphism algebra.

Fact. $\mathbf{A} \in \mathcal{HSP}\,\mathbf{R}$.

Proof. In $\mathbf{R}^{\mathbb{Z}}$, let B be the set of all $\mathbb{Z}$-tuples of the following form: For each $k \in \mathbb{Z}$, define α_k and β_k as follows:

$$\begin{array}{lccccccccc} \alpha_k := & \ldots & r & r & r & q & q & q & q & \ldots \\ \beta_k := & \ldots & 1 & 1 & 1 & 2 & 3 & 3 & 3 & \ldots \end{array}$$

The k is meant to designate the rightmost occurrence of r in α_k, and likewise the (only) occurrence of 2 in β_k. Notice that $\beta_k \cdot \alpha_{k+1} = \alpha_k$. No matter how we multiply some $\beta_n \cdot \alpha_m$, we'll either get some α_{m-1} or we'll get a string that contains a 0. Now let B consist of all the α_k's and all the β_k's as well as any member of $A^{\mathbb{Z}}$ that has at least one entry that is 0.

The point is that B is a subuniverse of $\mathbf{R}^{\mathbb{Z}}$. Let θ be the equivalence relation on B that lumps together all $\mathbb{Z}$-tuples containing 0 and isolates everything else. That is, θ is going to isolate each α_i with itself and it will isolate each β_j with itself.

We contend that θ is a congruence of the algebra B.

Evidently, θ an equivalence relation. It's also respects the operation. Just consider

$$\begin{array}{ccc} \delta & \theta & \gamma \\ \mu & \theta & \nu \\ \hline \mu\cdot\delta & \theta & \nu\cdot\gamma \end{array}$$

No matter how we pick δ, γ, μ, and ν in B, this will always work out. This is because our big lump will absorb everything containing 0.

But we have arranged matters so that $\mathbf{A} \cong \mathbf{B}/\theta$. Hence, $\mathbf{A} \in \mathcal{HSP}\,\mathbf{R}$, as desired. □

It follows that $\mathbf{R}$ is inherently nonfinitely based.

Our second example has only 3 elements. It is made from the graph below.

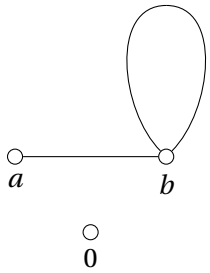

Murskiĭ's algebra $\mathbf{M}$

In graph algebras like $\mathbf{M}$ the universe consists of the vertices of the graph together with a default element 0 and the basic two place operation works are follows:

$$u\cdot v = \begin{cases} u & \text{if } u \text{ is adjacent to } v \\ 0 & \text{otherwise} \end{cases}$$

Once more, we construct a shift automorphism algebra $\mathbf{A} \in \mathcal{HSP}\,\mathbf{M}$. Here it is

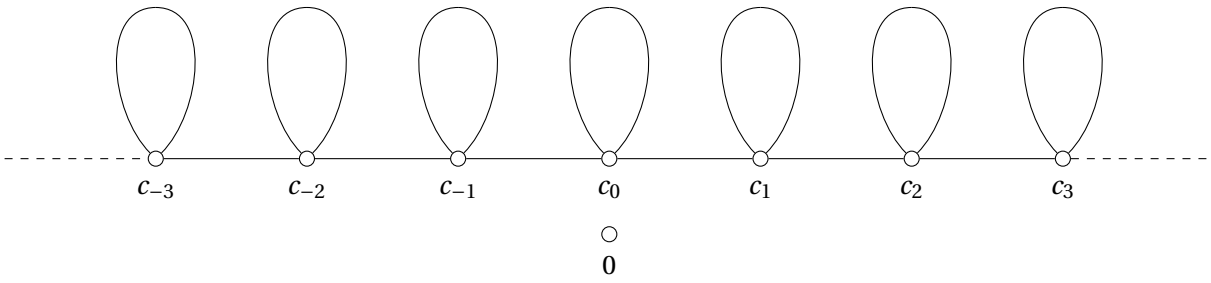

The Shift Automorphism Algebra $\mathbf{A}$

In this algebra, the set $A = \{c_k \mid k \in \mathbb{Z}\} \cup \{0\}$ and the operation is defined as it was in $\mathbf{M}$:

$$u\cdot v = \begin{cases} u & \text{if } u \text{ is adjacent to } v \\ 0 & \text{otherwise} \end{cases}$$

The automorphism σ works so that $\sigma(c_{k+1}) = c_k$ for all $k \in \mathbb{Z}$ and $\sigma(0) = 0$. Consequently, there are just two orbits: $\{c_k \mid k \in \mathbb{Z}\}$ and $\{0\}$. Finally, $\sigma(c_1) = c_0 = c_0 \cdot c_1 = f(c_1)$, where $f(x) = c_0 \cdot x$. So $\mathbf{A}$ is a shift automorphism algebra.

As in the first example, it remains to show that $\mathbf{A} \in \mathcal{HSP}\,\mathbf{M}$. We will make short work of this. In $M^{\mathbb{Z}}$, for each integer k,let

$$\gamma_k := \quad \ldots \quad b \quad b \quad b \quad a \quad b \quad b \quad a \quad b \quad b \quad b \quad a \quad b \quad \ldots$$

where the leftmost a occurs at position k. Observe that γ_k has an infinite string of b's running to the left, while to the right the occurrences of a are separated by longer and longer strings of b's. Now note

$$\begin{array}{rcccccccccccccc}
\gamma_{k+1} = & \ldots & b & b & b & a & b & b & a & b & b & b & a & b & \ldots \\
\gamma_k = & \ldots & b & b & a & b & a & b & b & a & b & b & b & a & \ldots \\
\hline
\gamma_k \cdot \gamma_{k+1} = \gamma_k = & \ldots & b & b & a & b & a & b & b & a & b & b & b & a & \ldots
\end{array}$$

and

$$\begin{array}{rcccccccccccccc}
\gamma_k = & \ldots & b & b & b & a & b & a & b & b & a & b & b & b & \ldots \\
\gamma_{k+1} = & \ldots & b & b & b & b & a & b & a & b & b & a & b & b & \ldots \\
\hline
\gamma_{k+1} \cdot \gamma_k = \gamma_{k+1} = & \ldots & b & b & b & b & a & b & a & b & b & a & b & b & \ldots
\end{array}$$

On the other hand, any other product of the γ_k's will produce a $\mathbb{Z}$-tuple with at least one entry that is 0. The rest of the argument proceeds as in the first example.

So Murskiĭ's algebra **M** is inherently nonfinitely based.

These two examples explain why we have called *shift* automorphism algebras what we did. The automorphisms given here are indeed shifts.

4.6 Problem Set 5

Problem Set Nonfinitely Based Finite Algebras

Problem 40.
Construct a locally finite algebra **A** so that the variety $\mathcal{HSP}\,\mathbf{A}$ is not locally finite.

Problem 41.
Prove that if **A** is uniformly locally finite, then $\mathcal{HSP}\,\mathbf{A}$ is locally finite.

Problem 42.
Prove that Lyndon's algebra **L** and Lee's algebra $\mathbf{L}^*$ generate the same variety.

Problem 43.
Add a couple of elements to Lyndon's Algebra **L** to obtain a finite algebra **B** so that **L** is a subalgebra of **B** and **B** is finitely based. In this way, establish that even though **L** fails to be finitely based, it also fails to be inherently nonfinitely based.

Problem 44.
Prove that the four-element algebra **P**, with just one basic operation and that atwo-place operation whose operation, table is displayed below is inherently nonfinitely based

·	0	a	b	c
0	0	0	0	0
a	0	a	b	0
b	0	b	b	c
c	0	0	c	c

LESSON 5

EQUATIONAL THEORIES THAT ARE FINITELY AXIOMATIZABLE

5.1 EVERY FINITE LATTICE IS FINITELY BASED

Ralph McKenzie proved, around 1967, that every finite lattice is finitely based **McKenzie1970** Subsequently, this theorem was extended, first by Kirby Baker **Baker1977** to finite algebras that belong to congruence distributive varieties, then by McKenzie **McKenzie1987** to finite algebras that belong to congruence modular varieties with a finite residual bound, and then by Ross Willard **2000** to finite algebras that belong to congruence meet-semidistributive varieites with a finite residual bound. And there have been alternative proofs to many of these results as well as various extensions to quasivarieties. The proofs of all these generalizations have a heavily algebraic character. In particular, they all depend on an analysis of subdirectly irreducible algebras in some way. It is striking, then, in retrospect, that McKenzie's original proof is completely syntactic. So it seems to me reasonable to take another look at McKenzie's original argument with an eye toward seeing if the argument itself can suggest a fresh more syntactic approach to finite basis theorems. What follows is, to all intents and purposes McKenzie's original published argument, although I have taken the liberty to frame it a little differently.

For every natural number n let $\ell(n) := n^m + 1$, where m is the smallest integer bigger than 1 so that $n < 2^{m+1}$. For example, $\ell(2) = 2^2 + 1 = 5$, $\ell(12) = 12^3 + 1$ and $\ell(120) = 120^6 + 1$. The function $\ell(n)$ is, roughly, $2^{(\log_2(n))^2} = n^{\log_2(n)}$ and it gets above and stays above any polynomial.

McKenzie's Finite Basis Theorem For Finite Lattices.
Every finite lattice **L** *of cardinality n has an equational basis using no more than* $\ell(n)$ *variables.*

Proof. We let **L** be a lattice of cardinality n and we put $\ell := \ell(n)$. We also take the number m described during the definition of ℓ as fixed.

Let $\mathcal{V}$ denote the variety generated by **L**. Our aim is to show that $\mathcal{V} = \mathcal{V}^{(\ell)}$.

The variety of all lattices has a base using just 3 variables. Since we want all the algebras we have to contend with to be lattices we insist that that n be large enough so that $3 \leq \ell(n)$.

Observe, $\ell(1) = 2$. So the proof below won't apply to lattices of cardinalities 1. But the trivial lattice is based on $x \approx y$, an equation of $2 = \ell(1)$ variables, as required by the theorem.

Let X be any finite set of variables. Below, we will define a (normal form) function η_X on terms whose variables are drawn from X that has the following properties for all terms s and t whose variables lie in X. In what follows, X can always be understood from the context so we use η in place of η_X.

(A) If $s \approx t$ is true in **L**, then $\eta(s) \approx \eta(t)$ is true in $\mathcal{V}^{(\ell)}$, and

(B) $s \approx \eta(s)$ is true in $\mathcal{V}^{(\ell)}$.

Once such normal form functions are in hand, the proof of the theorem will be complete.

Let Y be a set of variables, let A be any finite set, and let $f : Y \to A$. If s is a term whose variables are all drawn from Y we let $\mu_f(s)$ be the result of the substitution $y \mapsto \bigwedge_{f(z)=f(y)} z$, for all $y \in Y$. We also let $\mu_f^o(s)$ be the term resulting from s by the substitution $y \mapsto y_f$, where y_f is the variable in Y of least index such that $f(y_f) = f(y)$.

The normal form function we want is defined by

$$\eta(s) := \bowtie_{f:X\to L} (\mu_f(s)).$$

For any terms s and t whose variables all come from Y and for all $f : Y \to A$ all of the following hold.

(i) $\mu_f(s \wedge t) = \mu_f(s) \wedge \mu_f(t)$.

(ii) $\mu_f(s \vee t) = \mu_f(s) \vee \mu_f(t)$.

(iii) $\mu_f(s) = \mu_f(\mu_f^o(s))$.

The first two items follow since substitutions are endomorphisms of the algebra of terms, while the third follows directly from the definitions.

The first of the two things we have to verify, namely (A), is in reach. Suppose s and t are terms whose variables are drawn for X and that $s \approx t$ is true in $\mathbf{L}$. We need to prove that $\eta(s) \approx \eta(t)$ holds in every lattice in $\mathcal{V}^{(\ell)}$. Let $f : X \to L$. From $s \approx t$ we derive $\mu_f^o(s) \approx \mu_f^o(t)$. This is just a substitution instance of $s \approx t$ and it as at most n variables. Since $n \leq \ell$, we see that $\mu_f^o(s) \approx \mu_f^o(t)$ is true in $\mathcal{V}^{(\ell)}$. Any substitution instance of this equation must also be true in $\mathcal{V}^{(\ell)}$. In particular,

$$\mu_f(\mu_f^o(s)) \approx \mu_f(\mu_f^o(t)) \text{ is true in } \mathcal{V}^{(\ell)}.$$

But we saw above that $\mu_f(\mu_f^o(r)) = \mu_f(r)$ for any term r. So we have

$$\mu_f(s) \approx \mu_f(t) \text{ is true in } \mathcal{V}^{(\ell)}.$$

Now just form the joins of both sides as f runs through all the functions from X in L, to conclude that the equation below is true in $\mathcal{V}^{(\ell)}$.

$$\eta(s) \; = \bowtie_{f:X\to L} \mu_f(s) \; \approx \bowtie_{f:X\to L} \mu_f t) = \eta(t)$$

So it only remains to show (B), that is

$$s \approx \eta(s) \text{ is true in } \mathcal{V}^{(\ell)}.$$

Observe that $\mu_f(s) \leq s$ holds in every lattice. since the lattice operations are monotone. It follows that $\eta(s) \leq s$ in every lattice, as well. So we only need to show that $s \leq \eta(s)$ holds in $\mathcal{V}^{(\ell)}$. We prove this by induction on the complexity of s. The base step (when s is just a variable) is trivial. The induction step falls into two parts depending on whether $s = u \vee v$ or $s = u \wedge v$. The case with $\vee$ presents no difficulties. So we are left with establishing

$$\text{If } s = u \wedge v \text{ and both } u \leq \eta(u) \text{ and } v \leq \eta(v) \text{ hold in } \mathcal{V}^{(\ell)}, \text{ then } s \leq \eta(s) \text{ holds in } \mathcal{V}^{(\ell)}. \qquad (\star)$$

Before tackling ($\star$) we develop some further properties of the kind of substitutions defined above.

Fact. Let $f : X \to A$ and $g : X \to B$ so that $f(x) = f(y) \implies g(x) = g(y)$ for all $x, y \in X$. Then $\mu_g(s) \leq \mu_f(s)$ holds in all lattices for all terms s with variables drawn from X.

This is a simple consequence of the monotonicity of the lattice operations.

Fact. Suppose $f : Y \to A$. Let $Y_f := \{y_f \mid y \in Y\}$. Let $g : Y_f \to B$. Let $h : Y \to B$ be defined via $h(y) := g(y_f)$. The equation $\mu_h(t) \approx \mu_f(\mu_g(\mu_f^o(t)))$ is true in all lattices, for any term t with variables from Y.

Proof. We prove this by induction on the complexity of t.

The base step is that t is a variable, say x. Then

$$\begin{aligned}
\mu_f(\mu_g(\mu_f^o(x))) &= \mu_f(\mu_g(x_f)) \\
&= \mu_f(\bigwedge_{g(y_f)=g(x_f)} y_f) \\
&= \bigwedge_{g(y_f)=g(x_f)} \mu_f(y_f) \\
&= \bigwedge_{g(y_f)=g(x_f)} \bigwedge_{f(z)=f(y_f)} z \\
&= \bigwedge_{g(y_f)=g(x_f)} \bigwedge_{z_f=y_f} z \\
&= \bigwedge_{g(y_f)=g(x_f)} y \\
&= \bigwedge_{h(y)=h(x)} y \\
&= \mu_h(x).
\end{aligned}$$

The part of lattice theory that enters here is just the commutative and associative laws for $\wedge$.

The inductive step is immediate, since μ_f, μ_g, and μ_f^o are all endomorphisms of the term algebra. □

Key Lemma.
Let $g : X \to L^m$. Let s be any term with variables drawn from X and let y be a variable not belonging to X. The following inequalities all hold in $\mathcal{V}^{(\ell)}$.

(a) $y \wedge \mu_g(s) \leq \bowtie_{f:X\to L} (y \wedge \mu_f(s))$.

(b) $\mu_g(s) \leq \eta(s)$.

(c) $y \wedge \eta(s) \leq \bowtie_{f:X\to L} (y \wedge \mu_f(s))$.

Proof. Let $X' = \{x_g \mid x \in X\}$. Let t be any terms with variables drawn from X'. Notice that the inequality

$$y \wedge t \leq \bowtie_{h:X'\to L} (y \wedge \mu_h(t))$$

has no more than $n^m + 1 = \ell$ distinct variables. We will show that this inequality holds in $\mathbf{L}$ and therefore also in $\mathcal{V}^{(\ell)}$.

Let $\bar{a}$ be any X'-tuple of elements of L and let $b \in L$. In the left side of our inequality, plug b in for y and $\bar{a}$ in for the variables in t. The effect of μ_g^o is to identify certain variables.

Now the maps h appearing on the right side area assignments of elements of L to the variables in X'. One of these assignments h is precisely the assignment $\bar{a}$. This means that $h(x) = c \in A$ exactly when the entry on the tuple $\bar{a}$ associated with x is c. Hence, when $\bigwedge_{h(y)=h(x)} y$ is evaluated under the assignment $\bar{a}$ the

result will be $c \wedge c \wedge c \wedge \ldots \wedge c = c$. Consequently, the particular joinand on the right associated with this h, when evaluated at $\bar{a}$ is actually the value of the left side at $\bar{a}$. This verifies the inequality in **L**.

Now let $t = \mu_g^o(s)$. So

$$y \wedge \mu_g^o(s) \leq \bowtie_{h:X' \to L} (y \wedge \mu_h(\mu_g^o(s))$$

holds in $\mathcal{V}^{(\ell)}$. Now apply μ_g to both sides:

$$\begin{aligned}
\mu_g(y \wedge \mu_g^o(s)) &\leq \mu_g\left(\bowtie_{h:X' \to L} (y \wedge \mu_h(\mu_g^o(s)))\right) \\
\mu_g(y) \wedge \mu_g(\mu_g^o(s)) &\leq \bowtie_{h:X' \to L} (\mu_g(y) \wedge (\mu_g(\mu_h(\mu_g^o(s))))) \\
y \wedge \mu_g(s) &\leq \bowtie_{h:X' \to L} (y \wedge \mu_g(\mu_h(\mu_g^o(s)))) \\
y \wedge \mu_g(s) &\leq \bowtie_{h:X' \to L} (y \wedge \mu_{h^*}(s))
\end{aligned}$$

where $h^*(y) = h(y_g)$ for all $y \in X$, according to the Fact preceding the statement of the Key Lemma. Since

$$\bowtie_{h:X' \to L} (y \wedge \mu_{h^*}(s)) \leq \bowtie_{f:X \to L} (y \wedge \mu_f(s))$$

we are finished with (a).

Part (b) is an immediate consequence of part (a), obtained by setting y to the join of all the $\mu_f(s)$'s and $\mu_g(s)$.

To establish (c) we need a bit more groundwork. First consider the following equation:

$$y \wedge \bowtie_{i<m+1} x_i \approx \bowtie_{f:m \to m+1} \left(y \wedge \bowtie_{j<m} x_{f(j)}\right). \tag{$*$}$$

This equation has $m+2$ variables. Some simple calculations show that $m+2 \leq \ell$. To see that this equation is true in $\mathcal{V}^{(\ell)}$ we only need to see that it is true in **L**.

Let $A \subseteq L$ have $m+1$ elements. So A has 2^{m+1} subsets. Since $n < 2^{m+1}$ we see that two distinct subsets of A must have the same joins. It follows that A has a proper subset A' so that A and A' have the same joins. This means that under any assignment of members of L to the variables in $(*)$ the left side of the equation will be one of the joinands on the right side. On the other hand, each joinand on the right is dominated by the left side. Hence, the equation $(*)$ holds in **L** and hence in $\mathcal{V}^{(\ell)}$. Actually, from $(*)$, with the help of lattice theory, we can deduce

$$y \wedge \bowtie_{i<k} x_i \approx \bowtie_{f:m \to k} \left(y \wedge \bowtie_{j<m} x_{f(j)}\right). \tag{$**$}$$

for every $k > m$. So $(**)$ holds in $\mathcal{V}^{(\ell)}$ for every $k > m$.

Here is what we have to prove to finish this lemma.

$$y \wedge \bowtie_{f:X \to L} (\mu_f(s)) \leq \bowtie_{f:X \to L} (y \wedge \mu_f(s)) \tag{c}$$

holds in $\mathcal{V}^{(\ell)}$. Now the join on the left has $n^{|X|}$ joinands. It is easy to see that $n > m$, so we can use the equation $(**)$ to change (c) to

$$\bowtie_{h:m \to K} \left(y \wedge \bowtie_{j<m} \mu_{h(j)}(s)\right) \leq \bowtie_{f:X \to L} (y \wedge \mu_f(s)) \tag{c'}$$

where K is the set of functions from X into L. This reduces our task to demonstrating that each joinand on the left is dominated by the right side. So let $f_0, \ldots, f_{m-1}$ be functions from X to L. Let $g : X \to L^m$ be the function defined via

$$g(x) := (f_0(x), \ldots, f_{m-1}(x)).$$

Observe that $\mu_{f_i}(s) \leq \mu_g(s)$ holds in all lattices for all $i < m$, by a fact we proved above. So we get

$$y \wedge \bowtie_{i<m} \mu_{f_i}(s) \leq y \wedge \mu_g(s) \leq \bowtie_{f:X \to L} y \wedge \mu_f(s).$$

The rightmost inequality follows from part (a). This concludes the demonstration of part (c). □

With the help of the Key Lemma, we can complete the inductive step to establish (B) and finish the proof of the Theorem.

Recall that we were left with establishing

$$\text{If } s = u \wedge v \text{ and both } u \leq \eta(u) \text{ and } v \leq \eta(v) \text{ hold in } \mathcal{V}^{(\ell)}, \text{ then } s \leq \eta(s) \text{ holds in } \mathcal{V}^{(\ell)}. \qquad (\star)$$

Observe

$$\begin{aligned} s = u \wedge v \leq \eta(u) \wedge \eta(v) &\leq \bowtie_{f:X\to L} \eta(u) \wedge \mu_f(v) \\ &\leq \bowtie_{f,h:X\to L} \mu_h(u) \wedge \mu_f(v) \end{aligned}$$

by two applications of part (c) of the Key Lemma. So we need to show that

$$\mu_h(u) \wedge \mu_f(v) \leq \eta(s)$$

holds in $\mathcal{V}^{(\ell)}$ for all $f, h : X \to L$. Let $g : X \to A^m$ be the function defined via

$$g(x) := (h(x), f(x), f(x), \ldots, f(x)).$$

Here is where we need that $2 \leq m$. By the monotonicity fact we have that $\mu_f(w) \leq \mu_g(w)$ and $\mu_h(w) \leq \mu_g(w)$ for all terms w with variables from X.

Finally we see

$$\mu_h(u) \wedge \mu_f(v) \leq \mu_g(u) \wedge \mu_g(v) = \mu_g(u \wedge v) = \mu_g(s) \leq \eta(s),$$

where the rightmost inequality is part (b) of the Key Lemma. This is what we needed. □

5.2 Finite lalltice-ordered algebras

We will say that **A** is a **lattice-ordered algebra** provided

(a) **A** as among its basic operations two binary operations so that, reduced to just these two operations **A** is a lattice.

(b) Every basic operation of **A** is monotone with respect to the lattice order established by (a).

As McKenzie noted, his theorem applies to lattice-ordered algebras.

McKenzie's Theorem for Finite Lattice-Ordered Algebras.
Every finite lattice-ordered algebra of finite signature is finitely based.

Proof. The idea is simply to reprise the proof above, but in the inductive argument for (B) to replace the $\wedge$ by an arbitrary basic operation symbol Q. Were Q of rank 2 no particular, apart from replacing $\wedge$ by Q, would have to change. If, however, the rank of Q is larger then we are forced to change the value of m and hence that of ℓ. We need to require that m is an upper bound on the rank of any operation of our algebra.

Here is how the Key Lemma would go with Q being a 3-place operation symbol.

Key Lemma, take II.
Let $g : X \to L^m$. Let s and t be any terms with variables drawn from X and let y be a variable not belonging to X. The following inequalities all hold in $\mathcal{V}^{(\ell)}$.

(a) $Q(y, \mu_g(s), \mu_g(t)) \leq \bowtie_{f:X\to L} (Q(y, \mu_f(s), \mu_f(t))$.

(b) $\mu_g(s) \leq \eta(s)$ *and* $\mu_g(t) \leq \eta(t)$.

(c) $Q(y, \eta(s), \eta(t)) \leq \bowtie_{f:X\to L} (Q(y, \mu_f(s), \mu_f(t))$.

Proof. Let $X' = \{x_g \mid x \in X\}$. Let s and t be any terms with variables drawn from X'. Notice that the inequality

$$Q(y, s, t) \leq \bowtie_{h:X'\to L} (Q(y, \mu_h(s), \mu_h(t)))$$

has no more than $n^m + 1 = \ell$ distinct variables. We will show that this inequality holds in **L** and therefore also in $\mathcal{V}^{(\ell)}$.

Let $\bar{a}$ be any X'-tuple of elements of L and let $b \in L$. In the left side of our inequality, plug b in for y and $\bar{a}$ in for the variables in s and t. The effect of μ_g^o is to identify certain variables.

Now the maps h appearing on the right side area assignments of elements of L to the variables in X'. One of these assignments h is precisely the assignment $\bar{a}$. This means that $h(x) = c \in A$ exactly when the entry on the tuple $\bar{a}$ associated with x is c. Hence, when $\bigwedge_{h(y)=h(x)} y$ is evaluated under the assignment $\bar{a}$ the result will be $c \wedge c \wedge c \wedge \ldots \wedge c = c$. Consequently, the particular joinand on the right associated with this h, when evaluated at $\bar{a}$ is actually the value of the left side at $\bar{a}$. This verifies the inequality in **L**.

Now let $t = \mu_g^o(s)$. So

$$y \wedge \mu_g^o(s) \leq \bowtie_{h:X'\to L} (y \wedge \mu_h(\mu_g^o(s))$$

holds in $\mathcal{V}^{(\ell)}$. Now apply μ_g to both sides:

$$\begin{aligned}
\mu_g(y \wedge \mu_g^o(s)) &\leq \mu_g\left(\bowtie_{h:X'\to L} (y \wedge \mu_h(\mu_g^o(s)))\right)\\
\mu_g(y) \wedge \mu_g(\mu_g^o(s)) &\leq \bowtie_{h:X'\to L} (\mu_g(y) \wedge (\mu_g(\mu_h(\mu_g^o(s))))\\
y \wedge \mu_g(s) &\leq \bowtie_{h:X'\to L} (y \wedge \mu_g(\mu_h(\mu_g^o(s))))\\
y \wedge \mu_g(s) &\leq \bowtie_{h:X'\to L} (y \wedge \mu_{h^*}(s))
\end{aligned}$$

where $h^*(y) = h(y_g)$ for all $y \in X$, according to the Fact preceding the statement of the Key Lemma. Since

$$\bowtie_{h:X'\to L} (y \wedge \mu_{h^*}(s)) \leq \bowtie_{f:X\to L} (y \wedge \mu_f(s))$$

we are finished with (a).

Part (b) is an immediate consequence of part (a), obtained by setting y to the join of all the $\mu_f(s)$'s and $\mu_g(s)$.

To establish (c) we need a bit more groundwork. First consider the following equation:

$$y \wedge \bowtie_{i<m+1} x_i \approx \bowtie_{f:m\to m+1} \left(y \wedge \bowtie_{j<m} x_{f(j)}\right). \tag{$*$}$$

This equation has $m+2$ variables. Some simple calculations show that $m+2 \leq \ell$. To see that this equation is true in $\mathcal{V}^{(\ell)}$ we only need to see that it is true in **L**.

Let $A \subseteq L$ have $m+1$ elements. So A has 2^{m+1} subsets. Since $n < 2^{m+1}$ we see that two distinct subsets of A must have the same joins. It follows that A has a proper subset A' so that A and A' have the same joins. This means that under any assignment of members of L to the variables in $(*)$ the left side of the equation will be one of the joinands on the right side. On the other hand, each joinand on the right is dominated by the left side. Hence, the equation $(*)$ holds in **L** and hence in $\mathcal{V}^{(\ell)}$. Actually, from $(*)$, with the help of lattice theory, we can deduce

$$y \wedge \bowtie_{i<k} x_i \approx \bowtie_{f:m\to k} \left(y \wedge \bowtie_{j<m} x_{f(j)}\right). \tag{$**$}$$

for every $k > m$. So $(**)$ holds in $\mathcal{V}^{(\ell)}$ for every $k > m$.

Here is what we have to prove to finish this lemma.

$$y \wedge \bowtie_{f:X\to L} (\mu_f(s)) \leq \bowtie_{f:X\to L} (y \wedge \mu_f(s)) \tag{c}$$

holds in $\mathcal{V}^{(\ell)}$. Now the join on the left has $n^{|X|}$ joinands. It is easy to see that $n > m$, so we can use the equation $(**)$ to change (c) to

$$\bowtie_{h:m\to K} \left(y \wedge \bowtie_{j<m} \mu_{h(j)}(s)\right) \leq \bowtie_{f:X\to L} (y \wedge \mu_f(s)) \tag{c'}$$

where K is the set of functions from X into L. This reduces our task to demonstrating that each joinand on the left is dominated by the right side. So let $f_0, \ldots, f_{m-1}$ be functions from X to L. Let $g : X \to L^m$ be the function defined via

$$g(x) := (f_0(x), \ldots, f_{m-1}(x)).$$

Observe that $\mu_{f_i}(s) \leq \mu_g(s)$ holds in all lattices for all $i < m$, by a fact we proved above. So we get

$$y \wedge \bowtie_{i<m} \mu_{f_i}(s) \leq y \wedge \mu_g(s) \leq \bowtie_{f:X\to L} y \wedge \mu_f(s).$$

The rightmost inequality follows from part (a). This concludes the demonstration of part (c). □

□

5.3 Even More Algebraic Preliminaries

Let $\mathbf{A}$ be an algebra and $X \subseteq A \times A$. This collection $\mathcal{F}$ of congruences of $\mathbf{A}$ that include the set X is nonempty, since the largest congruence 1_A is the set $A \times A$. This makes $\bigcap \mathcal{F}$ the smallest congruence that includes X. We denote this congruence by $\mathrm{Cg}^{\mathbf{A}} X$ and call it the **congruence generated** by X. Of course, we would like to replace this shrink-wrap definition by a characterization of more constructive character.

By a **basic translation** of $\mathbf{A}$ we mean a function $\lambda : A \to A$ so that there is a basic operation symbol Q of positive rank r, elements $a_0, a_1, \ldots, a_{r-1} \in A$, and some $j < r$ so that

$$\lambda(a) = Q^{\mathbf{A}}(a_0, \ldots, a_{j-1}, a, a_{j+1}, \ldots, a_{r-1}) \quad \text{for all } a \in A.$$

That is, a basic translation is a polynomial obtained by taking a basic operation of $\mathbf{A}$ of positive rank, picking one input position of the basic operation as the input of λ, and then plugging in elements of A in the other positions to play the role of coefficients. By a **translation** of $\mathbf{A}$ we mean a function that is the composition of some finite number of basic translations. We take the composition of zero basic translations to be the identity function on A. A given translation might arise in many ways as a composition of basic translations. We say a translation λ is a k**-translation** provided λ is a composition of k or fewer basic translations. We will call the smallest such k the **complexity** of λ. It is easy to check that if a, b, c, and d belong to A and λ is any translation of $\mathbf{A}$ such that $\{a, b\} = \{\lambda(c), \lambda(d)\}$, then $(a, b) \in \mathrm{Cg}^{\mathbf{A}}(c, d)$. (Try induction on the complexity of λ....) The following theorem was proven by A. I. Mal'cev in 1954.

The Congruence Generation Theorem.
Let $\mathbf{A}$ be an algebra and a, b, c, and d be elements of A. Then $\langle a, b\rangle \in Cg^{\mathbf{A}}(c, d)$ if and only if there is a natural number n and elements $a_0, a_1, \ldots, a_n \in A$ and translations $\lambda_0, \lambda_1, \ldots, \lambda_{n-1}$ of $\mathbf{A}$ so that $a = a_0$ and $a_n = b$ and

$$\{a_i, a_{i+1}\} = \{\lambda_i(c), \lambda_i(d)\} \quad \textit{for all } i < n.$$

Here is a diagram of the condition in this theorem.

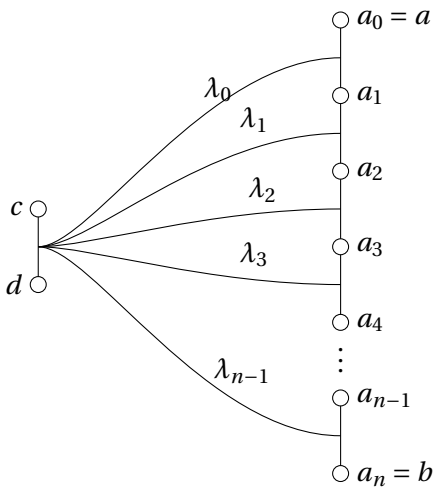

Proof. The condition given implies that $\langle a, b\rangle \in \mathrm{Cg}^{\mathbf{A}}(c, d)$ since each links is in the congruence and congruences are transitive relations.

For the converse, let

$$\theta = \{(u, v) \mid \text{ for some } n, \text{ some } a_0, a_1, \ldots, a_n \in A, \text{ and some translations } \lambda_0, \lambda_1, \ldots, \lambda_{n-1} \text{ so that } u = a_0 \text{ and } a_n = v \text{ and } \{a_i, a_{i+1}\} = \{\lambda_i(c), \lambda_i(d)\} \quad \text{for all } i < n.\}.$$

Taking $n = 1$ and λ_0 to be the identity function, we see that $\langle c, d\rangle \in \theta$. Also only a little work is required to see that $\theta \subseteq \mathrm{Cg}^{\mathbf{A}}(c, d)$. So we only need to proved that θ is a congruence relation. That θ is an equivalence relation is straightforward. So consider an operation symbol Q. To simplify notation, we suppose that the rank of Q is 3. So let u_0, u_0', u_1, u_1', and u_2, u_2' be elements of A such that

$$u_0 \,\theta\, u_0'$$
$$u_1 \,\theta\, u_1'$$
$$u_2 \,\theta\, u_1'$$

We need the conlcusion that $\qquad Q^{\mathbf{A}}(u_0, u_1, u_2) \,\theta\, Q^{\mathbf{A}}(u_0', u_1', u_2').$

The idea is the same one we used in the proof of the Completeness Theorem for Equational Logic. Let $a_0, \ldots, a_n$ be elements of A and $\lambda_0, \ldots, \lambda_{n-1}$ be translations witnessing that $u_0 \,\theta\, u_0'$. For each $i \leq n$ let $a_i^* := Q^{\mathbf{A}}(a_i, u_1, v_1)$ and let $\lambda^*(x) := Q^{\mathbf{A}}(\lambda_i(x), u_i, u_2)$. Then the elements $a_0^*, \ldots, a_n^*$ and the translations $\lambda_0^*, \ldots, \lambda_{n-1}^*$ witness that $Q^{\mathbf{A}}(u_0, u_1, u_2) \,\theta\, Q^{\mathbf{A}}(u_0', u_1, u_2)$. Now repeat this process for u_1 and, after that, for u_2. By concatenating the results, we arrive at a sequence of elements and a sequence of translations that witness $Q^{\mathbf{A}}(u_0, u_1, u_2) \,\theta\, Q^{\mathbf{A}}(u_0', u_1', u_2')$, as desired. □

While we have framed the Congruence Generation Theorem of Mal'cev just for principal congruences, it can be easily adapted two use a set X of ordered pairs of elements of A in place of the single ordered pair (c, d). This Congruence Generation Theorem was, in fact, the inspiration for our system of inference in equational logic, since Birkhoff had already shown that the equational theories where exactly the congruence relations of $\mathbf{T}^*$, the term algebra expanded by taking each endomorphism as a new one-place operation.

As a consequence of this theorem, we see that the congruences of $\mathbf{A}$ are just those equivalence relations that respect all the translations of $\mathbf{A}$. Let $\mathrm{Tr}\mathbf{A}$ denote the set of all translations of $\mathbf{A}$. Let $\mathbf{A}^* = \langle A, \lambda\rangle_{\lambda \in \mathrm{Tr}\mathbf{A}}$. Then $\mathrm{Con}\mathbf{A} = \mathrm{Con}\mathbf{A}^*$.

We use $\{c, d\} \hookrightarrow_{\ell}^{n} \{a, b\}$ to denote that there is a sequence $a_0, \ldots, a_n$ of $n+1$ elements of A and a sequence $\lambda_0, \ldots, \lambda_{n-1}$ of n translations of $\mathbf{A}$, each of complexity no more than ℓ, that witnesses $(a, b) \in \mathrm{Cg}^{\mathbf{A}}(c, d)$.

In case the sginature is finite, for variables x, y, z, and w, we can regard $\{x, y\} \not\hookrightarrow^n_\ell \{z, w\}$ as an elementary formula (a formula of first order logic) with free variables x, y, z, and w that has a string of existential quantifiers (here asserting the existence of $n+1$ elements, as well as the existence of lots of coefficients for the translations) followed by a disjuntion of a conjnction of a lot of equations.

5.4 Willard's Finite Basis Theorem

In 2000 Ross Willard proved the following theorem.

Willard's Finite Basis Theorem.
Every congruence meet-semidistributive variety of finite signature with a finite residual bound is finitely based.

Recall that a lattice is meet-semidistributive provided the following implication holds in the lattice:

$$x \wedge y \approx x \wedge z \Rightarrow x \wedge (y \vee z) \approx (x \wedge y) \vee (x \wedge z).$$

This is a weakening of the distributive property. A variety is congruence meet-semidistributive provided $\mathrm{Con}\, A$ is a meet-semidistributive lattice for each algebra $\mathbf{A}$ in the variety. There are several different characterizations of this property. We will use one that applies in case the variety is locally finite.

We will say a variety $\mathcal{V}$ has **half $\diamondsuit$-terms** if there are three terms $s_0(x, y, z), s_1(x, y, z)$, and $s_2(x, y, z)$ so that all the equations below are true in $\mathcal{V}$.

$$\begin{aligned}
s_0(x, y, y) &\approx s_2(x, y, y)\\
s_1(x, x, y) &\approx s_2(x, y, x)\\
s_0(x, x, y) &\approx s_0(y, x, y)\\
s_1(x, y, y) &\approx s_1(x, y, x)\\
s_2(x, x, y) &\approx s_1(x, x, y)\\
s_0(x, x, x) &\approx s_1(x, x, x) \approx s_2(x, x, x) \approx x
\end{aligned}$$

These equations are called the **half $\diamondsuit$-equations**.

A Characterization of Locally Finite Congruence Meet-Semidistributivity.
Every locally finite congruence meet-semidistributive variety has half $\diamondsuit$-terms. Every variety with half $\diamondsuit$-terms is congruence meet-semidistributive.

We will not prove this theorem here.

A proof of Willard's Finite Basis Theorem.
Our proof relies on three lemmas. The first, which lays out the organization of the proof, requires another notion.

Let $\mathcal{K}$ be a class of algebras of the same signature. We will say that $\mathcal{K}$ has the **definite atoms property** provided there is an elementary sentence σ so that

- The sentence σ is true in $\mathcal{K}$, and
- For all $\mathbf{A}$ such that $\mathbf{A} \models \sigma$ and for all $a, b, p, q \in A$ such that $p \neq q$, and for all natural numbers n and m such that $\{a, b\} \not\hookrightarrow^n_m \{p, q\}$ we have and

$$\mathbf{A} \models \exists u, v \left[\neg u \approx v \nLeftarrow \{a, b\} \not\hookrightarrow^1_m \circ \not\hookrightarrow^2_1 \{u, v\} \text{ and } \{p, q\} \not\hookrightarrow^1_1 \{u, v\} \right].$$

- There is a natural number ℓ such that for all $\mathbf{A}$ such that $\mathbf{A} \models \sigma$ and for all $a, b \in A$ and all atoms $\alpha \subseteq \mathrm{Cg}^{\mathbf{A}}(a,b)$, and for all $(p,q) \in \alpha$ such that $p \neq q$, we have

$$\mathbf{A} \models \exists u, v \left[\neg u \approx v \nwarrow \{a,b\} \nrightarrow^1_\ell \circ \nrightarrow^2_1 \{u,v\} \nwarrow \{p,q\} \nrightarrow^1_1 \{u,v\} \right].$$

We need three lemmas. The first traces its ancestry back to Bjarni Jónsson's proof of Kirby Baker's Finite Basis Theorem.

The Definite Atoms Finite Basis Lemma.
Let $\mathcal{V}$ be a variety in a finite signature.
If

- *the variety $\mathcal{V}$ has the definite atoms property,*
- *$\mathcal{V}$ is locally finite, and*
- *$\mathcal{V}_{\mathrm{fsi}}$ is finitely axiomatizable,*

then $\mathcal{V}$ is finitely based.

Proof. Let n and m be natural numbers and let $\pi_m(x_0, y_0, x_1, y_1)$ denote the formula

$$\exists u, v \left[\neg u \approx v \nwarrow \{x_0, y_0\} \nrightarrow^1_m \circ \nrightarrow^2_2 \{u,v\} \nwarrow \{x_1, y_1\} \nrightarrow^1_m \circ \nrightarrow^2_2 \{u,v\} \right].$$

Let θ^n_m be the following sentence

$$\forall x, y, z, w \left[\pi_m(x,y,z,w) \Rightarrow \pi^2_{\ell+2}(x,y,z,w) \right].$$

Contention. If $\mathbf{A} \in \mathcal{V}$ and $a, b, c, d \in A$ so that $\mathrm{Cg}^{\mathbf{A}}(a,b) \cap \mathrm{Cg}^{\mathbf{A}}(c,d)$ is not trivial, then $\mathbf{A} \models \pi_\ell(a,b,c,d)$.

Proof. Suppose that $\mathbf{A} \in \mathcal{V}$ and that $a, b, c, d \in A$ so that $\mathrm{Cg}^{\mathbf{A}}(a,b) \cap \mathrm{Cg}^{\mathbf{A}}(c,d)$ is not trivial. So pick $p', q' \in A$ with $p' \neq q'$ so that $(p', q') \in \mathrm{Cg}^{\mathbf{A}}(a,b) \cap \mathrm{Cg}^{\mathbf{A}}(c,d)$. The computations witnessing this can be carried out in a finitely generated subalgebra $\mathbf{B}$ of $\mathbf{A}$. But $\mathcal{V}$ is locally finite, so $\mathbf{B}$ is finite. Observe that $\mathbf{B} \in \mathcal{V}$. Since $\mathbf{B}$ is finite every nontrivial congruence of $\mathbf{B}$ lies above an atom. Now pick p, q with $p \neq q$ and $(p,q) \in \mathrm{Cg}^{\mathbf{B}}(p', q')$ and $\mathrm{Cg}^{\mathbf{B}}(p,q)$ is an atom. Using the definite atoms property, we pick $p_0, p_1 \in B$ with $p_0 \neq p_1$ so that

$$\mathbf{B} \models \{a,b\} \nrightarrow^1_\ell \circ \nrightarrow^2_1 \{p_0, p_1\} \nwarrow \{p,q\} \nrightarrow^1_1 \{p_0, p_1\}.$$

Notice that $(p_0, p_1) \in \mathrm{Cg}^{\mathbf{B}}(c,d)$. So invoke the definite atoms property again to obtain $r_0, r_1 \in A$ with $r_0 \neq r_1$ so that

$$\mathbf{B} \models \{c,d\} \nrightarrow^1_\ell \circ \nrightarrow^2_1 \{r_0, r_1\} \nwarrow \{p_0, p_1\} \nrightarrow^1_1 \{r_0, r_1\}.$$

Now $\{a,b\} \nrightarrow^1_\ell \circ \nrightarrow^2_1 \{p_0, p_1\} \nrightarrow^1_1 \{r_0, r_1\}$ yields $\{a,b\} \nrightarrow^1_\ell \circ \nrightarrow^2_2 \{r_0, r_1\}$. This means

$$\mathbf{B} \models \{a,b\} \nrightarrow^1_\ell \circ \nrightarrow^2_2 \{r_0, r_1\} \nwarrow \{c,d\} \nrightarrow^1_\ell \circ \nrightarrow^2_2 \{r_0, r_1\}.$$

That is $\mathbf{B} \models \pi_\ell(a,b,c,d)$. But this is an existential formula. So $\mathbf{A}$ is a model as well, as desired. □

Let ψ be the following sentence

$$\forall x_0, y_0, x_1, y_1 \left[\pi_{\ell+1}(x_0, y_0, x_1, y_1) \implies \pi_\ell(x_0, y_0, x_1, y_1) \right]$$

It follows from the contention above that $\mathcal{V} \models \psi$. Let $\mathcal{W} = \mathrm{Mod}\{\sigma, \psi\}$. So $\mathcal{W}$ is a finitely axiomatizable class that includes $\mathcal{V}$. Let φ be the sentence

$$\forall x_0, y_0, x_1, y_1 \left[(\neg x_0 \approx y_0 \nwarrow \neg x_1 \approx y_1) \Leftrightarrow \pi_\ell(x_0, y_0, x_1 . y_1) \right].$$

It is our contention that $\{\sigma, \psi, \varphi\}$ axiomatizes $\mathcal{W}_{\text{fsi}}$. Certainly every model of those three sentencs belongs to $\mathcal{W}_{\text{fsi}}$. So we will concern ourselves with the reverse inclusion. Let $\mathbf{A} \in \mathcal{W}_{\text{fsi}}$. So we know $\mathcal{A} \models \{\sigma, \psi\}$. We have to establish $\mathbf{A} \models \varphi$. To this end, let $a, b, c, d \in A$. First suppose that $a \neq b$ and $c \neq d$. Since $\mathbf{A}$ is finitely subdirectly irreducible, we can pick $r, s \in A$ with $r \neq s$ and $(r, s) \in \text{Cg}^{\mathbf{A}}(a, b) \cap \text{Cg}^{\mathbf{A}}(c, d)$. By Mal'cev's description of principal congruences, there are numbers n and m so that

$$\{a, b\} \not\hookrightarrow^n_m \{r, s, \} \text{ and } \{c, d\} \not\hookrightarrow^n_m \{r, s, \}$$

Since $\mathbf{A} \models \sigma$, we can apply the definite atom property (twice as we did above) to obtain

$$\mathbf{A} \models \pi_m(a, b, c, d).$$

This doesn't quite get us $\mathbf{A} \models \varphi$ since m might be larger than ℓ. But now we can invoke ψ to decrease m step by step to ℓ. So we have established the left-to-right direction within φ. To obtain the other implication, let us suppose $a = b$. $\pi_\ell(a, b, c, d)$ cannot hold since we would have to have some $r_0, r_1 \in A$ with $r_0 \neq r_1$ and $(r_0, r_1) \in \text{Cg}^{\mathbf{A}}(a, b)$. The case when $c \neq d$ is similar. This means that $\mathbf{A} \models \varphi$ as desired.

Let us gather together what we know. There is a finitely axiomatizable class $\mathcal{W}$ such that $\mathcal{V} \subseteq \mathcal{W}$ and $\mathcal{W}_{\text{fsi}}$ is finitely axiomatizable. We also know that $\mathcal{V}_{\text{fsi}}$ is finitely axiomatizable (its one of our hypotheses).

Now let τ be a sentence that axiomatizes $\mathcal{V}_{\text{fsi}}$. Then we have $\mathcal{V} \models (\sigma \curlywedge \psi \curlywedge \varphi) \implies \tau$. Let Γ be a finite set of equations true in $\mathcal{V}$ so that $\Gamma \models \Sigma \cup \{\theta \to \tau\}$ and put $\mathcal{V}' = \text{Mod}\,\Gamma$. So $\mathcal{V} \subseteq \mathcal{V}' \subseteq W$. But then $\mathcal{V}_{\text{fsi}} = \mathcal{V}'_{\text{fsi}}$. So $\mathcal{V} = \mathcal{V}'$. Since $\mathcal{V}'$ is finitely based, we conclude that $\mathcal{V}$ is finitely based. □

So we need to show that a variety of finite signature that has a finite residual bound and half $\diamondsuit$-terms, must have the definite atoms property and $\mathcal{V}_{\text{fsi}}$ is finitely axiomatizable. Our second lemma addresses the last point.

Folklore Lemma.
Let $\mathcal{V}$ be a variety such that $\mathcal{V}_{\text{si}}$ is axiomatizable by a set of elementary sentences. Every finitely subdirectly irreducible algebra in $\mathcal{V}$ is embeddable into some subdirectly irreducible algebra in $\mathcal{V}$.

Proof. Let $\mathbf{B}$ be a finitely subdirectly irreducible algebra belonging to $\mathcal{V}$ and let Φ be a set of elementary sentences which axiomatizes $\mathcal{V}_{\text{si}}$. Expand the signature by adding a new constant to name each element of B. We use $\mathbf{B}^*$ to denote the corresponding expansion of $\mathbf{B}$. Let Δ be the atomic diagram of $\mathbf{B}$. To prove the Lemma we need to show that $\Delta \cup \Phi$ has a model. In view of the Compactness Theorem, we need only show that $\Gamma \cup \Phi$ has a model whenever Γ is a finite subset of Δ. So consider such a Γ. Without loss of generality, we assume that the only negated equations in Γ are of the form $c \neq d$ were c and d are constant symbols. Let S be the set of all elements of B named by constants occurring in Γ. Since S is finite and $\mathbf{B}$ is finitely subdirectly irreducible, pick $p, q \in B$ with $p \neq q$ and so that $\langle p, q \rangle \in \text{Cg}^{\mathbf{B}}(r, s)$ whenever r and s are distinct elements of S. Let θ be a maximal congruence of $\mathbf{B}$ which separates r and s. Then $\mathbf{B}/\theta$ is subdirectly irreducible. So $\mathbf{B}^*/\theta \models \Phi$. Now the equations in Γ hold in $\mathbf{B}^*/\theta$ since equations are preserved in the passage to quotient algebras. The negated equations in Γ also hold in $\mathbf{B}^*/\theta$ since θ separates all the elements of S. Therefore, $\mathbf{B}^*/\theta \models \Gamma \cup \Phi$, as desired. □

Thus in a variety $\mathcal{V}$ of finite signature with a finite residual bound b, every finitely subdirectly irreducible algebra must have cardinality bounded by b as well. It follows that every finitely subdirectly irreducible algebra is finite and so it is subdirectly irreducible. Thus, $\mathcal{V}_{\text{fsi}} = \mathcal{V}_{\text{si}}$. In consequence, $\mathcal{V}_{\text{fsi}}$ is also finitely axiomatizable.

We still need to obtain the definite atoms property. The third lemma does that. It traces its heritage ultimately back to Kriby Baker's original proof of Baker's Finite Basis Theorem.

The Half $\diamondsuit$ Single Sequence Lemma.
Let $\mathcal{W}$ be a class of algebras of finite signature with half $diamondsuit$-terms. Let $\mathbf{A} \in \mathcal{W}$ and $a, b, p, q \in A$ with $p \neq q$ and n and ℓ be natural numbers such that $\{a, b\} \twoheadrightarrow^n_\ell \{p, q\}$. Then there are elements $e, f \in A$ so that

$$\begin{gathered}\{a, b\} \twoheadrightarrow^1_\ell \circ \twoheadrightarrow^2_1 \{e, f\} \\ \{p, q\} \twoheadrightarrow^1_1 \{e, f\}, \text{ and} \\ e \neq f.\end{gathered}$$

Moreover, if $\mathcal{W}$ is a variety with a finite residual bound, then there is a positive natural number ℓ such that for every $\mathbf{A} \in \mathcal{V}$ and $a, b \in A$ and every atom $\alpha \subseteq \mathrm{Cg}^{\mathbf{A}}(a, b)$ and every $(p, q) \in \alpha$ with $p \neq q$, there are elements $e, f \in A$ so that

$$\begin{gathered}\{a, b\} \twoheadrightarrow^1_\ell \circ \twoheadrightarrow^2_1 \{e, f\} \\ \{p, q\} \twoheadrightarrow^1_1 \{e, f\}, \text{ and} \\ e \neq f.\end{gathered}$$

Here is a diagram of this Lemma.

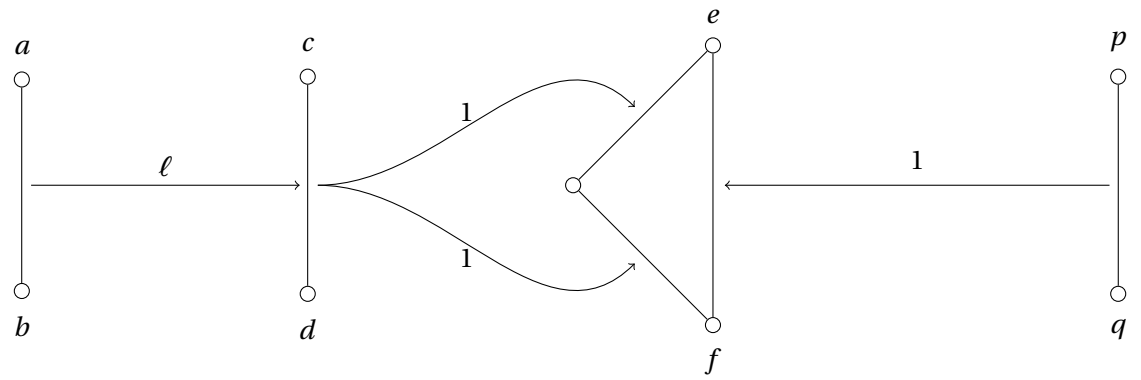

Proof of the Single Sequence Lemma.
We see that there is a sequence $p = r_0, r_1, \ldots, r_n = q$ of elements of A and translations of $\lambda_0, \ldots, \lambda_{n-1}$ of complexity no more than ℓ so that

$$\{\lambda_j(a), \lambda_j(b)\} = \{r_j, r_{j+1}\} \text{ for all } j < n.$$

Observe that $p \neq q$. Since the half $\diamondsuit$ equations hold in $\mathbf{A}$, we see

$$\mathbf{A} \models \forall x, y \left[\left(s_0(x, y, y) \approx y \twoheadrightarrow s_1(x, x, y) \approx x \twoheadrightarrow s_2(x, y, y) \approx s_2(x, y, x) \right) \Longrightarrow x \approx y \right].$$

Establishing this uses the first two of the half $\diamondsuit$-equations. This means that there are three alternatives:

$$s_0^{\mathbf{A}}(p, q, q) \neq q \quad \text{or} \quad s_1^{\mathbf{A}}(p, p, q) \neq p \quad \text{or both} \quad s_2^{\mathbf{A}}(p, q, q) \neq s_2^{\mathbf{A}}(p, q, p) \quad \text{and} \quad s_1^{\mathbf{A}}(p, p, q) = p.$$

Alternative I: $s_0^{\mathbf{A}}(p, q, q) \neq q = s_0^{\mathbf{A}}(q, q, q)$
Observe that $s_0^{\mathbf{A}}(p, p, q) = s_0^{\mathbf{A}}(q, p, q)$ by the third half $\diamondsuit$-equation.. This means that the equation $s_0^{\mathbf{A}}(p, x, q) = s_0^{\mathbf{A}}(q, x, q)$ holds at the left end of the sequence $p = r_0, r_1, \ldots, r_{n-1}, r_n = q$ but fails at the right end. There

must be a leftmost place in this sequence where the equation fail. Let this element at this position be $r_{i+1} = c$ and put $r_i = d$. Then

$$s_0^{\mathbf{A}}(p,c,q) \neq s_0^{\mathbf{A}}(q,c,q) \quad \text{and} \quad s_0^{\mathbf{A}}(p,d,q) = s_0^{\mathbf{A}}(q,d,q).$$

Let $e = s_0^{\mathbf{A}}(p,c,q)$ and $f = s_0^{\mathbf{A}}(q,c,q)$. This completes the proof of the first part of the Lemma under Alternative I.

Alternative II: $s_1^{\mathbf{A}}(p,p,q) \neq p = s_1^{\mathbf{A}}(p,p,p)$
Observe that $s_1^{\mathbf{A}}(p,q,q) = s_1^{\mathbf{A}}(p,q,p)$ by the fourth half $\diamondsuit$-equation. This means that the equation $s_1^{\mathbf{A}}(p,x,q) = s_1^{\mathbf{A}}(p,x,p)$ holds at the right end of the sequence $p = r_0, r_1, \ldots, r_{n-1}, r_n = q$ but fails at the left end. There must be a rightmost place in this sequence where the equation fails. Let this element at this position be $r_{i+1} = d$ and put $r_i = c$. Then

$$s_1^{\mathbf{A}}(p,c,q) \neq s_1^{\mathbf{A}}(p,c,p) \quad \text{and} \quad s_1^{\mathbf{A}}(p,d,q) = s^{\mathbf{A}}(p,d,p).$$

Let $e = s_1^{\mathbf{A}}(p,c,q)$ and $f = s_1^{\mathbf{A}}(p,c,p)$. This completes the proof of the first part of the Lemma under Alternative II.

Alternative III: $s_2^{\mathbf{A}}(p,q,q) \neq s_2^{\mathbf{A}}(p,q,p)$ **and** $s_1^{\mathbf{A}}(p,p,q) = p$
The equation $s_2^{\mathbf{A}}(p,x,q) = s_2^{\mathbf{A}}(p,x,p)$ fail at the right end of the sequence $p = r_0, r_1, \ldots, r_{n-1}, r_n = q$ but hold at the left end, since $s_2(x,x,y) \approx s_1(x,x,y)$ is from the fifth half $\diamondsuit$ equation. So $s_2^{\mathbf{A}}(p,p,q) = s_1^{\mathbf{A}}(p,p,q) = p = s_2^{\mathbf{A}}(p,p,p)$ by an idempotence equation among the half $\diamondsuit$-equations. There must be a leftmost place in this sequence where the equation fails. Let this element at this position be $r_{i+1} = c$ and put $r_i = d$. Then

$$s_2^{\mathbf{A}}(p,c,q) \neq s_2^{\mathbf{A}}(p,c,p) \quad \text{and} \quad s_2^{\mathbf{A}}(p,d,q) = s_2^{\mathbf{A}}(p,d,p).$$

Let $e = s_2^{\mathbf{A}}(p,c,q)$ and $f = s_2^{\mathbf{A}}(p,c,p)$.

This completes the proof of the first part of the Lemma under Alternative III.

For the second part of the Lemma, we have the additional assumption that $\mathcal{W}$ is a variety with a finite residual bound. So we know that, up to isomorphism, there are only finitely many subdirectly irreducible algebras in $\mathcal{W}$ and they are all finite. Each of these subdirectly irreducible algebras can have only finitely many translations. So there is a positive natural number ℓ that is an upper bound on the complexity of all the translations in all the subdirectly irreducible algebras in $\mathcal{W}$. The point of this part of the Lemma is that this bound applies to *all* $\mathbf{A}$ and *all* $a, b, p, q \in A$ with $p \neq q$ such that $(p,q) \in \mathrm{Cg}^{\mathbf{A}}(a,b)$ such that $\mathrm{Cg}^{\mathbf{A}}(p,q)$ is an atom in $\mathrm{Con}\mathbf{A}$.

Pick a congruence $\theta \in \mathrm{Con}\mathbf{A}$ that is maximal with respect to separating p and q. Then $\mathbf{A}/\theta$ is subdirectly irreducible and $(p/\theta, q/\theta)$ is a critical pair of $\mathbf{A}/\theta$. Notice that θ must also separate a and b

Now $(p/\theta, q/\theta) \in \mathrm{Cg}^{\mathbf{A}/\theta}(a/\theta, b/\theta)$. According to Mal'cev there is finite sequence $r'_0, \ldots, r'_k$ in $\mathbf{A}/\theta$ and translations $\lambda'_0, \ldots, \lambda'_k$ such that

$$p/\theta = r'_0$$
$$\{\lambda'_i(a/\theta), \lambda'_i(b/\theta)\} = \{r'_i, r'_{i+1}\} \text{ for all } i < k$$
$$r'_k = q/\theta$$

and, moreover, the complexity of the translations λ'_i never exceeds ℓ. We pick representatives of the congruence classes and denote this by removing the primes. So we get

$$p \,\theta\, r_0$$
$$\{\lambda_i(a), \lambda_i(b)\} \,\theta\, \{r_i, r_{i+1}\} \text{ for all } i < k$$
$$r_k \,\theta\, q$$

To make the middle line of the above display more specific, for each $i < k$ pick a_i and b_i so that $\{a_i, b_i\} = \{a, b\}$ and

$$\begin{aligned} \lambda_i(a_i)\,\theta\, r_i \\ \lambda_i(b_i)\,\theta\, r_{i+1} \end{aligned}$$

We obtain

$$\begin{aligned} p\,\theta\, r_0\,\theta\, \lambda_0(a_0) \\ \lambda_1(a_1)\,\theta\, r_1\,\theta\, \lambda_0(b_0) \\ \lambda_1(b_1)\,\theta\, r_2\,\theta\, \lambda_2(a_2) \\ \vdots \\ \lambda_{2j+1}(a_{2j+1})\,\theta\, r_{2j+1}\,\theta\, \lambda_{2j}(b_{2j}) \\ \lambda_{2j+1}(b_{2j+1})\,\theta\, r_{2j+2}\,\theta\, \lambda_{2j+2}(a_{2j+2}) \\ \vdots \\ \lambda_{k-1}(b_{k-1})\,\theta\, r_k\,\theta\, q \end{aligned}$$

where we have made the harmless assumption that k is even. Of course, this leads to

$$\begin{aligned} p\,\theta\, \lambda_0(a_0) \\ \lambda_1(a_1)\,\theta\, \lambda_0(b_0) \\ \lambda_1(b_1)\,\theta\, \lambda_2(a_2) \\ \vdots \\ \lambda_{2j+1}(a_{2j+1})\,\theta\, \lambda_{2j}(b_{2j}) \\ \lambda_{2j+1}(b_{2j+1})\,\theta\, \lambda_{2j+2}(a_{2j+2}) \\ \vdots \\ \lambda_{k-1}(b_{k-1})\,\theta\, q \end{aligned}$$

Consider the *single sequence* made by traversing the display above in a back-and-forth manner:

$$p, \lambda_0(a_0), \lambda_0(b_0), \lambda_1(a_1), \lambda_1(b_1), \ldots, \lambda_{k-1}(a_{k-1}), \lambda_{k-1}(b_{k-1}), q.$$

Again, there are two alternatives.

Alternative I: $s_0^{\mathbf{A}}(p,q,q) \neq q = s_0^{\mathbf{A}}(q,q,q)$
Observe that $s_0^{\mathbf{A}}(p,p,q) = s_0^{\mathbf{A}}(q,p,q)$. This means that the equation $s_0^{\mathbf{A}}(p,x,q) = s_0^{\mathbf{A}}(q,x,q)$ holds at the left end of the single sequence but fails at the right end. There must be a leftmost place in this sequence where the equation fails. There are three cases.

Case (a): $s_0^{\mathbf{A}}(p,\lambda_j(a_j),q) = s_0^{\mathbf{A}}(q,\lambda_j(a_j),q)$ **and** $s_0^{\mathbf{A}}(p,\lambda_j(b_j),q) \neq s_0^{\mathbf{A}}(q,\lambda_j(b_j),q)$
In this case, we can let $d = \lambda_j(a_j)$ and $c = \lambda_j(b_j)$ to reach the desired conclusion.

Case (b): $s_0^{\mathbf{A}}(p,\lambda_j(b_j),q) = s_0^{\mathbf{A}}(q,\lambda_j(b_j),q)$ **and** $s_0^{\mathbf{A}}(p,\lambda_{j+1}(a_{j+1}),q) \neq s_0^{\mathbf{A}}(q,\lambda_{j+1}(a_{j+1}),q)$
We will reject this case. To simplify notation, let $u = \lambda_{j+1}(a_{j+1})$ and let $v = \lambda_j(b_j)$. So we have

$$s_0^{\mathbf{A}}(p,v,q) = s^{\mathbf{A}}(q,v,q) \quad \text{and} \quad s_0^{\mathbf{A}}(p,u,q) \neq s_0^{\mathbf{A}}(q,u,q).$$

Notice that $\langle s_0^{\mathbf{A}}(p,u,q), s_0^{\mathbf{A}}(q,u,q)\rangle \in \mathrm{Cg}^{\mathbf{A}}(p,q)$. Since $\mathrm{Cg}^{\mathbf{A}}(p,q)$ is an atom we have

$$\mathrm{Cg}^{\mathbf{A}}(p,q) = \mathrm{Cg}^{\mathbf{A}}(s_0^{\mathbf{A}}(p,u,q), s_0^{\mathbf{A}}(q,u,q)).$$

But $v\,\theta\,u$. This leads to

$$s_0^{\mathbf{A}}(p,u,q)\;\theta\; s_0^{\mathbf{A}}(p,v,q) = s_0^{\mathbf{A}}(q,v,q)\;\theta\; s_0^{\mathbf{A}}(q,u,q).$$

As θ is not above $\mathrm{Cg}^{\mathbf{A}}(p,q)$ we reject this case.

Case (c): $s_0^{\mathbf{A}}(p,\lambda_{k-1}(b_{k-1}),q) = s_0^{\mathbf{A}}(q,\lambda_{k-1}(b_{k-1}),q)$ **and** $s_0^{\mathbf{A}}(p,q,q) \neq s_0^{\mathbf{A}}(q,q,q)$
We will reject this case as well. Again, to simplify notation we take $v = \lambda_{k-1}(b_{k-1})$. So we have

$$s_0^{\mathbf{A}}(p,v,q) = s^{\mathbf{A}}(q,v,q) \quad \text{and} \quad s_0^{\mathbf{A}}(p,q,q) \neq s_0^{\mathbf{A}}(q,q,q).$$

As above, we deduce that

$$\mathrm{Cg}^{\mathbf{A}}(p,q) = \mathrm{Cg}^{\mathbf{A}}(s_0^{\mathbf{A}}(p,u,q), q).$$

But $v\,\theta\,q$. This leads to

$$s_0^{\mathbf{A}}(p,u,q)\;\theta\; s_0^{\mathbf{A}}(p,v,q) = s^{\mathbf{A}}(q,v,q)\;\theta\; s_0^{\mathbf{A}}(q,q,q) = q.$$

As θ is not above $\mathrm{Cg}^{\mathbf{A}}(p,q)$ we reject this case.

Alternative II: $s_1^{\mathbf{A}}(p,p,q) \neq p = s_1^{\mathbf{A}}(p,p,p)$
Observe that $s_1^{\mathbf{A}}(p,q,q) = s_1^{\mathbf{A}}(p,q,p)$. This means that the equation $s_1^{\mathbf{A}}(p,x,q) = s_1^{\mathbf{A}}(p,x,p)$ holds at the right end of the single sequence but fails at the left end. There must be a rightmost place in this sequence where the equation fails. Again, there are three cases.

Case (a): $s_1^{\mathbf{A}}(p,\lambda_j(b_j),q) = s_1^{\mathbf{A}}(p,p,\lambda_j(b_j))$ **and** $s_1^{\mathbf{A}}(p,\lambda_j(a_j),q) \neq s_1^{\mathbf{A}}(p,p,\lambda_j(a_j))$
In this case, we let $c = \lambda_j(a_j)$ and $d = \lambda_j(b_j)$ to reach the desired conclusion.

Case (b): $s_1^{\mathbf{A}}(p,\lambda_{j+1}(a_{j+1}),q) = s_1^{\mathbf{A}}(p,\lambda_{j+1}(a_{j+1}),p)$ **and** $s_1^{\mathbf{A}}(p,\lambda_j(b_j),q) \neq s_1^{\mathbf{A}}(p,p,\lambda_j(b_j))$
We will reject this case. To reduce notation, let $u = \lambda_j(b_j)$ and $v = \lambda_{j+1}(a_{j+1})$. Then we have

$$s_1^{\mathbf{A}}(p,v,q) = s_1^{\mathbf{A}}(p,p,v) \quad \text{and} \quad s_1^{\mathbf{A}}(p,u,q) \neq s_1^{\mathbf{A}}(p,p,u).$$

Notice that $\langle s_1^{\mathbf{A}}(p,u,q), s_1^{\mathbf{A}}(p,u,p)\rangle \in \mathrm{Cg}^{\mathbf{A}}(p,q)$. Since $\mathrm{Cg}^{\mathbf{A}}(p,q)$ is an atom we have

$$\mathrm{Cg}^{\mathbf{A}}(p,q) = \mathrm{Cg}^{\mathbf{A}}(s_1^{\mathbf{A}}(p,u,q), s_1^{\mathbf{A}}(p,u,p)).$$

But $v\,\theta\,u$. This leads to

$$s_1^{\mathbf{A}}(p,u,q)\;\theta\; s_1^{\mathbf{A}}(p,v,q) = s_1^{\mathbf{A}}(p,p,v) = s_1^{\mathbf{A}}(p,v,p)\;\theta\; s_1^{\mathbf{A}}(p,u,p).$$

As θ is not above $\mathrm{Cg}^{\mathbf{A}}(p,q)$ we reject this case.

Case (c): $s_1^{\mathbf{A}}(p,\lambda_0(a_0),q) = s_1^{\mathbf{A}}(p,\lambda_0(a_0),p)$ **and** $s_1^{\mathbf{A}}(p,p,q) \neq s_1^{\mathbf{A}}(p,p,p)$
We will reject this case. Again, to simplify notation we take $v = \lambda_0(a_0)$. So we have

$$s_1^{\mathbf{A}}(p,v,q) = s_1^{\mathbf{A}}(p,v,p) \quad \text{and} \quad s_1^{\mathbf{A}}(p,p,q) \neq s_1^{\mathbf{A}}(p,p,p).$$

As above, we deduce that

$$\mathrm{Cg}^{\mathbf{A}}(p,q) = \mathrm{Cg}^{\mathbf{A}}(s_1^{\mathbf{A}}(p,v,q), p).$$

But $v\,\theta\,p$. This leads to

$$s_1^{\mathbf{A}}(p,v,q) = s_1^{\mathbf{A}}(p,v,p)\;\theta\; s_1^{\mathbf{A}}(p,p,p) = p$$

As θ is not above $\mathrm{Cg}^{\mathbf{A}}(p,q)$ we reject this case.

Alternative III: $s_2^{\mathbf{A}}(p,q,q) \neq s_2^{\mathbf{A}}(p,q,p)$ **and** $s_1^{\mathbf{A}}(p,p,q) = p$
The equation $s_2^{\mathbf{A}}(p,x,q) = s_2^{\mathbf{A}}(p,x,p)$ fail at the right end of the single sequence but hold at the left end, since $s_2(x,x,y) \approx s_2(x,y,x) \approx s_1(x,x,y)$ are among the half $\diamondsuit$ equations. There must be a leftmost place in this sequence where the equation fails. There are three cases.

Case (a): $s_2^{\mathbf{A}}(p,\lambda_j(a_j),q) = s_2^{\mathbf{A}}(p,p,\lambda_j(a_j))$ **and** $s_1^{\mathbf{A}}(p,\lambda_j(b_j),q) \neq s_1^{\mathbf{A}}(p,p,\lambda_j(b_j))$
In this case, we let $c = \lambda_j(b_j)$ and $d = \lambda_j(a_j)$ to reach the desired conclusion.

Case (b): $s_2^{\mathbf{A}}(p,\lambda_j(b_j),q) = s_2^{\mathbf{A}}(p,\lambda_j(b_j),p)$ **and** $s_2^{\mathbf{A}}(p,\lambda_{j+1}(a_{j+1}),q) \neq s_2^{\mathbf{A}}(p,p,\lambda_{j+1}(a_{j+1}))$
We will reject this case. To reduce notation, let $u = \lambda_j(b_j)$ and $v = \lambda_{j+1}(a_{j+1})$. Then we have

$$s_2^{\mathbf{A}}(p,u,q) = s_2^{\mathbf{A}}(p,u,p) \quad \text{and} \quad s_2^{\mathbf{A}}(p,v,q) \neq s_2^{\mathbf{A}}(p,v,p).$$

Notice that $\langle s_2^{\mathbf{A}}(p,v,q), s_2^{\mathbf{A}}(p,v,p)\rangle \in \mathrm{Cg}^{\mathbf{A}}(p,q)$. Since $\mathrm{Cg}^{\mathbf{A}}(p,q)$ is an atom we have

$$\mathrm{Cg}^{\mathbf{A}}(p,q) = \mathrm{Cg}^{\mathbf{A}}(s_2^{\mathbf{A}}(p,v,q), s_2^{\mathbf{A}}(p,v,p)).$$

But $v\,\theta\,u$. This leads to

$$s_2^{\mathbf{A}}(p,v,q)\;\theta\; s_2^{\mathbf{A}}(p,u,q) = s_2^{\mathbf{A}}(p,u,p) = s_2^{\mathbf{A}}(p,v,p).$$

As θ is not above $\mathrm{Cg}^{\mathbf{A}}(p,q)$ we reject this case.

Case (c): $s_2^{\mathbf{A}}(p,\lambda_{k-1}(b_{k-1}),q) = s_2^{\mathbf{A}}(p,\lambda_{k-1}(b_{k-1}),p)$ **and** $s_2^{\mathbf{A}}(p,p,q) \neq s_2^{\mathbf{A}}(p,p,p)$
We will reject this case. Again, to simplify notation we take $v = \lambda_{k-1}(b_{k-1})$. So we have

$$s_2^{\mathbf{A}}(p,v,q) = s_2^{\mathbf{A}}(p,v,p) \quad \text{and} \quad s_2^{\mathbf{A}}(p,p,q) \neq s_2^{\mathbf{A}}(p,p,p) = p.$$

As above, we deduce that

$$\mathrm{Cg}^{\mathbf{A}}(p,q) = \mathrm{Cg}^{\mathbf{A}}(s_2^{\mathbf{A}}(p,v,q), p).$$

But $v\,\theta\,p$. This leads to

$$s_2^{\mathbf{A}}(p,v,q) = s_2^{\mathbf{A}}(p,v,p)\;\theta\; s_2^{\mathbf{A}}(p,p,p) = p.$$

As θ is not above $\mathrm{Cg}^{\mathbf{A}}(p,q)$ we reject this case.

This completes the proof of the Half $\diamondsuit$ Single Sequence Lemma. □

Now we put the pieces together to prove Willard's Finite Basis Theorem. According to the Half $\diamondsuit$ Single Sequence Lemma, we see that $\mathcal{V}$ has the definite atoms property, where σ is the conjunction of the half $\diamondsuit$-equations. As a consequence of the Folklore Lemma, we see $\mathcal{V}_{\mathrm{fsi}}$ is finitely axiomatizable. So Willard's Finite Basis Theorem follows from the Definable Atoms Finite Basis Lemma. □

LESSON 6

THE LATTICE OF EQUATIONAL THEORIES

For a fixed signature, the equational theories are ordered by set-inclusion $\subseteq$. More is true. There is a largest equational theory $\top$: the set of all equations of the signature. The set $\{x \approx y\}$, where x and y are different variables, is a base for $\top$. Equational theories other than $\top$ are called **proper** equational theories. There is a smallest equational theory $\bot$: the set of all equations of the form $s \approx s$, where s is a term. These equations are called **tautologies**. The theory $\bot$ is the **trivial** equational theory and is based on $\{x \approx x\}$. Mostly we will be interested in proper nontrivial equational theories. Moreover every set $\mathcal{F}$ of equational theories has a greatest lower bound or meet, namely $\bigcap \mathcal{F}$, provided $\mathcal{F}$ is not empty and the largest equational theory if $\mathcal{F}$ is empty. In symbols we write

$$\bigwedge \mathcal{F} = \begin{cases} \bigcap \mathcal{F} & \text{if } \mathcal{F} \neq \varnothing \\ \top & \text{otherwise} \end{cases}$$

Also, every set $\mathcal{F}$ has a least upper bound or join, namely $\mathrm{Th}\,\mathrm{Mod}\,(\bigcup \mathcal{F})$. In symbols

$$\bigvee \mathcal{F} = \mathrm{Th}\,\mathrm{Mod} \bigcup \mathcal{F}.$$

Indeed, all these things hold for the closed sets on either side of any Galois connection. Any ordered set with all these properties is referred to as a **complete lattice**.

Suppose **L** is a complete lattice. An element $c \in L$ is said to be **compact** provide if $c \leq \bigvee \mathcal{F}$, then $c \leq \bigvee \mathcal{F}^*$ for some finite $\mathcal{F}^* \subseteq \mathcal{F}$. The complete lattice **L** is said to be **algebraic**, provided each element of L is the join of some set of compact elements. It is not hard to see that the compact elements of the lattice of all equational theories are just those equational theories that are finitely based. It is also clear that every equational theory is the join of all finitely based theories it contains. That is, for any fixed signature, the lattice of all equational theories is an algebraic lattice.

The ordering in the lattice of equational theories provides a way to compare the strength of equational theories—the larger the theory the stronger it is. So understanding the structure of the lattice of equational theories of a given signature is one of the routes to a deeper study of equational logic.

6.1 MAXIMAL AND MINIMAL EQUATIONAL THEORIES

Once a signature is given, the lattice of equational theories will always have a top element. An equational theory T is said to be **maximal** or **equationally complete** provided T is a proper theory and if T' is an equational theory such that $T \subseteq T'$, then either $T = T'$ or $T' = \top$. That is, there is no equational theory properly between T and $\top$.

Theorem 6.1.1. *Every proper equational theory is contained in a maximal equational theory.*

Proof. This is a more or less standard Zorn's Lemma argument. Let T be a proper equational theory and let

$$\mathfrak{F} = \{S : T \subseteq S \text{ and } S \text{ is a a proper equational theory}\}.$$

Suppose $\mathfrak{C} \subseteq \mathfrak{F}$ is linearly ordered by $\subseteq$. We claim $\bigcup \mathfrak{C} \in \mathfrak{F}$. Certainly, $T \subseteq \bigcup \mathfrak{C}$ since $T \subseteq S$ for all $S \in \mathfrak{C}$.

Notice that $\bigcup \mathfrak{C}$ is an equational theory since it's closed under logical consequence. To see this, suppose $\bigcup \mathfrak{C} \models s \approx t$. Then there is a finite $\Delta \subseteq \bigcup \mathfrak{C}$ so that $\Delta \models s \approx t$. Pick $S_0, S_1, \ldots, S_n \in \mathfrak{C}$ so that $\Delta \subseteq S_0 \cup \cdots S_n$. Without loss of generality, since $\mathfrak{C}$ is a chain, we have $S_0 \subseteq S_1 \subseteq \cdots \subseteq S_n$. So $S_n \models s \approx t$, so $s \approx t \in S_n$ since S_n is an equational theory.

Lastly, $\bigcup \mathfrak{C}$ is proper. To see this, suppose to the contrary that $\mathfrak{C} \models x \approx y$. This would imply that $\Delta \models x \approx y$ for some finite $\Delta \subseteq \bigcap \mathfrak{C}$, which would imply that $S \models x \approx y$ for some $S \in \mathfrak{C}$. □

On the other hand, apart from certain meager signatures, the lattice of equational theories has no minimal nontrivial elements. In the language of lattices theory we would say that the lattice of equational theories has no atoms.

Theorem 6.1.2. *If the signature provides an operation symbol of positive rank, then any nontrivial equational theory has a proper nontrivial subtheory.*

Proof. Let $s \approx t \in T$ with $s \neq t$. Let Q be an operation symbol of positive rank. Pick m larger than the length of s and t. Let $T' = \mathrm{Th_e}\{Q^m ss \ldots s \approx Q^m tt \ldots t\}$. So $T' \subseteq T$ since $Q^m ss \ldots s \approx Q^m tt \ldots t$ is a consequence of $s \approx t$. T' is not trivial, since $s \neq t$. Finally, $s \approx t \notin T'$, because the terms s and t are too short for the equation in the basis of T' to come into play in any deducation. So $T' \neq T$. That is, T' is a proper nontrivial subtheory of T. □

For signatures that do not provide operation symbols of positive rank, it is possible to describe the lattice of equational theories. If fewer than 2 constant symbols are provided, then there are only two equational theories: $\top$ and $\bot$. If at least 2 constant symbols are available, then the lattice of equational theories is isomorphic to the lattice of equivalence relations on the set of constant symbols with a new top element adjoined.

Theorem on the Number of Equational Complete Theories.
Suppose the signature provides at least two operation symbols of rank 1 *or at least one operation symbols of rank at least* 2*. Then there are at least* 2^ω *equationally complete theories.*

Proof.
First consider the case that the signature provides at least two one-place operations symbols F and G.

Let n be any natural number. Put the term

$$s_n := FGF^{n+1}G^2x$$

where x is a variable.

Let X be a set of natural numbers. Define

$$\Sigma_X = \{s_n \approx x : n \in X\} \cup \{s_n \approx s_n(y) : n \notin X\}$$

Recall that $s_n(y)$ is the result of substituting y for x in s_n. That first set says that $FGF^{n+1}G^2$ is the identity function when $n \in X$ and the second set says that $FGF^{n+1}G^2$ is a constant function when $n \notin X$.

Let T_X be the equational theory based on Σ_X. We establish two claims:

Claim 0: T_X is a proper theory.

Proof of Claim 0.
T_X is a proper theory is equivalent to saying that Σ_X has a nontrivial model which is equivalent to saying $\Sigma_X \nvdash x \approx y$.

Let A be the set of terms—this is an infinite set. Let $\mathbf{A}$ be the algebra on the set of terms where each operation symbol, except F, denotes the same basic operation that it does in the term algebra. So G, for example, denotes in $\mathbf{A}$ the operation

$$G^{\mathbf{A}}(t) = Gt$$

for any term t in A. We change only the basic operation denoted by F in $\mathbf{A}$. Let it be

$$F^{\mathbf{A}}(t) = \begin{cases} v_0 & \text{if } Ft = s_n(u) \text{ for some } n \notin X \text{ and for some term } u \\ u & \text{if } Ft = s_n(u) \text{ for some } n \in X \text{ and some term } u \\ Ft & \text{otherwise} \end{cases}$$

It is a consequence of unique readability of terms that $F^{\mathbf{A}}$ is well-defined.

Let n be a natural number and let p be a proper subterm of $FGF^{n+1}G^2x$. Our contention is that

$$p^{\mathbf{A}}(u) = p(u).$$

We proceed by induction. If $p = x$, then $x^{\mathbf{A}}(u) = u = x(u)$. For the inductive step, there are two cases. Suppose first that $p = Gq$. Then $p^{\mathbf{A}}(u) = G^{\mathbf{A}}\left(q^{\mathbf{A}}(u)\right) = G^{\mathbf{A}}(q(u)) = Gq(u) = p(u)$. Second, suppose $p = Fq$. Then $p^{\mathbf{A}}(u) = F^{\mathbf{A}}\left(q^{\mathbf{A}}(u)\right)$. By the inductive hypothesis, $q^{\mathbf{A}}(u) = q(u)$ so we get $p^{\mathbf{A}}(u) = F^{\mathbf{A}}\left(q(u)\right)$. Since p is a proper subterm of $FGF^{n+1}G^2x$ we must fall into the "otherwise" case of the definition of $F^{\mathbf{A}}$. So $p^{|mathbfA}(u) = F^{mathbfA}(q(u)) = Fq(u) = p(u)$. So the contention is established.

Now observe

$$(FGF^{n+1}G^2)^{\mathbf{A}}(t) = F^{\mathbf{A}}((GF^{n+1}G)^{\mathbf{A}}(t)) = F^{\mathbf{A}}(GF^{n+1}G^2t) = \begin{cases} v_0 & \text{if } n \notin X \\ t & \text{if } n \in X \end{cases}$$

This means that $FGF^{n+1}G^2$ denotes a constant function, if $n \notin X$ and it denotes the identity function, if $n \in X$. We conclude that $\mathbf{A}$ is a model (nontrivial) of Σ_X. So T_X is a proper equational theory. □

Claim 1: $\Sigma_X \cup \Sigma_Y$ has only trivial (1-element) models when X and Y are different.

Proof of Claim 1.
Well, if Σ_X and Σ_Y are different, then we can pick $n \in X$ with $n \notin Y$ (or the other way around). Without loss of generality, we'll pick $n \in X$ with $n \notin Y$. Then

$$\Sigma_X \cup \Sigma_Y \vdash x \approx s_n, s_n \approx s_n(y)$$

with the first equation belongs to Σ_X and the second to Σ_Y. So

$$\Sigma_X \cup \Sigma_Y \vdash x \approx s_n(y)$$

but then of course

$$\Sigma_X \cup \Sigma_Y \vdash x \approx y$$

□

We already saw that we can extend a proper theory to a maximal theory. So we'll extend T_X to M_X where $M_X \supseteq T_X$ and M_X is maximal. Then $M_X \cup M_Y$ has only trivial models if $X \neq Y$. Hence $M_X \neq M_Y$ if $X \neq Y$. Hence, the number of equationally complete theories is at least as large has the number of subsets of the set of natural numbers. This concludes the first case.

For the second case, we suppose that the signature provides an operation symbol Q of rank at least 2. Our plan is to follow the reasoning used in the first case. To simplify matters, we describe how to do this in case the rank of Q is 3. Then extending the reasoning to arbitrary ranks large than 1 is easy.

Let n be a natural number and let

$$s_n := QxxQ^{n+1}xx\ldots x$$

where that final string of x's has length $2n+3$. Here is the term s_2 depicted as a tree:

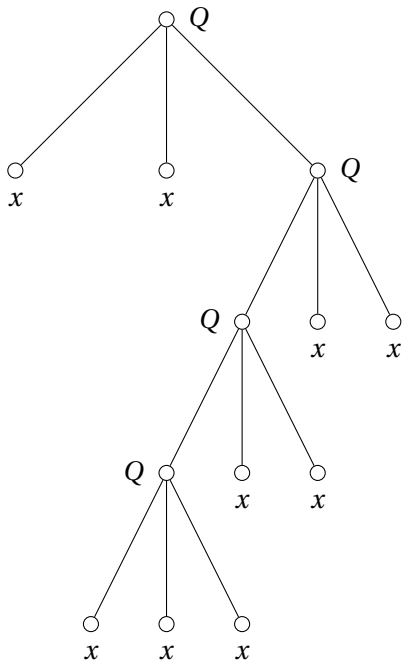

As this tree makes plain, s_n is unbalanced in favor of the rightmost node beneath the top. On the other hand, every proper subterm of s_n that is not a variable is either perfectly balanced, or unbalanced in favor of the leftmost node beneath its top.

Now the argument we used in the case of two one-place operation symbols goes through, provided we can establish here a version of Claim 0. We devise an infinite algebra by modifying the term algebra. So again let A be the set of all terms and take as basic operations of $\mathbf{A}$ all the basic operations of the term algebra with the exception of the basic operation $Q^{\mathbf{A}}$ associated with the operation symbol Q. We define that via

$$Q^{\mathbf{A}}(t_0, t_1, \ldots, t_{r-1}) := \begin{cases} v_0 & \text{if } Qt_0t_1\ldots t_{r-1} = s_n(u) \text{ for some } n \notin X \text{ and for some term } u \\ u & \text{if } Qt_0t_1\ldots t_{r-1} = s_n(u) \text{ for some } n \in X \text{ and for some term } u \\ Qt_0t_1\ldots t_{r-1} & \text{otherwise} \end{cases}$$

where r is the rank of Q and $t_0, t_1, \ldots, t_{r-1}$ are any terms. Now the crucial contention in the proof of Claim 0 above will hold in present case because none of the proper subterms of s_n are unbalanced to right—that is if p is a proper subterm of s_n that is not a variable, then no substitution instance of p is also a substitution

instance of s_m for any natural number m. This ensure that $p^{\mathbf{A}}$ is always evaluated as if it were being evaluated in the term algebra. In this way, a version of Claim 0 for the present case can be established and, with it, the remainder of the proof goes as before. □

6.2 Sublattices of the Lattice of Equational Theories

LESSON 7

SECOND INTERLUDE: THE RUDIMENTS OF COMPUTABILITY

LESSON 8

UNDECIDABILITY IN EQUATIONAL LOGIC

8.1 A FINITELY BASED UNDECIDABLE EQUATIONAL THEORY

Let T be an equational theory. How can we decide whether $s \approx t \in T$? If writing such a program is possible, we say T has a decidable equational theory. Sad to say, but most of the equational theories are undecidable.

Definition. An equational theory T is **decidable** provided it is possible to write a computer program to do the job. If this is not possible, we say that T is **undecidable**.

This idea really only works if you have a friendly-enough signature. If your signature is uncountable, for instance, you'd run into problems.

In order to formalize what it means to be a "computer program", we're going to talk about Turing Machines.

Definition. A **Turing Machine** is a device consisting of a finite tape alphabet Σ, a finite set of states Q, and a transition process. Given $a \in \Sigma$ and $q \in Q$, there is only one instruction $\langle a, q, \ldots \rangle$. The Turing Machine is just a finite list of 5-tuples subject to the restriction just mentioned.

Theorem 8.1.1 (The Halting Problem; Turing, Post, Kleene 1936, Markov 1944)**.**

- *There is a Turing Machine M so that the set of strings for which M halts is undecidable.*
- *The set of Turing Machines which halt when launched on the empty string (blank tape) is undecidable.*

We begin with some more about Turing Machines. We'll use the symbol $\flat$ ("flat") for the blank square. We will make the signature as follows:

- F_a will be a unary operation symbol for each tape symbol a;
- G_q will be a unary operation symbol for each state q;
- For housekeeping purposes, we'll make unary operation symbols H, K, and J.

We're going to make a finite set of equations Σ that will completely describe our Turing machine. We'll write

$$F_c G_q F_a x \approx G_r F_c F_b x.$$

This goes into Σ_m for every c in the tape alphabet.

Given an word w on the tape alphabet, there is a term t_w which looks like

$$HG_{q_0}F_aF_bF_aF_aF_b\cdots Hx$$

where q_0 is the initial state and w is $abaab\cdots$ corresponding to the subscripts of the F's.

Definition. A **Post** term looks like

$$HATBHx$$

where A and B are (possibly empty) strings of F_c's for various c's and T is one of J, K, G_q where q is a state.

Theorem 8.1.2. *Let M be a Turing Machine and w be a word on the tape alphabet of M. Then M halts on input w if and only if*

$$\Sigma_M \vdash t_w \approx HJHx$$

Proof. ($\Rightarrow$) Clear from the construction. If we're in state q_0 reading the letter a, and our instruction set has aq_0cRr, we overwrite a with c, move to the Right, and become state r. As far as the terms go, we're starting with

$$HG_{q_0}F_aF_b\cdots Hx$$

and we want to get

$$HF_cG_rF_b\cdots Hx.$$

We'll do this using the equation from our set of axioms that says

$$G_{q_0}F_aF_b \approx F_cG_rF_bx.$$

($\Leftarrow$) We begin in the middle of the proof from last time. We were still trying to reject the two cases that we had before.

Case: t is $HQt_0\cdots t_{r-1}$, where H is some string of unary operation symbols. ✓

We also need to see that d and $\hat{d}$ are ω-universal. Let $f, g : \omega \to \omega$. Since d^∂ and $\hat{d}^\partial$ are ω-universal, we can pick $\mathbf{A}$ so that

$$(d^\partial)^{\mathbf{A}} = f$$

and

$$(\hat{d}^\partial)^{\mathbf{A}} = g.$$

Expand $\mathbf{A}$ to $\mathbf{A}^*$ by making all the unary operation symbols name the identity function. So

$$d^{\mathbf{A}^*} = (d^\partial)^{\mathbf{A}} = f$$

and

$$\hat{d}^{\mathbf{A}^*} = \ldots = g$$

So d and $\hat{d}$ are ω-universal. □

8.2 ω-UNIVERSAL SYSTEMS OF DEFINITIONS

The finitely based undecidable equational theory described in the last section had a signature consisting of some large finite number of unary operation symbols. In this section we show how to construct such equational theories in any signature with at least two distinct unary operation symbols or with some operation symbol of rank at least two.

Let σ denote a signature with exclusively unary operation symbols. We suppose σ provides only countably many (perhaps finitely many) operation symbols $F_0, F_1, F_2, \ldots$. Let τ be some other countable signature.

Suppose that a system $d = \langle d_0, d_1, d_2, \ldots \rangle$ of distinct terms of signature τ has been given so that x is the only variable to occur in each of the d_i's. We call d a **system of definitions for σ in τ**. The system d offers us the means to translate terms of signature σ into terms of signature τ. We define $\mathrm{Tr}_s : Te_\sigma \to Te_\tau$ by the following recurision:

(a) When t is variable, let $\mathrm{Tr}_d(t) = t$.

(b) When $t = F_i \hat{t}$,let $\mathrm{Tr}_d(t) = d_i(\mathrm{Tr}_d(\hat{t}))$.

We apply Tr_d to equations by taking $\mathrm{Tr}_d(s \approx t)$ to be $\mathrm{Tr}_d(s) \approx \mathrm{Tr}_d(t)$. We apply Tr_d to sets Σ of equations by putting

$$\mathrm{Tr}_d(\Sigma) = \{\mathrm{Tr}_d(s \approx t) \mid s \approx t \in \Sigma\}.$$

Let $\mathbf{A}$ be an algebra of signature τ and d be a system of definitions for σ in τ. Let $\mathbf{A}(\leftarrow d)$ be the algebra of signature σ with universe A so that

$$F_i^{\mathbf{A}(\leftarrow d)} = d_i^{\mathbf{A}}$$

for all i. A straight forward induction of the complexity of a term t of signature σ shows

$$t^{\mathbf{A}(\leftarrow d)} = (\mathrm{Tr}_d(t))^{\mathbf{A}}.$$

So for any equation $s \approx t$ of signature σ we have

$$\mathbf{A}(\leftarrow d) \models s \approx t \text{ if and ony if } \mathbf{A} \models \mathrm{Tr}_d(s \approx t).$$

It is easy, as we shall see below, to prove that if $\Sigma \vdash s \approx t$ in signature σ, then $\mathrm{Tr}_d(\Sigma) \vdash \mathrm{Tr}_d(s \approx t)$ in signature τ. However, in general the converse fails without some further restriction on the system d of terms.

We will say that the system d of distinct terms of signature τ is ω-**universal** provided for every system $\langle g_0, g_1, g_2, \ldots \rangle$ of functions from ω into ω there is an algebra A of signature τ with universe ω such that $d_i^{\mathbf{A}} = g_i$ for all i.

The ω-Universal Translation Theorem.
Let σ be a countable unary signature and let τ be a countable signature. Let d be an ω-universal system of definitions for σ in τ. Then for any set $\Sigma \cup \{s \approx t\}$ of equations of signature σ,

$$\Sigma \vdash s \approx t \textit{ if and only if } \mathrm{Tr}_d(\Sigma) \vdash \mathrm{Tr}_d(s \approx t).$$

Proof. First, let us suppose that $\Sigma \vdash s \approx t$ and deduce that $\mathrm{Tr}_d(\Sigma) \vdash \mathrm{Tr}_d(s \approx t)$. So suppose that $\mathbf{A} \models \mathrm{Tr}_d(\Sigma)$. Then $\mathbf{A}(\leftarrow d) \models \Sigma$. So also $\mathbf{A}(\leftarrow d) \models s \approx t$. But then $\mathbf{A} \models \mathrm{Tr}_d(s \approx t)$ as desired.

Now let us suppose that $\Sigma \nvdash s \approx t$ and deduce that $\mathrm{Tr}_d(\Sigma) \nvdash \mathrm{Tr}_d(s \approx t)$. Let $\mathbf{B}\langle \omega, g_0, g_1, g_2, \ldots \rangle$ be a model of Σ in which $s \approx t$ fails. Let $\mathbf{A}$ be an algebra of signature τ with universe ω such that $d_i^{\mathbf{A}} = g_i$ for all i. Such an algebra exists since d is ω-universal. Then it is clear that $\mathbf{A}(\leftarrow d) = \mathbf{B}$. Consequently, $\mathbf{A}$ is a model of $\mathrm{Tr})d(\Sigma)$ in which $\mathrm{Tr}_d(s \approx t)$ fails, as desired. □

Of course the usefulness of this theorem depends on the availability of an ω-universal system d of definition for σ in τ.

Definition. A set Δ of terms is **nonoverlapping** provided

- No variable belongs to Δ; and
- For all $d, d' \in \Delta$ and p a subterm of d and $u_0, u_1, \ldots$ and $v_0, v_1, \ldots$ terms, if

$$p(u_0, u_1, \ldots) = d'(v_0, v_1, \ldots)$$

then $p = d = d'$ and $u_i = v_i$ for all i occurring in d.

Theorem 8.2.1. *Every nonoverlapping set of terms is ω-universal.*

Proof. Let Δ be a set of nonoverlapping terms. Let F assign to each term d in Δ a function $F_d : \omega^\omega \to \omega$ so that F_d depends on coordinate i only if x_i occurs in d. We want:

> "There is an algebra $\mathbf{A}$ with universe ω (or the set of terms T) so that $F_d = d^{\mathbf{A}}$ for all $d \in \Delta$."

This is what it meant to be ω-universal. It meant that when you took an assignment that assigned to each term in Δ a function that had any chance at all of being a term function, then there was an algebra that would make it work.

Let Q be an operation symbol and let r be the rank of Q. Then for terms t_i, we have

$$Q^{\mathbf{A}}(t_0, \ldots, t_{r-1}) := \begin{cases} F_d(\bar{u}) & \text{if } Qt_0 \cdots t_{r-1} = d(u_0, u_1, \ldots) \text{ for some } d \in \Delta \text{ and some } \bar{u} \in T^\omega \\ Qt_0 \cdots t_{r-1} & \text{otherwise} \end{cases}$$

Is this definition okay? It looks like there are only two cases, but there really are a bunch more: there's a case for each $d \in \Delta$ and Δ could be really big. In other words, we need to make sure that our cases do not conflict.

Suppose $d, d' \in \Delta$ and $\bar{u}, \bar{v} \in T^\omega$ so that

$$d(\bar{u}) = Qt_0 \ldots t_{r-1} \text{ and } d'(\bar{v}) = Qt_0 \ldots t_{r-1}$$

that is, $d(\bar{u}) = d'(\bar{v})$. By nonoverlapping of Δ, we get $d' = d$ and $u_i = v_i$ if x_i occurs in d. So $F_d(\bar{u}) = F_{d'}(\bar{v})$.

Our contention:

> "If $d \in \Delta$ and s is a proper subterm of d, then $s^{\mathbf{A}}(u_0, u_1, \ldots) = s(u_0, u_1, \ldots) = s^{\mathbf{T}}(u_0, u_1, \ldots)$ for all $\bar{u} \in T^\omega$."

We will prove this by induction on the complexity of s. The base step is too easy.

For the inductive steps, say $s = Qs_0 \cdots s_{r-1}$. Then

$$\begin{aligned} s^{\mathbf{A}} &= Q^{\mathbf{A}}\left(s_0^{\mathbf{A}}(\bar{u}), \ldots, s_{r-1}^{\mathbf{A}}(\bar{u})\right) \\ &= Q^{\mathbf{A}}(s_0(\bar{u}), \ldots, s_{r-1}(\bar{u})) \\ &= Qs_0(\bar{u}) \cdots s_{r-1}(\bar{u}) = s(\bar{u}) \end{aligned}$$

So $Q^{\mathbf{A}}(s_0(\bar{u}), \ldots, s_{r-1}(\bar{u})) = s(\bar{u}) = s^{\mathbf{A}}(\bar{u})$.

Now we claim that $F_d = d^{\mathbf{A}}$ for all $d \in \Delta$. To prove this, let $d \in \Delta$ and $\bar{u} \in T^\omega$. Well,

$$\begin{aligned} d^{\mathbf{A}}(\bar{u}) &= Q^{\mathbf{A}}\left(t_0^{\mathbf{A}}(\bar{u}), \ldots, t_{r-1}^{\mathbf{A}}(\bar{u})\right) \\ d = Qt_0 \cdots t_{r-1} &= Q^{\mathbf{A}}\left(t_0^{\mathbf{A}}(\bar{u}), \ldots, t_{r-1}^{\mathbf{A}}(\bar{u})\right) \\ &= \bar{T}_d(\bar{u}) \end{aligned}$$

So we'll just test

$$Qt_0(\bar{u}) \cdots t_{r-1}(\bar{u}) = d(\bar{u})$$

□

Example. Let B and D be two binary operation symbols. Consider the term $DDxyBxy$. It looks like

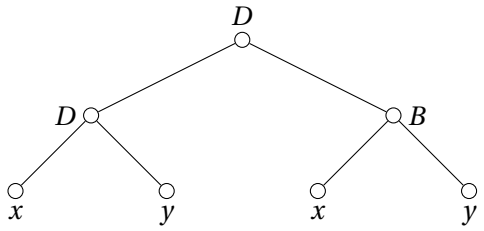

Notice that Dxy is a proper subterm of d. But

$$x \mapsto Dxy \text{ and } y \mapsto Bxy$$

which sends $D(Dxy)(Bxy)$ to d.

We claim that $\{d\}$ is ω-universal but fails to be nonoverlapping.

Let $F : \omega \times \omega \to \omega$. Define $\mathbf{A}$ so that $A = \omega$ and

$$B^{\mathbf{A}} = F(a,b) \text{ and } D^{\mathbf{A}}(a,b) = b$$

Notice that

$$\begin{aligned} d^A(a,b) &= D^A\left(D^A(a,b), B^A(a,b)\right) \\ &= D^A\left(b, F(a,b)\right) \\ &= F(a,b) \end{aligned}$$

So $d^{\mathbf{A}} = F$.

Example. Let Q be a binary operation symbol. Let

$$d_n = QxQ^n x^{n+1} = QxQQ\cdots Qxxx\cdots x$$

We claim that $\{d_n : n \in \omega\}$ is nonoverlapping and hence is ω-universal. So let s be a subterm of d_n with $s \neq x$. Suppose u and v are terms and $m \in \omega$ and $s(u) = d_m(v)$. To have nonoverlapping, we need to check that $s = d_n = d_m$ and $u = v$.

There were some cases that went along with this that required rather intricate tree diagrams.

8.3 Base undecidability: the set up

Here we will use σ to denote a fixed signature which provides just two operation symbols D and E, both unary. We reserve Ψ to denote a finite set of equations in the variable x of signature σ that is the base of an undecidable equational theory. Let H and K be two new unary operation symbols.

Now let τ be another signature and let $d = \langle d_D, d_E, d_H, d_K\rangle$ be a system of definitions for σ in τ.

For any equation $s \approx t$ of signature σ we put

$$\Psi(s \approx t) := \Psi \cup \{Hs(Kx) \approx Hs(Ky), Ht(Kx) \approx x\}.$$

Lemma.
Let s, t, r, and q be terms of signature σ in the variable x such that $\Psi \not\vdash s \approx t$. Then

$$\Psi(s \approx t) \vdash q \approx r \text{ if and only if } \Psi \vdash q \approx r.$$

Conseuqently, if s and t are terms of signature σ in the variable x, then

$$\Psi(s \approx t) \vdash s \approx t \text{ if and only if } \Psi \vdash s \approx t.$$

Proof. The implication from right to left is clear. So suppose that $\Psi \nvdash q \approx r$ and $\Psi \nvdash s \approx t$. We have to show $\Psi(s \approx t) \nvdash q \approx r$. Let $\mathbf{A}$ be a countably infinite model of Ψ in which both $q \approx r$ and $s \approx t$ fail. Pick $a, b \in A$ so that $q^{\mathbf{A}}(a) \neq r^{\mathbf{A}}(a)$ and $s^{\mathbf{A}}(b) = c \neq d = t^{\mathbf{A}}(b)$. Now $\mathbf{A}$ is an algebra of signature σ. We start by improving $\mathbf{A}$ so that $s \approx t$ will have, loosely speaking, infinitely many disjoint failures. As a first step, let $\infty \notin A$ and put $A' = A \cup \{\infty\}$. Make $\mathbf{A}'$ by extending the operations of $\mathbf{A}$ so that $D^{\mathbf{A}'}\infty = \infty = E^{\mathbf{A}'}(\infty)$. Since Ψ consists of regular equations we will have $\mathbf{A}' \models \Psi$. Let

$$B = \{\bar{e} \mid \bar{e} \in A'^{\omega} \text{ and at most one coordinate of } \bar{e} \text{ differs from } \infty\}.$$

It is easy to see that B is a subuniverse of $\mathbf{A}'^{\omega}$ and that B is countably infinite. Let $\mathbf{B}$ be the subalgebra of $\mathbf{A}'^{\omega}$ with universe B. So we have $\mathbf{B} \models \Psi$.

For $e \in A$ and $i \in \omega$, let $e[i]$ be the ω-tuple with e at the i^{th} position and ∞ at all the other positions. So we have $s^{\mathbf{B}}(b[i]) = c[i] \neq d[i] = t^{\mathbf{B}}(b[i])$ for all $i \in \omega$.

Now let $\mathbf{B}^*$ be the expansion of $\mathbf{B}$ by taking $K^{\mathbf{B}^*}$ to be a one-to-one function from B onto $\{b[i] \mid i \in \omega\}$ and by defining $H^{\mathbf{B}^*}$ as follows:

$$H^{\mathbf{B}^*}(u) = \begin{cases} (K^{\mathbf{B}^*})^{-1}(b[i]) & \text{if } u = d[i] \text{ for some } i \in \omega \\ d[0] & \text{otherwise} \end{cases}$$

We contend that $\mathbf{B}^* \models \Psi(s \approx t)$. We already have that $\mathbf{B}^* \models \Psi$, so what we need is that $\mathbf{B}^* \models Hs(Kx) \approx Hs(Ky)$ and $\mathbf{B}^* \models Ht(Kx) \approx x$.

Consider the first equation. Because $K^{\mathbf{B}^*}$ produces only the values $b[i]$ for various $i \in \omega$ and because $s^{\mathbf{B}^*}(b[i]) = c[i]$, we will end up evaluating $H^{\mathbf{B}^*}(c[i])$. But sinces $c[i] \neq d[j]$ for any choices of i and j, we see that $H^{\mathbf{B}^*}(c[i]) = d[0]$, no matter what value i has. Thus both sides of $Hs(Kx) \approx Hs(Ky)$ evaluate to $d[0]$ and the equation holds.

Now consider the second equation. Let $u \in B^*$. Pick $i \in \omega$ so that $K^{\mathbf{B}^*}(u) = b[i]$. Then $t^{\mathbf{B}^*}(b[i]) = d[i]$. Now the definition of $H^{\mathbf{B}^*}$ gives us $H^{\mathbf{B}^*}(d[i]) = u$. Putting this together, we have $\mathbf{B}^* \models Ht(Kx) \approx x$, as desired. □

The Base Undecidability Lemma.
Let τ be any recursive signature. Let Γ be any finite set of equations of signature τ. If there are two distinct terms d and $\hat{d}$ in the variable x such that

- *$\{d, \hat{d}\}$ is ω-universal.*
- *$\Gamma \vdash d \approx x, \hat{d} \approx x$.*

Then the equational theory based on Γ is base undecidable.

Proof. First put

$$\begin{aligned} d_D &:= d^2(\hat{d}(d(\hat{d}))) \\ d_E &:= d^2(\hat{d}^2(d(\hat{d}))) \\ d_H &:= d^2(\hat{d}^3(d(\hat{d}))) \\ d_K &:= d^2(\hat{d}^4(d(\hat{d}))) \end{aligned}$$

It routine to verify that $\Gamma \vdash d_D \approx x, d_E \approx x, d_H \approx x, d_K \approx x$. That $\{d_D, d_E, d_H, d_K\}$ is ω-universal follows from the fact that $\{d, \hat{d}\}$ is ω-universal and the fact that $\{F^2G^{k+1}FGx \mid k \in \omega\}$ is nonoverlapping and hence ω-universal. To simpify notation we use Tr to denote the translation function with respect to $\{d_D, d_E, d_H, d_K\}$.

We put

$$B(\Gamma, s \approx t) := \mathrm{Tr}(\Psi) \cup \{\mathrm{Tr}(Hs(Kx))(p) \approx \mathrm{Tr}(Hs(Kx))(q) \mid p \approx q \in \Gamma\} \cup \{\mathrm{Tr}(Ht(Kx)) \approx x\}$$

We reduce the decision problem for the equational theory based on Ψ to the base decidability problem for the equational theory based on Γ. What we need is to show, for each equation $s \approx t$ of signature σ, where x is the only variable to occur in $s \approx t$ and neither s nor t is itself a variable,

$$\Psi \vdash s \approx t \text{ if and only if } B(\Gamma, s \approx t) \text{ and } \Gamma \text{ are bases for the same equational theory.}$$

By the previous lemma and the ω-Universal Translation Lemma, we have

$$\begin{array}{ccc} \Psi \vdash s \approx t & \iff & \Psi(s \approx t) \vdash s \approx t \\ \Updownarrow & & \Updownarrow \\ \mathrm{Tr}(\Psi) \vdash \mathrm{Tr}(s) = \mathrm{Tr}(t) & \iff & \mathrm{Tr}(\Psi(s \approx t)) \vdash \mathrm{Tr}(s) \approx \mathrm{Tr}(t) \end{array}$$

First, let us suppose that $\Psi \vdash s \approx t$. Then $\mathrm{Tr}(\Psi) \vdash \mathrm{Tr}(s) \approx \mathrm{Tr}(t)$. Consequently, $\mathrm{Tr}(\Psi) \vdash \mathrm{Tr}(Hs(Kx)) \approx \mathrm{Tr}(Ht(Kx))$. From the definition of $B(\Gamma, s \approx t)$ we find that $B(\Gamma, s \approx t) \vdash \Gamma$. On the other hand, since $\Gamma \vdash d_Q \approx x$ for all $Q \in \{D, E, H, K\}$, we have that $\Gamma \vdash B(\Gamma, s \approx t)$. So we conclude that Γ and $B(\Gamma, s \approx t)$ are bases for the same equational theory.

Now let us suppose that $\Psi \nvdash s \approx t$. So $\mathrm{Tr}(\Psi(s \approx t)) \nvdash \mathrm{Tr}(s) \approx \mathrm{Tr}(t)$. All the equations in $B(\Gamma, s \approx t)$ are substitution instances of equations in $\mathrm{Tr}(\Psi(s \approx t))$, so we have that $\mathrm{Tr}(\Psi(s \approx t)) \vdash B(\Gamma, s \approx t)$. So $B(\Gamma, s \approx t) \nvdash \mathrm{Tr}(s) \approx \mathrm{Tr}(t)$. But we know that $\Gamma \vdash \mathrm{Tr}(s) \approx \mathrm{Tr}(t) \approx x$. This means that $B(\Gamma, s \approx t)$ and Γ cannot be bases for the same equational theory. □

8.4 The Base Undecidability Theorem

The Base Undecidability Theorem.
Let T be a finitely based equational theory in a recursive signature such that there is a term t such that $t \approx x \in T$ and either two different unary operation symbols occur in t or some operation symbol of rank at least 2 occurs in t. Then T is base undecidable.

Proof. We can assume that x occurs in t and x is the only variable which occurs in t. We need to construct two distinct terms d and $\hat{d}$ so that $\{d, \hat{d}\}$ is ω-universal and $t \approx x \vdash d \approx x, \hat{d} \approx x$.

We first note that we may assume that x occurs in t. Otherwise $t \approx x \vdash x \approx y$. In this case every equation belongs to T so we may select a suitable replacement for t. We may also suppose that x is the only variable to occur in t, since it does no harm to substitute x for every variable the the equation $t \approx x$.

We consider three cases depending on the term t.

CASE: t **has only unary operation symbols.**
The term t is $F^{n+1}GMF^kx$ where F and G are unary operation symbols, M is a (possibly empty) string of unary operation symbols not ending in F, and n and k are natural numbers.

Claim. $F^{n+1}GMF^kx \approx x \vdash F^{n+k+1}GMx \approx x$.

Proof. Suppose $k > 0$, as there is nothing to prove otherwise. In any model of the equation $F^{n+1}GMF^kx \approx x$ the function denoted by F must be bijective and so it is invertible. Hence F^k denotes an invertible function, as well. Thus, $F^{n+k+1}GMx \approx x$ must also be true in the model. □

This claim allows us to assume, without loss of generality, that $k = 0$ and t is $F^{n+1}GMx$ where M is a (possibly empty) string of unary operation symbols not ending in F. Let $p-1$ be the number of occurences of F in M. Then define

$$d := (F^{n+1})^{3p}(GM)^{3p}x$$
$$\hat{d} := (F^{n+1})^{2p}(GM)^{2p}(F^{n+1})^{p}(GM)^{p}x$$

It is evident that $t \approx x \vdash d \approx x, \hat{d} \approx x$. To see that d and $\hat{d}$ are nonoverlapping, observe that d begins with a string of $3p(n+1)$ occurences of F, but after this initial string no strings of consecutive F's can be longer than $p-1$. Likewise $\hat{d}$ begins with a string of $2p(n+1)$ occurences of F, but after this there is one string of $p(n+1)$ consecutive F's. Elsewhere in $\hat{d}$ no string of consecutive F's is longer than $p-1$. Thus no proper subterm q of d such that q is not x can have a substitution instance $q(u)$ which is of the form $d(w)$ or $\hat{d}(w)$. And likewise for any proper subterm of $\hat{d}$. It is also plain that $d(u)$ and $\hat{d}(w)$ are never the same, regardless of how u and w are chosen. So d and $\hat{d}$ are nonoverlapping and therefore also ω-universal.

CASE: t **begins with an operation symbol of rank at least** 2**.**
We suppose that t is $Qt_0t_1\dots t_{r-1}$ where r is the rank of Q. So $r \geq 2$. Without loss of generality we will suppose that x occurs in t_0.

Our plan is to construct to terms from t which which will be nonoverlapping or at least ω-universal. We will need some easily established symbol counting principles. We are able to limit our attention to those terms in which no variable other than x occurs. For a term q we use $|q|_x$ to denote the number of times the variable x occurs in q and $|q|_c$ to denote the numberof occurences of constant symbols in q. We use $q(w)$ to denote the result of substituting the term w for the variable x in q. Last q^n is defined recursively by $q^0 = x$ and $q^{n+1} = q(q^n)$. The following are easily established:

$$|q(w)|_x = |q|_x|w|_x$$
$$|q^n|_x = |q|_x^n$$
$$|q(w)|_c = |q|_c + |q|_x|w|_c$$
$$|q^n|_c = \begin{cases} |q|_c & \text{if } |q|_x = 0 \\ n|q|_c & \text{if } |q|_x = 1 \\ \frac{|q|_x^n - 1}{|q|_x - 1}|q|_c & \text{if } |q|_x > 1. \end{cases}$$

Let $\bar{n}$ be any r-tuple of natural numbers. We say that d is the **associate** of $Qt_0t_1\dots t_{r-1}$ **corresponding to** $\bar{n}$ provide d is $Qt^{n_0}(t_0)t^{n_1}(t_1)\dots t^{n_{r-1}}(t_{r-1})$. We see immediately that $t \approx x \vdash s \approx x$ for every associate d of t. Our plan is to produce two nonoverlapping associates of t by chosing the r-tuples of natural numbers with care.

Let $\bar{n}'$ be any r-tuple of natural numbers and let $\hat{d}$ be the associate of t corresponding to $\bar{n}'$. Suppose q is a proper subterm of $\hat{d}$ and q is not a variable. So q is a subterm of $t^{n_i'}(t_i)$ for some $i < r$. It follows that q is a substitution instance of a subterm of t. Now let d be the associate of t corresponding to the r-tuple $\bar{n}$. Below we will have to exclude the possibility that there are terms w and u so that $q(w) = d(u)$. We can achieve this by instead letting q range over the (finite) collection of subterms of t which are not variables. It is also clear that since d is an associate of t we need only consider those subterms q of t where $q = Qq_0q_1\dots q_{r-1}$.

We consider two subcases.

SUBCASE: NO CONSTANT SYMBOLS OCCUR IN t.
We need to find two r-tuple $\bar{n}$ and $\bar{n}'$ such that the corresponding associates d and $\hat{d}$ of t are nonoverlapping. There are two situations which must be rejected:

(a) q is a proper subterm of $\hat{d}$ which is not a variable and $q(w) = d(u)$ for some terms w and u.

(b) $\hat{d}(u) = d(w)$ for some terms w and u.

In either of these situations, if constant symbols occurred in w or u we could change all the constant symbols to the variable x and obtain a situation without constants. So we assume no constants occur. Moreover, in situation (a) we can instead consider that q is a subterm of t of the form $Qq_0 \dots q_{r-1}$.

Consider situation (a). For each $i < r$ we have

$$\begin{aligned} |q_i(w)|_x &= |t^{n_i}(t_i(u))|_x \\ |q_i|_x|w|_x &= |t|_x^{n_i}|t_i|_x|u|_x \\ \frac{|q_0|_x|w|_x}{|q_1|_x|w|_x} &= \frac{|t|_x^{n_0}|t_0|_x|u|_x}{|t|_x^{n_1}|t_1|_x|u|_x} \\ \frac{|q_0|_x}{|q_1|_x} &= \frac{|t|_x^{n_0}|t_0|_x}{|t|_x^{n_1}|t_1|_x} \\ \frac{|q_0|_x|t_1|_x}{|q_1|_x|t_0|_x}|t|_x^{n_1} &= |t|^{n_0} \end{aligned}$$

In this subcase $|t|_x \geq 2$, so let m be a natural number large enough so that

$$\frac{|q_0|_x|t_1|_x}{|q_1|_x|t_0|_x}|t|_x^2 < |t|^m$$

for all subterms q_0 and q_1 of t. Here are the desired r-tuples:

$$\begin{aligned} \bar{n} &= \langle m+1, 1, 1, \dots, 1\rangle \\ \bar{n}' &= \langle m, 2, 1, \dots, 1\rangle \end{aligned}$$

With these choices, the situation (a) is rejected.

Now consider situation (b). With our $\bar{n}$ and $\bar{n}'$ were $\hat{d}(u) = d(w)$ we would have

$$\begin{aligned} t^m(t_0(u)) &= t^{m+1}(t_0(w)) \\ t^2(t_1(u)) &= t^1(t_1(w)) \end{aligned}$$

Since $t^{m+1}(t_0)$ is longer than $t^m(t_0)$, the first equation entails that u is longer than w. However the second equation entails that u is shorter than w, which is impossible since $|t|_x \geq 2$. So the associates $\hat{d}$ and d of t are indeed nonoverlapping, as desired.

SUBCASE: CONSTANT SYMBOLS OCCUR IN t.
As in the last subcase, we need to find two r-tuple $\bar{n}$ and $\bar{n}'$ such that the corresponding associates d and $\hat{d}$ of t are nonoverlapping. Again, there are two situations which must be rejected:

(a) q is a proper subterm of $\hat{d}$ which is not a variable and $q(w) = d(u)$ for some terms w and u.

(b) $\hat{d}(u) = d(w)$ for some terms w and u.

Moreover, in situation (a) we can instead consider that q is a subterm of t of the form $Qq_0 \dots q_{r-1}$. Recall that x occurs in t and we have assumed (without loss of generality) that x occurs in t_0.

Since x occurs in t we know that $|t^k(t_0)|_c$ and $|t^k(t_1)|_c$ are strictly increasing functions of k. Pick m so large that

$$|q|_c < |t^m(t_0)|_c \text{ and } |q|_c < |t^m(t_1)|_c \text{ for all subterms } q \text{ of } t.$$

Next pick ℓ so large that

$$|t^\ell(t_0)|_c > \frac{|q_0|_x}{|q_1|_x}|t^{m+1}(t_1)|_c + |q_0|_c$$
$$|t^\ell(t_0)|_x > \frac{|q_0|_x}{|q_1|_x}|t_x^{m+1}|t_1|_x$$

for all subterms q_0 and q_1 of t such that $|q_1|_x \neq 0$.

We take

$$\bar{n} := \langle \ell+1, m, 1, \dots, 1\rangle$$
$$\bar{n}' := \langle \ell, m+1, 1, \dots, 1\rangle$$

and let d and $\hat{d}$ be the associates of t corresponding to $\bar{n}$ and $\bar{n}'$.

Consider situation (a). We have

$$|q_0(w)|_c = |t^{\ell+1}(t_0(u))|_c$$
$$|q_1(w)|_c = |t^m(t_1(u))|_c$$
$$|q_0|_c + |q_0|_x|w|_c = |t^{\ell+1}(t_0)|_c + |t^{\ell+1}(t_0)|_x|u|_c$$
$$|q_1|_c + |q_1|_x|w|_c = |t^m(t_1)|_c + |t^m(t_1)|_x|u|_c$$

In the event that $|q_1|_x = 0$ we get

$$|q_1|_c = |t^m(t_1)|_c + |t^m(t_1)|_x|u|_c$$

But this violates the choice of m. So $|q_1|_x \neq 0$, allowing us to solve the next to the last displayed equation for $|w|_c$. After some manipulation, we obtain

$$\frac{|q_0|_x}{|q_1|_x}|t^m(t_0)|_c + |q_0|_c + \frac{|q_0|_x}{|q_1|_x}|t^m(t_1)|_x|u|_c = |t^{\ell+1}(t_0)|_c + |t^{\ell+1}(t_0)|_x|u|_c + \frac{|q_0|_x}{|q_1|_x}|q_1|_c.$$

But this violates the choice of ℓ. A similar analysis works with $\hat{d}$ in place of d. So the situation (a) is rejected.

Situation (b) is rejected in this subcase the same way it was rejected in the first subcase. Some we see that the associates d and $\hat{d}$ of t are nonoverlapping.

CASE: t **begins with a unary operation symbol and has operation symbols of rank at least** 2**.**
In this case, it turns out not to be possible always to obtain two associates of t which are nonoverlapping. But we can still find two associates of t which are ω-universal.

For any term q let q^∂ be the term obtained from q by deleting all the unary operation symbols. Now t^∂ is a term that falls into the last subcase. Let $\bar{n}$ and $\bar{n}'$ be two distinct r-tuples whose corresponding associates are nonoverlapping. Let d and $\hat{d}$ be the associates of t corresponding to $\bar{n}$ and $\bar{n}'$ respectively. Then d^∂ and $\hat{d}^\partial$ are the corresponding associates of t^∂. To see that d and $\hat{d}$ are ω-universal, let $f : \omega \to \omega$ and $g : \omega \to \omega$

be any functions on ω. Since d^∂ and $\hat{d}^\partial$ are nonoverlapping, they are ω-universal. Let $\mathbf{A}$ be an algebra with universe ω and a basic operation for each operation symbol in t^∂ so that $(d^\partial)^{\mathbf{A}} = f$ and $(\hat{d}^\partial)^{\mathbf{A}} = g$. Expand $\mathbf{A}$ to $\mathbf{B}$ by interpreting all the unary operation symbols as the identity function on ω. Then $d^{\mathbf{B}} = f$ and $\hat{d}^{\mathbf{B}} = g$, establishing that d and $\hat{d}$ are ω-universal associates of t.

□

LESSON 9

RESIDUAL BOUNDS

Let $\mathcal{K}$ be a class of algebras of the same signature. The **residual bound** of $\mathcal{K}$ is the least cardinal κ so that every subdirectly irreducible algebra in $\mathcal{K}$ has fewer than κ elements. We also refer to this cardinal as the **residual character** of $\mathcal{K}$. If no such κ exists, then we say $\mathcal{K}$ is **residually large** and we put ∞ as the residual bound. Most varieties are residually large.

We say $\mathcal{K}$ is **residually finite** if the residual bound κ is either finite or equals ω. That is, $\mathcal{K}$ is residually finite provided its residual bound is countable. We say $\mathcal{K}$ is **residually very finite** if the residual bound κ is finite. We say $\mathcal{K}$ is **residually small** if the residual bound κ for some cardinal κ.

The next two theorems are offered for information. Their proof lie beyond the scope of this exposition.

Theorem 9.0.1 (Walter Taylor, 1970)**.** *A variety $\mathcal{V}$ of countable signature is residually small if and only if $\mathcal{V}$ has residual bound of $\leq (2^{\omega})^{+}$.*

A variety with residual bound $\aleph_0$ has arbitrarily large finite subdirectly irreducible algebras but no infinite subdirectly irreducible algebra. A variety with residual bound $\aleph_1$ must have a countably infinite sudirectly irreducible algebra but no uncountable subdirectly irreducible algebra.

Theorem 9.0.2 (McKenzie and Shelah, 1970)**.** *Every variety of countable signature that has uncountable subdirectly irreducible algebras must have a subdirectly irreducible algebra of cardinality at least $2^{\aleph_0}$.*

Notice the following:

I. The variety of one-element algebras has no subdirectly irreducible algebra. This variety will have residual bound 0, since 0 is certainly larger than the cardinality of any subdirectly irreducible algebra in the variety.

II. 1 cannot be a residual bound of a variety, since this entails that the variety has a subdirectly irreducible algebra of cardinality 0. Algebras cannot have cardinality 0.

III. 2 cannot be the residual bound of a variety because there are no 1-element subdirectly irreducible algebras.

IV. There are varieties of countable signature with residual bounds of $3, 4, 5, \ldots, \aleph_0, \aleph_1, (2^{\aleph_0})^{+}$, as well as varieties that are residually large.

9.1 The Variety Generated by McKenzie's Algebra **R** is Resdually Large

Recall McKenzie's algebra from Section 4.5:

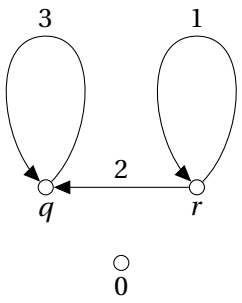

Figure 9.1: McKenzie's Automatic Algebra **R**

This is an automatic algebra. It has just one operation and that operation is a two-place operation, which we denote be juxtaposition. The elements of this algebra fall into the set $\{1,2,3\}$ of letters and the disjoint set $\{q,r\}$ of states, with an additional default element 0. The operation is defined so that

$$1r = r \qquad 2r = q \qquad 3q = q$$

with all other products resulting in 0. We proved that the variety generated by **R** contains a shift automorphism algebra. As a consequence, we know that $\mathcal{HSP}$ **R** has an infinite subdirectly irreducible algebra, according to the Shift Automophism Theorem. However it is easy to construct subdirectly irreducible algebras in this variety that have any cardinality greater than 1.

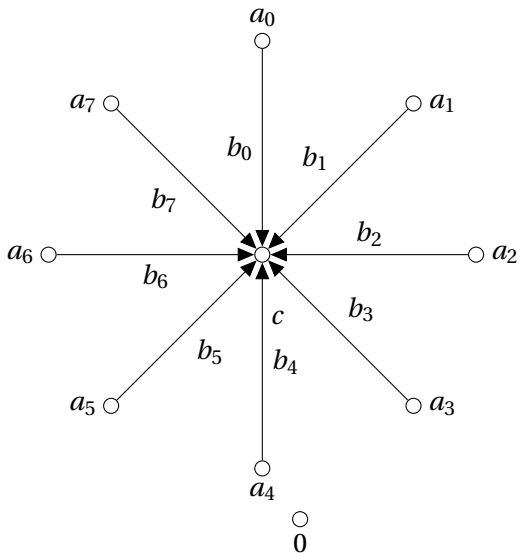

Figure 9.2: The Subdirectly Irreducible Algebra $\mathbf{S}_8$

Figure 9.2 gives a diagram of a subdirectly irreducible algebra with 18 elements in this variety and it is plainly possible to replace 8 with any cardinal greater than 1 and also to vary and elaborate this diagram to obtain quite complicated subdirectly irreducible algebras of arbitrary infinite cardinalities. The node in the middle of the diagram is c. Moreover, the operation is defined by following the arrows, just as it was for **R**.

Fact. $\mathbf{S}_8$ is subdirectly irreducible.

Proof. We argue that $(c,0)$ is a critical pair. What we must show is that if $u \neq v$ is S_8, then any congruence θ that collapse u and v must also collapse c and 0. Notice that $c = b_i a_i$ for all i and, moreover, that this product is the only nonzero product that involves either b_i or a_i. Without loss of generality, we can suppose that $u \neq 0$. There are three cases: $u = a_i$ for some i, $u = b_i$, for some i, and $u = c$. We consider just the first case and the last case, leaving the middle case in the hands of the diligent graduate students.

Case: $u = a_i$
We know that $c = b_i a_i$ and $0 = b_i v$, since $v \neq a_i = u$. Thus $c = b_i a_i \,\theta\, b_i v = 0$, as desired.

Case: $u = c$
If $v = 0$ there is nothing to prove. So consider that $c = u \neq v \neq 0$. So $v = a_i$ for some i or $v = b_i$ for some i. We only deal with first alternative. So we have $c = a_i b_i = v b_i \,\theta\, u b_i = c b_i = 0$, as desired. □

Fact. $\mathbf{S}_8 \in \mathcal{HSP}\,\mathbf{R}$.

Proof. We ω-tuples. α_i, β_i, and γ, for $i < 8$, like we did in Section 4.5. Were we to replace 8 by another cardinal κ, this construction can be modified (we would have to replace ω by a suitable ordered set). Actually, in the present case, we could replace ω by 8.

$$\begin{array}{lccccccccc} \alpha_i := & q & q & q & q & r & q & q & q & \ldots \\ \beta_i := & 3 & 3 & 3 & 3 & 2 & 3 & 3 & 3 & \ldots \\ \gamma := & q & q & q & q & q & q & q & q & \ldots \end{array}$$

where the r and the 2 occur at the i^{th} position. As in Section!4.5, we consider the subalgebra $\mathbf{B}$ of $\mathbf{R}^\omega$ generated by the tuples displayed above. The equivalence relation θ that lumps together all tuple in B that contain a 0 and isolates everthing else is a congruence. Moreover, $\mathbf{S}_8 \cong \mathbf{B}/\theta$. So $\mathbf{S}_8 \in \mathcal{HSP}\,(\mathbf{R})$. □

So $\mathcal{HSP}\,\mathbf{R}$ is residually large. To devise finite algebras that generate residually small varieties we need a method to limit the subdirectly irreducible algebras in the variety.

9.2 Finite Subdirectly Irreducibles Generated by Finite Flat Algebras

In this section we will suppose that $\mathbf{A}$ is a finite *flat* algebra (that is, an algebra among whose operations $\wedge$ and 0 can be found which provide the algebra with the structure of a meet-semilattice of height one with least element 0) and that $\mathbf{S}$ is a finite subdirectly irreducible algebra in the variety generated by $\mathbf{A}$.

Now according to Birkhoff's HSP Theorem, $\mathbf{S}$ will always arise as a quotient of some $\mathbf{B}$, which is in turn a subalgebra of $\mathbf{A}^T$ for some T. Since $\mathbf{S}$ is subdirectly irreducible, we know that there is a strictly meet irreducible $\theta \in \operatorname{Con}\mathbf{B}$ such that $\mathbf{S} \cong \mathbf{B}/\theta$. The restriction of strictly meet irreducibility means that there is a congruence μ of $\mathbf{B}$ so that any congruence of properly above θ must include μ. This μ is just the inverse image under that quotient map established by θ of the monolith of $\mathbf{S}$. It is more convenient to work with $\mathbf{B}$ than with $\mathbf{S}$. Since $\mathbf{S}$ is finite, we can choose T to be finite. Indeed, in this section we assume the following:

- $\mathbf{B} \subseteq \mathbf{A}^T$
- $\theta \in \operatorname{Con}\mathbf{B}$
- θ is strictly meet-irreducible in $\mathbf{Con}\,\mathbf{B}$.
- $\mathbf{S} \cong \mathbf{B}/\theta$
- T is as small as possible for representing $\mathbf{S}$ in this way.

In particular this last condition entails that if $t \in T$, then there must be $u, v \in B$ so that $(u, v) \notin \theta$ but $u(s) = v(s)$ for all $s \in T - \{t\}$. Our effort at understanding the finite subdirectly irreducible $\mathbf{S}$ is largely focussed on θ.

First, we locate an element in B element b_0 whose image in $\mathbf{S}$ will be part of a critical pair. Since $\mathbf{B}$ has a semilattice operation, there are elements $u, v \in B$ with $u < v$ and (u, v) critical over θ, that $(u, v) \in \mu$. Using the finiteness of B pick p to be minimal, with respect to the semilattice order, among all those $v \in B$ such that (u, v) is critical over θ for some $u < v$.

Fact 0. If $w < p$, then $(w, p) \notin \theta$.

Proof. Suppose $w < p$ but $w \; \theta \; p$. Pick $u < p$ with (u, p) critical over θ. Then $w = p \wedge w \; \varphi \; u \wedge w$, for all $\varphi \in \mathrm{Con}\,\mathbf{B}$ with $\theta < \varphi$. But this means that either $(w, u \wedge w) \in \theta$ or that $(w, u \wedge w)$ is critical over θ. So by the minimality of p, we have $u \wedge w \; \theta \; w$. But then $u = u \wedge p \; \theta \; u \wedge w \; \theta \; w \; \theta \; p$, contradicting $(u, p) \notin \theta$. □

Now for each $t \in T$ pick $(x, y) \in B^2 - \theta$ so that $x(t) \neq y(t)$ but $x(s) = y(s)$ for all $s \in T - \{t\}$. Pick $u < p$ so that (u, p) is critical over θ. So $(u, p) \in \theta \vee \mathrm{Cg}^{\mathbf{B}}(x, y)$. Then according to Mal'cev's Congruence Generation Theorem there is a finite sequence $e_0, e_1, \ldots, e_n$ of elements of B, of translations $\lambda_0, \ldots, \lambda_{n-1}$ of $\mathbf{B}$, and of two-element subsets $\{z_0, w_0\}, \ldots, \{z_{n-1}, w_{n-1}\}$ each belonging the $\theta \cup \{x, y\}$ such that

$$u = e_0 \quad \{e_i, e_{i+1}\} = \{\lambda_i(z_i), \lambda_i(w_i)\} \text{ for all } i < n \quad e_n = p.$$

But now, meeting every element in the sequence with p, we have

$$u = u \wedge p = e_0 \wedge p \quad \{e_i \wedge p, e_{i+1} \wedge p\} = \{\lambda_i(z_i) \wedge p, \lambda_i(w_i) \wedge p\} \text{ for all } i < n \quad e_n \wedge p = p \wedge p = p$$

Since $u < p$ there must be some $i < n$ so that $p \in \{\lambda_i(z_i) \wedge p, \lambda_i(w_i) \wedge p\}$ where $\lambda_i(z_i) \wedge p \neq \lambda_i(w_i) \wedge p$. Let χ_t denote the element of $\{\lambda_i(z_i) \wedge p, \lambda_i(w_i) \wedge p\}$ which is different from p. Evidently $\chi_t < p$. By Fact 0 we see that $(\chi_t, p) \notin \theta$. Hence, $(z_i, w_i) = (x, y)$ and $\{p, \chi_t\} = \{\lambda_i(x) \wedge p, \lambda_i(y) \wedge p\}$. From this construction we obtain:

- $\chi_t(s) = p(s)$ for all $s \in T - \{t\}$.
- $\chi_t(t) < p(t)$ for all $t \in T$.
- $\chi_t(t) = 0$ and $0 < p(t)$ for all $t \in T$.

The last item listed above is a consequence of the flatness of $\mathbf{A}$. Thus, χ_t agrees with p at all coordinates with the exception of t, where χ_t is 0 while p is not 0. So χ_t is uniquely determined by p and t (and is independent of the choices of x, y, and λ_i made above). We will eventually see—once enough is specified about $\mathbf{A}$—that p is also uniquely determined.

Fix $t_0 \in T$ so that $u \leq \chi_{t_0}$ for some $u < p$ for which (u, p) is critical over θ. Let $q = \chi_{t_0}$.

Fact 1. p is a maximal element of $\mathbf{A}^T$. $\chi_t \in B$ and p covers χ_t in $\mathbf{A}^T$ for all $t \in T$. (q, p) is critical over θ. Finally, if $u \in A^T$ and $u < p$, then $u \in B$.

Proof. Essentially, Fact 1 gathers the conclusions we drew above. To see that (q, p) is critical, notice $(q, p) \notin \theta$ according to Fact 0. Let $u \leq q < p$ with (u, p) critical over θ. Then we have $p \varphi u = q \wedge u \varphi q \wedge p = q$, for all $\varphi > \theta$. The elements of A^T less than or equal to p form a Boolean algebra in which every element is a meet of the coatoms χ_t. □

Fact 2. If $p \; \theta \; x$, then $p = x$

Proof. Suppose $p\ \theta\ x$. Meeting both sides with p we also get $p\ \theta\ p \wedge x$. From Fact 0, we conclude that $p \not> p \wedge x$. Thus $p \leq x$. But since p is a maximal element, we arrive at $p = x$. □

Fact 3. *$x\ \theta\ y$ if and only if $\mu(x) = p \Leftrightarrow \mu(y) = p$ for all translations μ of $\mathbf{B}$.*

Proof. In the forward direction the result follows from Fact 2.

Now for the converse direction, suppose $(x, y) \notin \theta$. By Fact 1, we know $(q, p) \in \theta \vee \mathrm{Cg}^{\mathbf{B}}(x, y)$. Now repeating the analysis that led to the χ_t's we obtain a translation $\mu = \lambda \wedge p$ so that $\mu(x) \neq \mu(y)$ but $p \in \{\mu(x), \mu(y)\}$. □

Fact 4. If $x < p$, then $(x, x \wedge q) \in \theta$.

Proof. (q, p) is critical over θ by Fact 1, so $x = x \wedge p \varphi x \wedge q$, for all $\varphi > \theta$. Hence, either $(x, x \wedge q) \in \theta$ or $(x, x \wedge q)$ is critical over θ. Since $p > x \geq x \wedge q$, it follows from the minimality of p that $x\ \theta\ x \wedge q$. □

Suppose that x, y, and $z \in B$. Then $(x \wedge y)$ and $(x \wedge z)$ also belong to B and the element x is a common upper bound. Recalling that $\mathbf{B}$ has the structure of a finite $\wedge$-semilattice, it follows that $(x \wedge y)$ and $(x \wedge z)$ must have a least upper bound—we denote it by $(x \wedge y) \vee (x \wedge z)$.

Fact 5. *$\mathbf{S} \in \mathcal{HS}\mathbf{A}$ or $(x \wedge y) \vee (x \wedge z)$ is not a polynomial of $\mathbf{B}$.*

Proof. Suppose $\mathbf{S} \notin \mathcal{HS}\mathbf{A}$. Then T has at least two elements. Let $t_1 \in T$ with $t_0 \neq t_1$. Let $q' = \chi_{t_1}$. Since $q' < p$ we have by Fact 4 that $q'\ \theta\ q' \wedge q$. But then, were $(x \wedge y) \vee (x \wedge z)$ a polynomial of $\mathbf{B}$, we would have $p = (p \wedge q) \vee (p \wedge q')\ \theta\ (p \wedge q) \vee (p \wedge q \wedge q') = q$. Since $(p, q) \notin \theta$, we conclude that $(x \wedge y) \vee (x \wedge z)$ is not a polynomial. □

Fact 5 reveals that our investigation of (finite) subdirectly irreducible algebras can be split in two. Since $\mathbf{A}$ is finite, a complete description of the subdirectly irreducible algebras in $HS\mathbf{A}$ can be devised given a description of $\mathbf{A}$. We only note the obvious upper bound on their cardinality. Most of our effort will concern the alternative case when $(x \wedge y) \vee (x \wedge z)$ is not a polynomial of $\mathbf{B}$. It is th subdirectly irreducible algebras arising from these algebras that we want to show must be isomorphic to our $\mathbf{Q}_n$'s.

Here is a lemma that simply gathers together the most salient of the facts just listed.

Lemma 9.2.1. *Suppose that $\mathbf{A}$ is a finite flat algebra and that $\mathbf{S}$ is a finite subdirectly irreducible algebra in $\mathcal{HSP}\mathbf{A}$. Choose T, $\mathbf{B}$, and $\theta \in \mathrm{Con}\,\mathbf{B}$ so that*

- *$\mathbf{B}$ is a subalgebra of $\mathbf{A}^T$,*
- *θ is (strictly) meet irreducible in $\mathrm{Con}\,\mathbf{B}$.*
- *$\mathbf{S} \cong \mathbf{B}/\theta$, and*
- *T is as small as possible subject to fulfilling the conditions above.*

Then there is an element $p \in B$ such that

i. *$(0, p)$ is critical over θ,*

ii. *$p/\theta = \{p\}$,*

iii. *p is a maximal element of A^T (so $p(s) > 0$ for all $s \in T$), and*

iv. *for all $x, y \in B$, $x\ \theta\ y$ if and only if $\mu(x) = p \Leftrightarrow \mu(y) = p$ for all translations μ of $\mathbf{B}$.*

Moreover, $\mathbf{S} \in \mathcal{HS}\mathbf{A}$ or $(x \wedge y) \vee (x \wedge z)$ is not a polynomial of $\mathbf{B}$.

LESSON 10

THE EIGHT ELEMENT ALGEBRA A

The six element algebra **R** which was constructed at the end of Lecture **??** generates a variety with a lot of finite subdirectly irreducible algebras in addition to the $\mathbf{Q}_n$'s. An example is the flat automatic algebra $\mathbf{S}_8$ described next. For each $i < 8$, define the following elements of the 8-fold direct power of the algebra **R**:

$$\begin{aligned} c &= \langle q,q,\ldots,q,q,q,\ldots\rangle \\ d_i &= \langle q,q,\ldots,q,\ r,q,\ldots\rangle \\ r_i &= \langle 3,3,\ldots,3,\ 2,3,\ldots\rangle \end{aligned}$$

where the sole r in d_i and the sole 1 in r_i occur at position i. Let B_8 be the subset of the 8-fold direct power consisting of c, all the d_i's, all the r_i's, and all the 8-tuples which have 0 in at least one position. Then

$$r_i \cdot d_i = c$$

for every $i < 8$, but any other product of elements of B_8 results in an 8-tuple with 0 in at least one position. This means that B_8 is a subuniverse of the eightfold direct power of **R**. Let θ_8 be the equivalence relation on $\mathbf{B}_8$ which collapses into one big block all the 8-tuples in B_8 which have 0 in at least one position, but which isolates all the other members into singletons. It is easy to see that θ_8 is a congruence relation of $\mathbf{B}_8$. Let $\mathbf{S}_8 = \mathbf{B}_8/\theta_8$. $\mathbf{S}_8$ is displayed below in Figure 10.1.

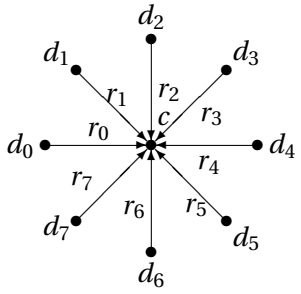

Figure 10.1: The directed graph for $\mathbf{S}_8$

Plainly, $\mathbf{S}_8$ is an algebra in the variety generated by the flat automatic algebra **R**. A routine calculation shows that $(c,0)$ is a critical pair in $\mathbf{S}_8$. Consequently, $\mathbf{S}_8$ is subdirectly irreducible, as desired. Evidently,

there is nothing special in the choice of 8. A similar construction can be carried out for any cardinal in place of 8. A little reflection reveals that there must be, in the variety generated by **R**, other subdirectly irreducible algebras, each having the appearance of a tree with branches directed toward the root.

We must modify our little 6-element algebra from Lecture **??** to eliminate subdirectly irreducible algebras like $\mathbf{S}_8$, whose diagrams are not (finite) directed paths. Evidently, for our subdirectly irreducible algebras, we need a kind of unique factorization property:

$$a \cdot b = c \cdot d \neq 0 \Rightarrow a = c \text{ and } b = d.$$

To accomplish this we are going to add some new basic operations and some new elements to the algebra **R**, but we need to have some care since we want $\mathbf{Q}_{\mathbb{Z}}$ to remain essentially unchanged and still to belong to the variety generated by the finite algebra we are trying to devise.

To obtain the unique factorization property we introduce the new basic 4-place operation U^0:

$$U^0(x,y,z,w) = \begin{cases} xy & \text{if } xy = zw \neq 0 \text{ and } x = z \text{ and } y = w, \\ x\bar{y} & \text{if } xy = zw \neq 0 \text{ and either } x \neq z \text{ or } y \neq w, \\ 0 & \text{otherwise.} \end{cases}$$

At the moment, we should understand that the first case corresponds to the situation when the unique factorization property prevails, the second case corresponds to the failure of the unique factorization property, and the remaining case is just a default. For the moment, $\overline{xy}$ is simply a reminder that the output in this case should depend on xy but differ from xy. Our hope is to obtain the unique factorization property by forcing the first case to happen. In essence, this means preventing the second case. For this purpose we introduce a new basic 5-place operation S_2:

$$S_2(u,v,x,y,z) = \begin{cases} (x \wedge y) \vee (x \wedge z) & \text{if } u = \bar{v}, \\ 0 & \text{otherwise.} \end{cases}$$

Recall the algebra **B** from the preceding lecture. In **B** we know from Fact 5 that $(x \wedge y) \vee (x \wedge z)$ cannot be a polynomial. So S_2 is designed to prevent B from having elements u and v so that $u = \bar{v}$. This in turn will prevent the second case in the definition of U^0 from arising.

To give more sense to this, notice that in six element algebra **R**, a product xy could have only q, r, or 0 as a value. So we introduce two elements $\bar{q}$ and $\bar{r}$ in addition to the six with which we have been dealing. Further, we stipulate that $\bar{u} = q$ if $u = \bar{q}$ and likewise $\bar{u} = r$ if $u = \bar{r}$. In this way, both U^0 and S_2 have unambiguous definitions, once the product and meet have been extended to operations on the new set with eight elements.

These two additional operations and two additional elements are not quite enough.

$$S_1(u,v,x,y,z) = \begin{cases} (x \wedge y) \vee (x \wedge z) & \text{if } u \in \{1,3\}, \\ 0 & \text{otherwise.} \end{cases}$$

The role of S_1, as we will see, is ensure that our finite subdirectly irreducible algebra **S** has another property that each $\mathbf{Q}_n$ has—namely, that the labels of the edges are not repeated. Last, here are the operations J and J' which are 3=-place operations:

$$J(x,y,z) = \begin{cases} x & \text{if } x = y \neq 0, \\ x \wedge z & \text{if } x = \bar{y}, \\ 0 & \text{otherwise.} \end{cases} \qquad J'(x,y,z) = \begin{cases} x \wedge z & \text{if } x = y \neq 0, \\ x & \text{if } x = \bar{y}, \\ 0 & \text{otherwise.} \end{cases}$$

The role of these operations is less forthright. Since we are really working inside a subalgebra of a direct power, we have to contend with coordinate-wise properties. The role of these last two operations is to ensure that we fall into the "good" case at every coordinate.

We are led to a flat algebra **A** with eight elements and eight basic operations.

The universe is $A = \{0\} \cup \{1,2,3\} \cup \{q, \bar{q}, r, \bar{r}\}$. We set $U = \{1,2,3\}$ and $W = \{q, \bar{q}, r, \bar{r}\}$. We regard $\bar{}$ as an involution on W. The basic operations of **A** are denoted by $0, \wedge, \cdot, J, J', U^0, S_1$, and S_2. $\langle A, \wedge, 0\rangle$ is a flat semilattice with least element 0. The operation $\cdot$ is defined to give the default value 0 except when

$$\begin{aligned} 1 \cdot r &= r & 1 \cdot \bar{r} &= \bar{r} \\ 2 \cdot r &= q & 2 \cdot \bar{r} &= \bar{q} \\ 3 \cdot q &= q & 3 \cdot \bar{q} &= \bar{q} \end{aligned}$$

This is an automatic operation. Ordinarily, we represent the product $\cdot$ simply by juxtiposition. Here is the diagram of the automatic algebra **A**:

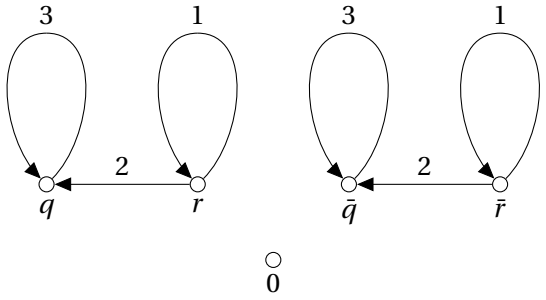

The following fact is evident from the definition of the product.

Fact 6. If λ is a basic translation on **A** associated with the product $\cdot$, and $\lambda(a) = \lambda(b) \neq 0$, then $a = b$. The same is true for every translation built using only the product.

LESSON 11

PROPERTIES OF B BASED ON THE EIGHT ELEMENT ALGEBRA A

With the description of our eight element algebra **A** in hand, we continue to develop facts about **B** and its congruence θ. Denote by B_1 the set consisting of p and all its factors with respect to the product $\cdot$. That is

$$B_1 = \{u : \lambda(u) = p \text{ for some nonconstant translation } \lambda \text{ of } \mathbf{B} \text{ built only from the product}\}$$

So $u \in B_1$ if and only if $u = p$ or $u = c_i$ for some factorization $p = c_0 c_1 \dots c_m$ (where this latter product is associated to the right).

Let B_0 denote that set of those tuples in B which contain at least one 0. Plainly $B_0 \subseteq B - B_1$. It is also clear that if $\mathbf{S} \notin HS\mathbf{A}$, then the ranges of the operations S_1 and S_2 are contained in B_0 and hence in $B - B_1$.

The basic operation J of **A** is **monotone** in the sense that if $a \le a', b \le b'$ and $c \le c'$ where all these elements belong to A, then $J(a,b,c) \le J(a',b',c')$.

Fact 7. Let f be a monotone unary polynomial of **B**. If $x < p$ and $f(x) = p$, then $f(q) = p$.

Proof. By Fact 4 we have $x \,\theta\, x \wedge q$. This entails $p = f(x) \,\theta\, f(x \wedge q)$. So by Fact 2 we get $p = f(x \wedge q)$. But then $p \le f(q)$ by the monotonicity of f. Thus $p = f(q)$ by the maximality of p. □

PROVISO: *The facts below are established under the assumption that the ranges of S_1 and S_2 are contained in $B - B_1$.*

Fact 8. If $u \in B_1$ and $v \in B$ so that for all $s \in T$ either $u(s) = v(s)$ or $u(s) = \overline{v(s)} \in W$, then $u = v$.

Proof. First suppose $u = p$.

Let $Y = \{s : p(s) = \overline{v(s)}\}$.

CLAIM: Y is empty.

Proof of the Claim: Since the range of the operation S_2 is disjoint from B_1, it follows that $T \neq Y$. Pick $t' \in T - Y$ and let $q' = \chi_{t'}$. So for each $s \in T$ we have

$$J(p(s), v(s), q'(s)) = \begin{cases} p(s) & \text{if } s \notin Y, \\ p(s) \wedge q'(s) & \text{if } s \in Y. \end{cases}$$

But this entails $J(p, v, q') = p$, since $q'(s) = p(s)$ for all $s \in Y$ because $t' \notin Y$. Therefore, by Fact 7 and the monotonicity of J, we have $J(p, v, q) = p$. But then the definition of J gives us $q(s) = p(s)$ for all $s \in Y$. Since $q(t_0) = 0$, it follows that $t_0 \notin Y$.

Now observe that $J'(p(t_0), v(t_0), q(t_0)) = p(t_0) \wedge 0 = 0$. Hence, $J'(p, v, q) \neq p$. So by Fact 7 and the monotonicity of J', we conclude that $J'(p, v, \chi_t) \neq p$ for all $t \in T$. But for all $s, t \in T$

$$J'(p(s), v(s), \chi_t(s)) = \begin{cases} p(s) \wedge \chi_t(s) & \text{if } s \notin Y, \\ p(s) & \text{if } s \in Y. \end{cases}$$

It follows that $t \notin Y$ for all $t \in Y$. This means Y is empty. So the Claim is established.

Since Y is empty, we also know that $p(s) = v(s)$ for all $s \in T$. Hence $u = v$ as desired.

Now suppose $u \in B_1 - \{p\}$. There are two kinds of elements in $B_1 - \{p\}$—those in U^T and those in W^T. Clearly, we can restrict our attention to the case when $u \in W^T$. Let λ be a translation built from the product such that $\lambda(u) = p$. Set $p' = \lambda(v)$. Since the product respects bars on elements, we see that for each $s \in T$, either $p(s) = p'(s)$ or $p(s) = \overline{p'(s)}$. So by the claim just established, we have $\lambda(u) = p = p' = \lambda(v)$. But then $u = v$ by Fact 6 □

Our basic strategy calls for θ to isolate the members of B_1 and to lump all the elements of $B - B_1$ together. To see that this really does happen, in view of Fact 3 we need the following.

Fact 9. If $u \in B$ and $\lambda(u) \in B_1$ for some nonconstant translation λ, then $u \in B_1$.

Proof. The proof is by induction on the complexity of λ. The initial step of the induction is obvious, since the identity function is the only simplest nonconstant translation. The inductive step breaks down into seven cases, one for each basic operation of positive rank.

CASE $\wedge$: $\lambda(x) = \mu(x) \wedge r$, where $r \in B$.

We have $\lambda(u) \leq \mu(u)$. But every element of B_1 is maximal with respect to the semilattice order. So $\lambda(u) = \mu(u) \in B_1$. Now μ must be nonconstant. Invoking the induction hypothesis, we get $u \in B_1$.

CASE $\cdot$: $\lambda(x) = \mu(x)r$ or $\lambda(x) = r\mu(x)$.

Under the first alternative we have $\mu(u)r = \lambda(u) \in B_1$. So $\mu(u), r \in B_1$. Since μ must be nonconstant, we can invoke the induction hypothesis to conclude that $u \in B_1$. The other alternative is similar.

CASE J: $\lambda(x) = J(\mu(x), r, s)$ or $\lambda(x) = J(r, \mu(x), s)$ or $\lambda(x) = J(r, s, \mu(x))$.

Consider the first alternative. We have $\lambda(u) = J(\mu(u), r, s) \leq \mu(u)$. By the maximality of $\lambda(u)$ we get

$$\lambda(u) = J(\mu(u), r, s) = \mu(u) \in B_1.$$

Now μ cannot be constant. Hence we can invoke the inductive hypothesis to conclude that $u \in B_1$. The second alternative is similar, except that Fact 8 comes into play. Under the last alternative, since $r \geq J(r, s, \mu(u)) = \lambda(u)$ is maximal, we see that r and s fulfill the hypotheses of Fact 8. Consequently, $r = s \in B_1$. But then, $\lambda(x) = J(r, s, \mu(x)) = r$ according to the definition of J. This means the third alternative is impossible, since $\lambda(x)$ is not constant.

CASE J': $\lambda(x) = J'(\mu(x), \nu(x), \rho(x))$.

This case is easier than the last one and is omitted.

CASES S_1 AND S_2: Too easy.

CASE U^0: $\lambda(x) = U^0(\mu(x), s, r', s')$ or $\lambda(x) = U^0(r, \mu(x), r', s')$ or $\lambda(x) = U^0(r, s, \mu(x), s')$ or $\lambda(x) = U^0(r, s, r', \mu(x))$.

Consider the first alternative. We have $\lambda(u) = U^0(\mu(u), s, r', s') \in B_1$. Evidently, $\lambda(u)$ and $\mu(u)s$ satisfy the hypotheses of Fact 8. So $\lambda(u) = \mu(u)s$. Since $\lambda(u) \in B_1$, we know that $\mu(u) \in B_1$ by the definition of B_1. Now μ is nonconstant. So $u \in B_1$ by the inductive hypothesis. The second alternative is similar.

Consider the third alternative. We have $\lambda(u) = U^0(r,s,\mu(u),s')$. Evidently, $\lambda(u)$ and rs satisfy the hypotheses of Fact 8. So $\lambda(u) = rs$. Then by the definition of T, we have $\lambda(u) = rs = \mu(u)s'$. But then $\mu(u) \in B_1$ and the induction hypotheses applies to yield $u \in B_1$. The fourth alternative is similar. □

Fact 10. $u/\theta = \{u\}$ for each $u \in B_1$ and $0/\theta = B - B_1$.

Proof. Suppose $u \in B_1$ and that $u\,\theta\,v$. Let $\lambda(u) = p$ for some translation λ built just using $\cdot$. It follows that $\lambda(v) = p$ by Fact 3. By Fact 6, we conclude that $u = v$.

Fact 9 says that $B - B_1$ is closed with respect to nonconstant translations. Since $p \in B_1$, we have that $\lambda(u) \neq p$ for all $u \in B - B_1$ and all nonconstant translations λ. Hence, by Fact 3, $B - B_1$ is collapsed by θ. But, as we just saw, B_1 is the union of (singleton) θ-classes. Hence $B - B_1$ is a θ-class. Clearly, $0 \in B - B_1$. □

To establish that $\mathbf{S} \cong \mathbf{Q}_n$ for some natural number n we need to analyze each of our basic operations. We deal with the product first.

Here is the unique factorization property for the product that we require.

Fact 11. If $ab = cd \in B_1$, then $a = c$ and $b = d$.

Proof. Let $u = ab$ and $v = U^0(a,b,c,d)$. From the definition of the operation U^0, we see that u and v satisfy the hypotheses of Fact 8. Hence, $ab = U^0(a,b,c,d)$. But then the definition of U^0 gives $a = c$ and $b = d$. □

In $\mathbf{Q}_\mathbb{Z}$ none of the labels of the edges were repeated. We need this property as well. It is the reason why we introduced the operation S_1. The relevant fact is next.

Fact 12. No factorization of p has repeated factors.

Proof. It is clear that if $d_0 d_1 \ldots d_{m-1} e = p$ then $e \in W^T$ while $d_0, \ldots, d_{m-1} \in U^T$. Suppose that $d_i = d_j$ with $i < j$. Since the range of the operation S_1 is disjoint from B_1, we conclude that B contains no elements from $\{1,2\}^T$. So pick $s \in T$ so that $d_i(s) = d_j(s) = H$. Now we see

$$p(s) = d_0(s) \ldots d_{i-1}(s) H d_{i+1}(s) \ldots d_{j-1}(s) H d_{j+1}(s) \ldots d_{m-1}(s) e(s)$$

So $p(s) = 0$, violating the maximality of p. □

We are now in a position to describe B_1 more explicitly. Consider the following factorization of p:

$$\begin{aligned} p &= b_0 \\ &= a_0 b_1 \\ &= a_0 a_1 b_2 \\ &\vdots \\ &= a_0 a_1 \ldots a_{n-1} b_n \\ &\vdots \end{aligned}$$

Evidently, $a_i \in U^T$ for all i and according to Fact 12 all the a_i's are distinct. But B_1 is finite, so we suppose without loss of generality that b_n cannot be factored. But the unique factorization property Fact 11 entails that the factorization of p displayed above is the only way p can be factored. Consequently,

$$B_1 = \{a_0, a_1, \ldots, a_{n-1}\} \cup \{b_0, b_1, \ldots, b_n\}$$

It is also evident that $b_i \in W^T$ for all i. Were $b_i = b_j$ for some $i \neq j$, it would be easy to construct a factorization of p with repeated factors, in violation of Fact 12. This means that B_1 has $2n+1$ elements, and that $b_i = a_i b_{i+1}$ for each $i < n$. That all the other products of elements chosen from B_1 will belong to B_0, follows easily from the unique factorization property Fact 11. Consequently, at least with respect to the product operation, **S** and $\mathbf{Q}_n$ are isomorphic.

Now consider the operation $\wedge$. Since $\wedge$ is obviously a semilattice operation on **S**, what we need is that **S** is flat.

Fact 13. If $x, y \in B_1$ and $x \neq y$, then $x \wedge y \in B - B_1$.

Proof. Since $x \neq y$ there is $t \in T$ with $x(t) \neq y(t)$. But then $((x \wedge y)(t) = 0$. So $x \wedge y \in B - B_1$. □

Finally, we need to know that the remaining basic operations on **S** can be construed as term operations built up from $\cdot, \wedge$, and 0 in a manner dependent only on the hypotheses that **S** is a finite subdirectly irreducible algebra in $HSP\mathbf{A}$ and that $\mathbf{S} \notin HS\mathbf{A}$. That is the content of the next sequence of facts.

Fact 14. $U^0(x, y, z, w)\ \theta\ (xy) \wedge (zw)$ for all $x, y, z, w \in B$.

Proof. We must show that either $U^0(x, y, z, w)$ and $(xy) \wedge (zw)$ both belong to $B - B_1$ or else $U^0(x, y, z, w) = (xy) \wedge (zw) \in B_1$. Since $B - B_1$ is a θ-class, Fact 8 forces $U^0(x, y, z, w) \in B - B_1$ except in the case that $xy = zw \in B_1$. In that case, $U^0(x, y, z, w) = xy = zw = (xy) \wedge (zw) \in B_1$. But also, $(xy) \wedge (zw) \in B - B_1$ except in the case that $xy = zw \in B_1$. In that case, $U^0(x, y, z, w) = xy = (xy) \wedge (zw) \in B_1$. Therefore, $U^0(x, y, z, w)\ \theta\ (xy) \wedge (zw)$. □

Fact 15. $J(x, y, z)\ \theta\ x \wedge y$ for all $x, y, z \in B$.

Proof. Again, we must show that either $J(x, y, z)$ and $x \wedge y$ both belong to $B - B_1$ or else $J(x, y, z) = x \wedge y \in B_1$. Now again using that $B - B_1$ is a θ-class and Fact 8, $J(x, y, z) \in B - B_1$, except in the case that $x = y \in B_1$. In that case, $J(x, y, z) = x = y = x \wedge y \in B_1$. But also, $x \wedge y \in B - B_1$, except in the case that $x = y \in B_1$. In that case, $x \wedge y = x = J(x, y, z) \in B_1$. Therefore, $J(x, y, z)\ \theta\ x \wedge y$. □

Fact 16. $J'(x, y, z)\ \theta\ x \wedge y \wedge z$ for all $x, y, z \in B$.

Proof. This is too easy. □

Fact 17. $S_1(u, v, x, y, z)\ \theta\ 0\ \theta\ S_2(u, v, x, y, z)$ for all $u, v, x, y, z \in B$.

For each natural number n, we take $\mathbf{Q}_n$ to be an algebra on $2n+2$ elements with the basic operations $\cdot, \wedge$, and 0 as described in Lecture 0, and the remaining basic operations determined by the stipulation that the following equations are true in $\mathbf{Q}_n$:

$$U^0(x, y, z, w) \approx (xy) \wedge (zw)$$

$$\begin{aligned} J(x, y, z) &\approx x \wedge y & \qquad S_1(u, v, x, y, z) &\approx 0 \\ J'(x, y, z) &\approx x \wedge y \wedge z & \qquad S_2(u, v, x, y, z) &\approx 0 \end{aligned}$$

Thus we arrive at the desired conclusion.

Lemma 11.0.1. *Let* **S** *be a finite subdirectly irreducible algebra in* $HSP\mathbf{A}$. *Either* $\mathbf{S} \in HS\mathbf{A}$ *or else there is a natural number n such that* $\mathbf{S} \cong \mathbf{Q}_n$.

What we haven't done in this lecture is prove that any of these expanded $\mathbf{Q}_n$'s belong to the variety generated by our 8-element algebra **A**.

LESSON 12

A IS INHERENTLY NONFINITELY BASED AND HAS RESIDUAL CHARACTER ω_1

The algebra $\mathbf{Q}_\mathbb{Z}$ and its subalgebras $\mathbf{Q}_\omega$, and $\mathbf{Q}_n$ for each $n \in \omega$, were introduced in Lecture **??**. The operations $0, \wedge$, and $\cdot$ were examined in detail, but the only stipulation about any remaining operations was that they must be defined as term operations of these first three. In Lecture **??**, five more operation symbols were introduced: U^0, J, J', S_1, and S_2. In the algebras $\mathbf{Q}_\mathbb{Z}, \mathbf{Q}_\omega$, and $\mathbf{Q}_n$ these five further basic operations are defined so that the following equations are true:

$$U^0(x,y,z,w) \approx (xy) \wedge (zw)$$

$$\begin{aligned} J(x,y,z) &\approx x \wedge y & \qquad S_1(u,v,x,y,z) &\approx 0 \\ J'(x,y,z) &\approx x \wedge y \wedge z & \qquad S_2(u,v,x,y,z) &\approx 0 \end{aligned}$$

The whole discussion of these algebras in Lecture **??** goes through in this expanded setting, with the exception of the last phase. The five new operations were not defined on the six element algebra **R** in Lecture 0. We now want to replace that algebra with the eight element algebra **A** introduced in Lecture **??**. What we need is the following theorem to replace Theorem **??** of Lecture **??**.

Theorem 12.0.1. $\mathbf{Q}_\mathbb{Z}$ *belongs to the variety generated by* **A**.

Proof. We retrace the proof of Theorem **??**. First, for each $p \in \mathbb{Z}$ we designate elements α_p and β_p of $A^\mathbb{Z}$ as before:

$$\begin{array}{lccccccccc} \alpha_p := & \ldots & 1 & 1 & 1 & 2 & 3 & 3 & 3 & \ldots \\ \beta_p := & \ldots & r & r & r & q & q & q & q & \ldots \end{array}$$

where the change is taking place at the p^{th} position. Next we let $B_1 = \{\alpha_p : p \in \mathbb{Z}\} \cup \{\beta_p : p \in \mathbb{Z}\}$ and we let $\mathbf{B}$ be the subalgebra of $\mathbf{A}^\mathbb{Z}$ generated by B_1. B_0 be the set of all elements of B in which 0 occurs. Now let Φ be the map defined from B to $A_\mathbb{Z}$ via

$$\Phi(x) = \begin{cases} a_p & \text{if } x = \alpha_p \text{ for some } p \in \mathbb{Z}, \\ b_p & \text{if } x = \beta_p \text{ for some } p \in \mathbb{Z}, \\ 0 & \text{otherwise} \end{cases}$$

We contend that $B_0 \cup B_1$ is a subuniverse of $\mathbf{B}$ (and so $B = B_0 \cup B_1$) and also that Φ is a homomorphism from $\mathbf{B}$ onto $\mathbf{Q}_{\mathbb{Z}}$. Checking either of these contentions can be done by examining the behavior of each baisc operation case by case. We will do this simultaneously.

CASE 0. Plainly, $0 \in B_0$ and $\Phi(0) = 0$. So this case is secure.

CASE $\wedge$. Suppose that $u, v \in B_0 \cup B_1$. Then either $u = v$ and $u \wedge v = u \in B_0 \cup B_1$ or else $u \neq v$ and $u \wedge v \in B_0$. Hence, $B_0 \cup B_1$ is closed under $\wedge$. But also, $\Phi(u \wedge v) = \Phi(u) \wedge \Phi(v)$.

CASE $\cdot$. Suppose that $u, v \in B_0 \cup B_1$. Then either $uv \in B_0$ or for some p we have $u = \alpha_p, v = \beta_{p+1}$ and $uv = \beta_p \in B_1$. It follows that $B_0 \cup B_1$ is closed under $\cdot$ and that Φ preserves $\cdot$.

To handle the remaining cases, the following property of $x, y \in B_0 \cup B_1$ proves useful:

$$\text{If for each } s \in \mathbb{Z}, \text{ either } x(s) = y(s) \neq 0 \text{ or } x(s) = \overline{y(s)}, \text{ then } x = y. \qquad (\star)$$

Since any $x, y \in B_0 \cup B_1$ for which the hypothesis of $\star$ holds must both belong to B_1, and since bars occur in no member of B_1, $\star$ is true.

CASE J. For $J(x, y, z) \notin B_0$, observe that the inputs x and y must satisfy the hypothesis of $\star$. Hence, either $J(x, y, z) \in B_0$ or $x = y \in B_1$ and $J(x, y, z) = x \in B_1$. So $B_0 \cup B_1$ is closed under J and Φ preserves J.

CASE J'. This case is very similar to the last case.

CASE U^0. Let $x, y, z, w \in B_0 \cup B_1$. Let $u = U^0(x, y, z, w)$ and $v = xy$. Then $U^0(x, y, z, w) \in B_0$ unless u and v fulfill the hypothesis of $\star$. In that case, we must have $U^0(x, y, z, w) = u = v \in B_1$. Consequently, $B_0 \cup B_1$ is closed under U^0 and Φ preserves U^0.

CASE S_1. Let $u, v, x, y, z \in B_0 \cup B_1$. Then $S_1(u, v, x, y, z) \in B_0$ unless $u \in \{1, 2\}^{\mathbb{Z}}$. But $\{1, 2\}^{\mathbb{Z}}$ and $B_0 \cup B_1$ are disjoint. Consequently, $B_0 \cup B_1$ is closed under S_1 and Φ preserves S_1.

CASE S_2. Let $u, v, x, y, z \in B_0 \cup B_1$. Then $S_2(u, v, x, y, z) \in B_0$ unless $u(s) = \overline{v(s)}$ for all $s \in \mathbb{Z}$. But no element of B_1 has a bar at any of its entries. Consequently, $B_0 \cup B_1$ is closed under S_2 and Φ preserves S_2.

□

At this point we know that the eight element algebra $\mathbf{A}$, which has eight basic operations, is inherently nonfinitely based, that the finite subdirectly irreducible algebras in the variety generated by $\mathbf{A}$ are the subdirectly irreducible algebras in $HS\mathbf{A}$ and the algebras $\mathbf{Q}_n$ for each $n \in \omega$, and that $\mathbf{Q}_\omega$ is a countably infinite subdirectly irreducible member of the variety.

We will demonstrate that our variety has no other infinite subdirectly irreducible algebras.

Let $\mathbf{S}$ be any infinite subdirectly irreducible algebra in the variety generated by $\mathbf{A}$. According to the Theorem of Dziodiak and Quackenbush (see the Toolbox), any finite subalgebra of $\mathbf{S}$ can be embedded into arbitrarily large finite subdirectly irreducible algebras in the variety generated by $\mathbf{A}$, i.e. into $\mathbf{Q}_n$ for all large enough n. This means that every finitely generated (= finite) subalgebra of $\mathbf{S}$ is embeddable into $\mathbf{Q}_\omega$. Consequently, every universal sentence true in $\mathbf{Q}_\omega$ must be true in $\mathbf{S}$.

Here are some interesting properties of $\mathbf{Q}_\omega$ which can be expressed with universal sentences:

- Any equation true in $\mathbf{Q}_{\mathbb{Z}}$. For example: $\quad U^0(x, y, z, w) \approx (xy) \wedge (zw)$.
- The height is no bigger than 1: $x \not\approx y \to x \wedge y \approx 0$.
- $xy \approx zw \not\approx 0 \to (x \approx z \,\&\, y \approx w)$.
- $xy \not\approx 0 \not\approx xz \to y \approx z$.
- $xy \not\approx 0 \not\approx zy \to x \approx z$.

- $xy \not\approx 0 \rightarrow zx \approx 0 \approx yw$.

Consequently, in **S**, the operations U^0, J, J', S_1, and S_2 are term functions (using the same terms as in $\mathbf{Q}_\omega$) in $0, \wedge$, and $\cdot$. We ignore them from now on. With respect to $\wedge$ and 0, **S** is a height 1 meet-semilattice with least element 0. So the balance of our analysis depends primarily on the product $\cdot$. Since $(xy)z \approx 0$ is true in $\mathbf{Q}_\omega$, we see that in **S**, just as in $\mathbf{Q}_\omega$, only right-associated products can differ from 0. The last four properties itemized above put further and severe restrictions on the product in **S**.

We make $S - \{0\}$ into a labelled directed graph as follows. We take as the vertex set those elements which are right factors, outputs or do not occur in nonzero products. We take as the set of labels those elements which are left factors in nonzero products. Our itemized properties entail that the set of vertices and the set of labels are disjoint. We put an edge from b to c and label it with a provided $ab = c$ in **S**. Our itemized assertions ensure that a vertex can have outdegree at most 1, indegree at most 1, and that every edge has a uniquely determined label which occurs as a label of exactly one edge in the whole graph.

Let C be a connected component of our graph. Let θ_C be the equivalence relation that collapses all the vertices and labels in C to 0, but which isolates every other point. θ_C is a congruence of **S**. Since **S** is subdirectly irreducible, it follows that our graph has only one component. This already implies that **S** is countably infinite. But more is true. There are only three possible countable connected graphs of this kind: the one associated with $\mathbb{Z}$ (and then we would have $\mathbf{S} \cong \mathbf{Q}_\mathbb{Z}$), the one associated with ω (and then we would have $\mathbf{S} \cong \mathbf{Q}_\omega$), and the one associated with the set of nonnegative integers (and then **S** would be isomorphic to an algebra we might as well call $\mathbf{Q}_{-\omega}$). But neither $\mathbf{Q}_\mathbb{Z}$ nor $\mathbf{Q}_{-\omega}$ is subdirectly irreducible. So **S** must be isomorphic to $\mathbf{Q}_\omega$.

We summarize the results in the following theorem.

Theorem 12.0.2. *The eight element algebra* **A***, which has only eight basic operations, is inherently nonfinitely based. The subdirectly irreducible algebras in the variety generated by* **A** *are, up to isomorphism, exactly the subdirectly irreducible homomorphic images of subalgebras of* **A***, the algebra* $\mathbf{Q}_\omega$*, and the algebra* $\mathbf{Q}_n$ *for each* $n \in \omega$.

This theorem settles in the negative some outstanding problems. We will say that a variety is *finitely generated* provided it is generated by a finite algebra with only finitely many fundamental operations. It is *residually small* if there is an upper bound on the cardinalities of its subdirectly irreducible algebras. It is *residually finite* if all its subdirectly irreducible algebras are finite. It is *residually very finite* if there is a finite upper bound on the cardinalities of its subdirectly irreducible algebras.

The R-S Conjecture: *Every finitely generated residually small variety is residually very finite.*

The Broader Finite Basis Speculation: *Every finitely generated residually small variety is finitely based.*

Theorem 12.0.2 is a counterexample to both of these. However, the two problems below are closely related and still open.

The Quackenbush Conjecture: *Every finitely generated residually finite variety is residually very finite.*

Park's Conjecture: *Every finitely generated residually finite variety is finitely based.*

LESSON 13

HOW A($\mathcal{T}$) ENCODES THE COMPUTATIONS OF $\mathcal{T}$

In this lecture we describe, in part, McKenzie's machine algebras and show how they capture the computations of Turing machines. Turing machines are finite objects, but the computations that they produce can be endless. So it is reasonable to expect to use a finite algebra to convey the information of any particular Turing machine. However, finite algebras are too small to hold arbitrary computations. The algebra $\mathbf{Q}_{\mathbb{Z}}$, however, suggests a way to grapple with arbitrary computations. The idea is to designate certain elements of the algebra as configurations of a Turing machine and draw labeled directed edges between configurations to represent the transitions of the machine computation. Then we try to realize these directed edges by new operations applied to certain elements. Next we try to find a finite algebra so that the whole thing is happening coordinatewise inside a big direct power. Finally, we will have to add further operations to control all the finite subdirectly irreducible algebras.

For a Turing machine $\mathcal{T}$, we devise a finite algebra $\mathbf{A}(\mathcal{T})$ which enlarges A (in order to have enough distinct elements to code configurations) by adding finitely many elements and which expands $\mathbf{A}$ by adjoining operations to emulate the transitions between configurations, as well as to keep control of the finite subdirectly irreducible algebras. But the analysis of computation itself will go on in $\mathbf{A}(\mathcal{T})^X$ for some large set X [think of $X = \mathbb{Z}$].

We conceive of a Turing machine $\mathcal{T}$ as having finitely many internal states $0, 1, \ldots, m$. The machine is always launched in state 1 and we take 0 to be the unique halting state. The Turing machine $\mathcal{T}$ has a tape alphabet consisting of the symbols 0 and 1. The Turing machine itself is a finite collection of 5-tuples each of the form:

$$[i, \gamma, \delta, M, j]$$

This 5-tuple is the instruction, "If you are in state i and you are examining a tape square containing the symbol γ, then write the symbol δ on that square, move one square in the direction M (M must be either L for left or R for right), and pass into internal state j". We insist that no 5-tuple begin with 0 and that otherwise the machine must have exactly one instruction which begins $[i, \gamma, \ldots]$ for each state i other than the halting state 0 and each tape symbol γ.

We say Q is a *configuration* for a Turing machine $\mathcal{T}$ provided $Q = \langle t, n, i \rangle$ where $t \in \{0,1\}^{\mathbb{Z}}$, $n \in \mathbb{Z}$, and i is one of the states of $\mathcal{T}$. The idea is that at some stage of a computation, the tape of the machine looks like t, the machine is focussed on square n and is itself in state i.

A significant problem we have to resolve comes from the fact that machine computations, at any given stage, happen at a particular location on the tape, and that these locations are arranged in a sequence with only the adjacent locations available for the next step in the computation. Thus some elements of our "computation algebra" which are used to label those directed edges must also fall into a sequence of "tape locations". To make short work of this point we take the elements a_p of $\mathbf{Q}_{\mathbb{Z}}$ as a model of how elements fall into sequence. Looking at what we had to have in A to get these a_p's we recall:

$$\begin{array}{rccccccccc} \alpha_p\colon & \ldots & 1 & 1 & 1 & 2 & 3 & 3 & 3 & \ldots \\ \alpha_{p+1}\colon & \ldots & 1 & 1 & 1 & 1 & 2 & 3 & 3 & \ldots \\ \alpha_{p+2}\colon & \ldots & 1 & 1 & 1 & 1 & 1 & 2 & 3 & \ldots \end{array}$$

So in all our machine algebras we want a subset $U = \{1,2,3\}$ making elements like the ones above available in direct powers. To impose the precedence above in the direct power, we impose $3 \prec 3 \prec 2 \prec 1 \prec 1$ on U. We also use $\prec$ to denote the coordinatewise relation in any direct power of a machine algebra. Suppose $\mathbf{B} = \mathbf{A}(\mathcal{T})^X$. A subset $F \subseteq B$ is *sequentiable* provided

- $F \subseteq U^X$,
- 2 occurs at least once in f, for each $f \in F$, and
- $\prec$ gives F a structure isomorphic to some convex substructure of the ordered set of integers.

Since 2 may occur at several places in such an f, sequentiable sets can be more complex than $\{\alpha_p : p \in \mathbb{Z}\}$. For a fixed sequentiable set F the index set X falls into natural pieces that help us see the structure. Look at the following display of the four element sequentiable set $F = \{f_0, f_1, f_2, f_3\}$.

$$\begin{array}{rcccccccccccc} f_0\colon & 1 & 1 & 2 & 3 & 3 & 3 & 2 & 3 & 3 & 3 & 1 & 2 & 1 \\ f_1\colon & 1 & 1 & 1 & 2 & 3 & 3 & 1 & 3 & 3 & 3 & 1 & 1 & 1 \\ f_2\colon & 1 & 1 & 1 & 1 & 3 & 3 & 1 & 2 & 3 & 2 & 1 & 1 & 1 \\ f_3\colon & 1 & 1 & 1 & 1 & 3 & 2 & 1 & 1 & 3 & 1 & 1 & 1 & 1 \end{array}$$

Examining the 13 columns, we see that several are exactly the same. In this example the set X has 13 elements and some unspecified arrangement of these thirteen elements underlies the display above. But the particular arrangement of X is immaterial from the point of view of the algebra $\mathbf{A}(\mathcal{T})^X$. Thus we are free to rearrange X to make the precedence on F more transparent. Below is the result of such a rearrangement:

$$\begin{array}{rcccccccccccc} f_0\colon & 1 & 1 & 1 & 1 & 2 & 2 & 2 & 3 & 3 & 3 & 3 & 3 & 3 \\ f_1\colon & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 2 & 3 & 3 & 3 & 3 & 3 \\ f_2\colon & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 2 & 2 & 3 & 3 & 3 \\ f_3\colon & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 2 & 3 & 3 \end{array}$$

We have put all the columns consisting entirely of 1's to the left. Next we put all the columns beginning with 2 in position 0, then all columns with 2 in position 1, and so on. At the right we have placed all columns consisting entirely of 3's. Doing this, we see that there are only $6 = 4+2$ different kinds of columns possible:

$$\begin{array}{cccccc} 1 & 2 & 3 & 3 & 3 & 3 \\ 1 & 1 & 2 & 3 & 3 & 3 \\ 1 & 1 & 1 & 2 & 3 & 3 \\ 1 & 1 & 1 & 1 & 2 & 3 \end{array}$$

This means our sequentiable set F partitions the index set X into 6 blocks. The blocks can be labeled X_L for the set of all indices of columns that are constantly 1, X_R for the set of all indices of columns that are constantly 3, and X_n for the set of all indices where the necessarily unique 2 occurs at the n^{th} position.

To simplify the presentation a bit and make the pictures understandable, once a sequentiable set F has been specified, we will assume that X is arranged in such a line so that the set X_L is an initial (or left) segment, X_R is a final segment (or right) segment, and the each X_n is placed at the obvious position on the line. Since at its biggest, F can be indexed only by $\mathbb{Z}$, we can accommodate such a line like picture if we are willing to place X_L at $-\infty$ and X_R at $+\infty$.

Now let F be the four element sequentiable set above but with the columns collapsed to 6 and arranged as in the last display, and let $Q = \langle t, 2, i\rangle$ be a configuration. We code Q by

$$\begin{array}{rcccccc} \beta\text{:} & q^{0}_{i,t(2)} & q^{t(0)}_{i,t(2)} & q^{t(1)}_{i,t(2)} & M^{t(2)}_{i} & r^{t(3)}_{i,t(2)} & r^{0}_{i,t(2)} \\ \text{block:} & X_L & X_0 & X_1 & X_2 & X_3 & X_R \end{array}$$

This gives a real forest of superscripts and subscripts and the truth is that we will need a few more to get to full generality. However, we can decode it a bit. The q's mean "left of the reading head". The r's mean "to the right of the reading head". M locates where the machine reading head is. The index i specifies the state of the machine. The subscript $t(2)$ tells what symbol is written on the tape square scanned by the reading head. Finally, the indices $t(j)$ tell us what is printed on the corresponding square of the tape, unless it is too far off to the left (in X_L) or too far off to the right (in X_R), in which case we have used 0 as a default value (other choices would be okay). So reading across the superscripts is like reading across the tape. In this way, each component of β carries a lot of information about the configuration.

Now X in this example had 13 elements rather than 6, so the β above is too short. However, by duplicating the entries in β the correct number of times (e.g. the first entry $q^{0}_{i,t(2)}$ should occur 4 times while the last entry $r^{0}_{i,t(2)}$ should occur twice) we would get a β of the correct length. That $|X| = 13$ is immaterial. But our particular sequentiable set had only four elements, it was indexed with the convex set $\{0,1,2,3\}$, and we took $n = 2$ in our configuration. To get the general case, let I be any convex subset of $\mathbb{Z}$ and suppose that F is a sequentiable set indexed by I. Let $n \in I$ and let $Q = \langle t, n, i\rangle$ be a configuration. Then we use the β below as a code for Q and we say that β *codes* Q *over* F.

$$\beta(x) = \begin{cases} q^{0}_{i,t(n)} & \text{if } x \in X_L. \\ q^{t(j)}_{i,t(n)} & \text{if } x \in X_j \text{ and } j < n \text{ and } j \in I. \\ M^{t(n)}_{i} & \text{if } x \in X_j \text{ and } j = n \in I. \\ r^{t(j)}_{i,t(n)} & \text{if } x \in X_j \text{ and } n < j \in I. \\ r^{0}_{i,t(n)} & \text{if } x \in X_R. \end{cases}$$

Capturing the Transitions Between Configurations

To get a grip on how to handle the transition between configurations let $\mathbf{B} = \mathbf{A}(\mathcal{T})^{\mathbb{Z}}$ and let $F = \{\alpha_p : p \in \mathbb{Z}\}$. Then F is a sequentiable set indexed by $\mathbb{Z}$, and the partition imposed on $\mathbb{Z}$ by F consists of singleton sets $\{p\}$. Let $Q = \langle t, n, i\rangle$ be a configuration of $\mathcal{T}$, let $t(n) = \gamma$, and suppose that $[i, \gamma, \delta, L, j]$ is an instruction in $\mathcal{T}$. It also proves convenient to let $t(n-1) = \varepsilon$. Then $\mathcal{T}(Q) = \langle s, n-1, j\rangle$ is the configuration following Q in the computation of $\mathcal{T}$, where

$$s(k) = \begin{cases} \delta & \text{if } k = n, \\ t(k) & \text{otherwise.} \end{cases}$$

The configuration Q is coded over F by

$$\beta = \ldots \quad q_{i,\gamma}^{t(n-3)} \quad q_{i,\gamma}^{t(n-2)} \quad q_{i,\gamma}^{\varepsilon} \quad M_i^{\gamma} \quad r_{i,\gamma}^{t(n+1)} \quad r_{i,\gamma}^{t(n+2)} \quad r_{i,\gamma}^{t(n+3)} \quad \ldots$$

whereas the configuration $\mathcal{T}(Q)$ is coded over F by

$$\mathcal{T}(\beta) = \ldots \quad q_{j,\varepsilon}^{t(n-3)} \quad q_{j,\varepsilon}^{t(n-2)} \quad M_j^{\varepsilon} \quad r_{j,\varepsilon}^{\delta} \quad r_{j,\varepsilon}^{t(n+1)} \quad r_{j,\varepsilon}^{t(n+2)} \quad r_{j,\varepsilon}^{t(n+3)} \quad \ldots$$

$\mathcal{T}(\beta)$ differs from β in several ways. First, the two positions indexed by $n-1$ and n undergo a change of character from q to M and from M to r. Second, the remaining changes amount to changing γ to ε and i to j in various subscripts and superscripts. The idea is to effect this transition with a new operation for the machine instruction $[i,\gamma,\delta,L,j]$. Changes of the first kind have to do with two tape locations. Our new operation must combine the two location elements, α_{n-1} and α_n, with the configuration element β to produce the new configuration element $\mathcal{T}(\beta)$—our "instruction" operation should be ternary. To see what is needed to accomplish this, look at

$$\begin{array}{rccccccccc}
\alpha_{n-1} = \ldots & 1 & 1 & 2 & 3 & 3 & 3 & 3 & \ldots \\
\alpha_n = \ldots & 1 & 1 & 1 & 2 & 3 & 3 & 3 & \ldots \\
\beta = \ldots & r_{i,\gamma}^{t(n-3)} & r_{i,\gamma}^{t(n-2)} & r_{i,\gamma}^{\varepsilon} & M_i^{\gamma} & q_{i,\gamma}^{t(n+1)} & q_{i,\gamma}^{t(n+2)} & q_{i,\gamma}^{t(n+3)} & \ldots \\
\hline
\mathcal{T}(\beta) = \ldots & r_{j,\varepsilon}^{t(n-3)} & r_{j,\varepsilon}^{t(n-2)} & M_j^{\varepsilon} & q_{j,\varepsilon}^{\delta} & q_{j,\varepsilon}^{t(n+1)} & q_{j,\varepsilon}^{t(n+2)} & q_{j,\varepsilon}^{t(n+3)} & \ldots
\end{array}$$

The instruction $[i,\gamma,\delta,L,j]$ makes no reference to ε (the symbol written on square $n-1$ of the tape). Since our operation must act coordinatewise, we will build ε into the operation itself. So to each machine instruction we will associate two ternary operations, one for each of the two possible values of ε. Since the machine instructions for a fixed Turing machine $\mathcal{T}$ are determined by their first two components we will denote the operations corresponding to the machine instruction above by $F_{i\gamma\varepsilon}$. What must happen in $\mathbf{A}(\mathcal{T})$ to accomplish the transition above is

$$\begin{aligned}
F_{i\gamma\varepsilon}(1,1,r_{i,\gamma}^{\nu}) &= r_{j,\varepsilon}^{\nu} \\
F_{i\gamma\varepsilon}(3,3,q_{i,\gamma}^{\nu}) &= q_{j,\varepsilon}^{\nu} \\
F_{i\gamma\varepsilon}(2,1,r_{i,\gamma}^{\varepsilon}) &= M_j^{\varepsilon} \\
F_{i\gamma\varepsilon}(3,2,M_i^{\gamma}) &= q_{j,\varepsilon}^{\delta}
\end{aligned}$$

We would like to declare that in $\mathbf{A}(\mathcal{T})$ the operation $F_{i\gamma\varepsilon}$ results in the default value 0 except in the cases above. Ultimately, this won't do since we will find it necessary to introduce barred versions of all those q's, r's, and M's with all the attached subscripts and superscripts in order to control the finite subdirectly irreducible algebras. So we will have to revisit the definition of $F_{i\gamma\varepsilon}$. For the present, it is no great distortion to think that all the other values are 0.

A similar analysis of right-moving instructions leads the ternary operations $F_{i\gamma\varepsilon}$ being defined (with caveats about barred elements) in $\mathbf{A}(\mathcal{T})$ via

$$\begin{aligned}
F_{i\gamma\varepsilon}(1,1,r_{i,\gamma}^{\nu}) &= r_{j,\varepsilon}^{\nu} \\
F_{i\gamma\varepsilon}(3,3,q_{i,\gamma}^{\nu}) &= q_{j,\varepsilon}^{\nu} \\
F_{i\gamma\varepsilon}(2,1,M_i^{\gamma}) &= r_{j,\varepsilon}^{\delta} \\
F_{i\gamma\varepsilon}(3,H,q_{i,\gamma}^{\varepsilon}) &= M_j^{\varepsilon}
\end{aligned}$$

With this definition, in $\mathbf{A}(\mathcal{T})^{\mathbb{Z}}$

$$F_{i\gamma\varepsilon}(\alpha_n,\alpha_{n+1},\beta) = \mathcal{T}(\beta)$$

provided β is as above, ε is the symbol on tape square $n+1$, and $[i,\gamma,\delta,R,j]$ is an instruction of $\mathfrak{T}$. For a given Turing machine $\mathfrak{T}$, the definition of $F_{i\gamma\varepsilon}$ is unambiguous, since whether $F_{i\gamma\varepsilon}$ should be left or right moving can be determined from $\mathfrak{T}$, i, and γ.

These operations can be envisioned as edge operations, where, however, the edges representing a particular operation now have two labels.

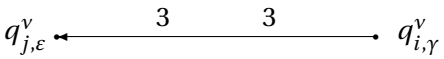

Left-Moving Case

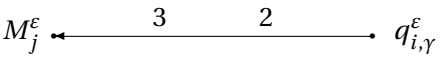

Right-Moving Case

Here is a useful fact, apparent in the diagrams above.

Fact 18. If λ basic translation on **A**($\mathfrak{T}$) associated with one of the operations $F_{i\gamma\varepsilon}$, and $\lambda(a)=\lambda(b)\neq 0$, then $a=b$. The same is true for every translation built only using the basic operations $F_{i\gamma\varepsilon}$, various choices of i,γ, and ε allowed.

On the basis of these definitions, we obtain the following very useful conclusion.

The Key Coding Lemma: *Let $\mathfrak{T}$ be a Turing machine, and let X be a set. Let F be a sequentiable set for* **A**$(\mathfrak{T})^X$ *and let i be a nonhalting state of $\mathfrak{T}$. Finally, let $\gamma,\varepsilon\in\{0,1\}$ and let f,g, and β be any elements of $A(\mathfrak{T})^X$.*

Then $F_{i\gamma\varepsilon}(f,g,\beta)=\mathfrak{T}(\beta)$ if

- *β codes a configuration Q over F,*
- *i and γ are the first two components of the $\mathfrak{T}$ instruction determined by Q,*
- *$f,g\in F$ with $f\prec g$ and these two elements refer to the two adjacent tape squares involved in the motion called for in the instruction,*
- *ε is the symbol in the square to which the reading head is being moved, and*
- *$\mathfrak{T}(\beta)$ codes the configuration $\mathfrak{T}(Q)$ over F;*

Otherwise 0 *occurs in $F_{i\gamma\varepsilon}(f,g,\beta)$.* □

LESSON 14

$\mathbf{A}(\mathcal{T})$ AND WHAT HAPPENS IF $\mathcal{T}$ DOESN'T HALT

The basic plan is to do for $\mathbf{A}(\mathcal{T})$ what we did for $\mathbf{A}$. We were able to prove for $\mathbf{A}$ three crucial things:

1. $\mathbf{Q}_{\mathbb{Z}}$ is in the variety generated by $\mathbf{A}$ (and hence that variety was inherently nonfinitely based and had a countably infinite subdirectly irreducible member).
2. Any finite subdirectly irreducible in the variety, except possibly a few very small ones, had a very well determined structure (in fact they were all embeddable into $\mathbf{Q}_{\mathbb{Z}}$).
3. There were no other infinite subdirectly irreducible algebras in the variety.

It was the second point that compelled us to adjoin additional elements and operations to our original 6-element algebra. Having done that, we had to revisit the first point to assure ourselves that the new elements and operations were innocuous. The third point depended on the first two and the Dziobiak-Quackenbush Theorem.

Proceeding along the same lines with $\mathbf{A}(\mathcal{T})$ we are able to do the following:

1. $\mathbf{Q}_{\mathbb{Z}}$ is in the variety generated by $\mathbf{A}(\mathcal{T})$, provided $\mathcal{T}$ does not halt.
2. In the event that $\mathcal{T}$ halts, the cardinality of any finite subdirectly irreducible can be bounded by a function of the size of $\mathcal{T}$ and the number of tape squares it visits before halting.
3. In the event that $\mathcal{T}$ halts, the variety generated by $\mathbf{A}(\mathcal{T})$ has no infinite subdirectly irreducible algebras.
4. In the event that $\mathcal{T}$ halts, the variety generated by $\mathbf{A}(\mathcal{T})$ is finitely based.

In the second point, at the cost of adding more elements and more operations to our 8-element algebra $\mathbf{A}$, we can ensure that any sequentiable set arising in the construction of a finite subdirectly irreducible cannot be large enough to accommodate the full halting computation. (The idea is that being able to reach a "halting configuration" would force the forbidden $(x \wedge y) \vee (x \wedge z)$ to be a polynomial.) Then we need to argue that bounding the size of sequentiable sets entails a bound on the subdirectly irreducible algebra itself. In the first point, after making an inessential modification to $\mathbf{Q}_{\mathbb{Z}}$ to make it into an algebra of the correct similarity type, it is the inaccessibility of the codes of halting configurations that ensures that the extra operations we had to add to accomplish the second point are innocuous. The third point is an

immediate consequence of Quackenbush's Theorem. The fourth point requires a tough proof due to Ross Willard.

The Algebra $\mathbf{A}(\mathcal{T})$

Let $\mathcal{T}$ be a Turing machine with states $0,1,\ldots,m$. The universe of the algebra $\mathbf{A}(\mathcal{T})$ is easiest to describe in pieces. For each of the $4m+4$ choices of $i = 0,1,\ldots,m$ and $\gamma,\delta \in \{0,1\}$, we need four distinct elements denoted by $q^{\delta}_{i,\gamma}$, $\overline{q^{\delta}_{i,\gamma}}$, $r^{\delta}_{i,\gamma}$, and $\overline{r^{\delta}_{i,\gamma}}$. For each of the $2m+2$ choices of $i = 0,1,\ldots,m$ and $\gamma \in \{0,1\}$, we need two elements denoted by M^{γ}_{i} and $\overline{M^{\gamma}_{i}}$. The unbarred versions were needed to code configurations. The barred versions help us control the finite subdirectly irreducible algebras. Let V be the set comprised of all $20m+20$ of these elements. We also let V_i denote the set of 20 elements of V whose first lower index is i. In particular, V_0 contains all the elements used in coding halting configurations. The universe of $\mathbf{A}(\mathcal{T})$ is just

$$A(\mathcal{T}) = \{0\} \cup U \cup W \cup V$$

where $U = \{1,2,3\}$ and $W = \{q, \bar{q}, r, \bar{r}\}$. Thus the size of $\mathbf{A}(\mathcal{T})$ is $20m+28$ where m is the number of nonhalting states of $\mathcal{T}$.

The old algebra $\mathbf{A}$ will be a subreduct of $\mathbf{A}(\mathcal{T})$. Indeed, we insist that $\wedge$ make $\mathbf{A}(\mathcal{T})$ into a height 1 meet-semilattice with least element 0, and that any product involving a new element results in 0. The definitions of the remaining old operations are changed little or not at all. Here are the J's:

$$J(x,y,z) = \begin{cases} x & \text{if } x = y \neq 0 \\ x \wedge z & \text{if } x = \bar{y} \in V \cup W \\ 0 & \text{otherwise.} \end{cases}$$

$$J'(x,y,z) = \begin{cases} x \wedge z & \text{if } x = y \neq 0 \\ x & \text{if } x = \bar{y} \in V \cup W \\ 0 & \text{otherwise.} \end{cases}$$

Along with the old S's we insert one more:

$$S_0(u,v,x,y,z) = \begin{cases} (x \wedge y) \vee (x \wedge z) & \text{if } u \in V_0, \\ 0 & \text{otherwise.} \end{cases}$$

$$S_1(u,v,x,y,z) = \begin{cases} (x \wedge y) \vee (x \wedge z) & \text{if } u \in \{1,3\}, \\ 0 & \text{otherwise.} \end{cases}$$

$$S_2(u,v,x,y,z) = \begin{cases} (x \wedge y) \vee (x \wedge z) & \text{if } u = \bar{v} \in V \cup W, \\ 0 & \text{otherwise.} \end{cases}$$

Along with the old U^0 we insert two new operations $U^1_{i\gamma\varepsilon}$ and $U^2_{i\gamma\varepsilon}$ for each of the $4m$ choices of i, γ, and ε, where i is a nohalting state:

$$U^0(x,y,z,w) = \begin{cases} xy & \text{if } xy = zw \neq 0 \text{ and } x = z \text{ and } y = w \\ \overline{xy} & \text{if } xy = zw \neq 0 \text{ and } x \neq z \text{ or } y \neq w \\ 0 & \text{otherwise.} \end{cases}$$

$$U^1_{i\gamma\varepsilon}(x,y,z,w) = \begin{cases} F_{i\gamma\varepsilon}(x,y,w) & \text{if } x \prec z \text{ and } F_{i\gamma\varepsilon}(x,y,w) \neq 0 \text{ and } y = z \\ \overline{F_{i\gamma\varepsilon}(x,y,w)} & \text{if } x \prec z \text{ and } F_{i\gamma\varepsilon}(x,y,w) \neq 0 \text{ and } y \neq z \\ 0 & \text{otherwise.} \end{cases}$$

$$U^2_{i\gamma\varepsilon}(x,y,z,w) = \begin{cases} F_{i\gamma\varepsilon}(y,z,w) & \text{if } x \prec z \text{ and } F_{i\gamma\varepsilon}(y,z,w) \neq 0 \text{ and } x = y \\ \overline{F_{i\gamma\varepsilon}(y,z,w)} & \text{if } x \prec z \text{ and } F_{i\gamma\varepsilon}(y,z,w) \neq 0 \text{ and } x \neq y \\ 0 & \text{otherwise.} \end{cases}$$

Finally, we need the $4m$ ternary operations $F_{i\gamma\varepsilon}$ introduced in Lecture 13 (but extended to accommodate the barred elements of V) and one further unary operation which serves to set up initial configurations:

$$I(x) = \begin{cases} q^0_{1,0} & \text{if } x = 1, \\ M^0_1 & \text{if } x = 2, \\ r^0_{1,0} & \text{if } x = 3, \\ 0 & \text{otherwise.} \end{cases}$$

Notice that for outputs other than 0, the operation I is one-to-one. In this way, the next fact is an extension of Fact 18

Fact 19. If λ is any translation of **A**($\mathcal{T}$) build only from the basic operations I and $F_{i\gamma\varepsilon}$, various choices of i, γ, and ε allowed, and $\lambda(a) = \lambda(b) \neq 0$, then $a = b$.

While all this is relatively intricate, the F's and the I plainly help us emulate the computations of the Turing machine. The role of the S's is to prevent certain kinds of elements from getting into the picture during the construction of finite subdirectly irreducible algebras. U^0 was crucial to get a kind of unique decomposition result for $\cdot$ in the finite subdirectly irreducible algebras. The U^1 and U^2 operations play a similar role in connection with the F operations.

What Happens If $\mathcal{T}$ Does Not Halt

Now we expand $\mathbf{Q}_{\mathbb{Z}}$ to the similarity type appropriate to $\mathcal{T}$ by insisting that all the following equations hold in the expansion:

$$U^0(x,y,z,w) \approx (xy) \wedge (zw) \qquad S_0(u,v,x,y,z) \approx 0$$
$$J(x,y,z) \approx x \wedge y \qquad S_1(u,v,x,y,z) \approx 0$$
$$J'(x,y,z) \approx x \wedge y \wedge z \qquad S_2(u,v,x,y,z) \approx 0$$
$$F_{i\gamma\varepsilon}(x,y,w) \approx 0 \qquad I(x) \approx 0$$
$$U^1_{i\gamma\varepsilon}(x,y,z,w) \approx 0 \qquad U^2_{i\gamma\varepsilon}(x,y,z,w) \approx 0$$

for all choices of i, γ and ε.

This sort of inessential expansion leaves its key properties intact: any locally finite variety to which (this expanded) $\mathbf{Q}_{\mathbb{Z}}$ belongs will be inherently nonfinitely based, and $\mathbf{Q}_{\mathbb{Z}}$ has a countably infinite subalgebra $\mathbf{Q}_{\omega}$ which is subdirectly irreducible.

Theorem 14.0.1. *If $\mathcal{T}$ does not halt, then $\mathbf{Q}_{\mathbb{Z}}$ belongs to the variety generated by* **A**($\mathcal{T}$). *In particular, if $\mathcal{T}$ does not halt, then* **A**($\mathcal{T}$) *is inherently nonfinitely based and the variety it generates has a countably infinite subdirectly irreducible algebra.*

Proof. We follow the pattern set in the proofs of Theorems **??** and 12.0.1. For each $p \in \mathbb{Z}$ we take $\alpha_p, \beta_p \in A(\mathcal{T})^{\mathbb{Z}}$ to be the same elements we used before:

$$\begin{array}{lccccccccc} \alpha_p := & \ldots & 1 & 1 & 1 & 2 & 3 & 3 & 3 & \ldots \\ \beta_p := & \ldots & r & r & r & q & q & q & q & \ldots \end{array}$$

where the change is taking place at the p^{th} position. Next we let $B_1 = \{\alpha_p : p \in \mathbb{Z}\} \cup \{\beta_p : p \in \mathbb{Z}\}$ and we take **B** to be the subalgebra of $\mathbf{A}(\mathcal{T})^{\mathbb{Z}}$ generated by B_1. Let B_0 denote the subset of B consisting of all those $\mathbb{Z}$-tuples in B which contain at least one 0. The set $\{\alpha_p : p \in \mathbb{Z}\}$ is sequentiable and consists of all the tuples in B belonging to $U^{\mathbb{Z}}$, since none of the operations of **A**($\mathcal{T}$) ever produces an element of U. Now for every $p \in \mathbb{Z}$

$$I(\alpha_p) := \quad \ldots \quad r^0_{1,0} \quad r^0_{1,0} \quad r^0_{1,0} \quad M^0_1 \quad q^0_{1,0} \quad q^0_{1,0} \quad q^0_{1,0} \quad \ldots$$

which gives the code of a configuration (the all-0 tape with the machine in state 1 reading square p). The $F_{i\gamma\varepsilon}$'s may now be applied, step by step, to produce the codes of further configurations reached as the computation of $\mathcal{T}$ proceeds. Plainly, all these codes of configurations belong to B. Let C denote the set of all these configuration codes. We will prove that $C \cup B_0 \cup B_1$ is a subuniverse of $\mathbf{A}(\mathcal{T})^{\mathbb{Z}}$, and therefore $B = C \cup B_0 \cup B_1$.

Now let Φ be the map defined from B to $A_{\mathbb{Z}}$ via

$$\Phi(x) = \begin{cases} a_p & \text{if } x = \alpha_p \text{ for some } p \in \mathbb{Z}, \\ b_p & \text{if } x = \beta_p \text{ for some } p \in \mathbb{Z}, \\ 0 & \text{otherwise} \end{cases}$$

We contend that Φ is a homomorphism from **B** onto $\mathbf{Q}_{\mathbb{Z}}$. To verify this, as well as that $C \cup B_0 \cup B_1$ is a subuniverse, requires us to examine the behavior of each of our operations on $C \cup B_0 \cup B_1$. For each operation in turn, we show that this set in closed and that Φ preserves the operation.

CASE 0: Evidently $0 = \ldots, 0, 0, 0, 0, \cdots \in B_0$ and so $\Phi(0) = 0$.

CASE $\wedge$: Evidently, $u \wedge v = u$ if $u = v$ and $u \wedge v \in B_0$ if $u \neq v$, for all $u, v \in C \cup B_0 \cup B_1$. Hence, our set is closed under $\wedge$ and $\Phi(u \wedge v) = \Phi(u) \wedge \Phi(v)$.

CASE $\cdot$: Clearly, $\alpha_p \cdot \beta_{p+1} = \beta_p$ for all $p \in \mathbb{Z}$, with all other $\cdot$-products resulting in elements of B_0. So our set is closed under $\cdot$ and Φ preserves $\cdot$.

CASE $F_{i\gamma\varepsilon}$: According to the Key Coding Lemma, the results of applying $F_{i\gamma\varepsilon}$ to members of $C \cup B_0 \cup B_1$ lie in $C \cup B_0$. Hence, $C \cup B_0 \cup B_1$ is closed under this operation and Φ preserves the operation.

CASE I: Applied to elements of $C \cup B_0 \cup B_1$, I produces only elements of $C \cup B_0$. Hence, $C \cup B_0 \cup B_1$ is closed with respect to I, and Φ preserves I.

Observe that no barred elements occur in any of the members of $C \cup B_1$. It follows that

$$\text{if } u, v \in C \cup B_0 \cup B_1 \text{ with } u(p) = v(p) \neq 0 \text{ or } u(p) = \overline{v(p)} \in V \cup W \text{ for all } p \in \mathbb{Z}, \text{ then } u = v. \tag{$\star$}$$

CASE J: Evidently, $J(x,y,z) \in B_0$ if $x \in B_0$ or $y \in B_0$ or $x \neq y$, according to $(\star)$. Otherwise, $J(x,y,z) = x$. This entails that $C \cup B_0 \cup B_1$ is closed under J and Φ preserves J.

CASE J': Likewise, $J'(x,y,z) \in B_0$ if $x \in B_0$ or $y \in B_0$ or $x \neq y$, according to $(\star)$. Otherwise, $J'(x,y,z) = x \wedge z$. This entails that $C \cup B_0 \cup B_1$ is closed under J and Φ preserves J.

CASE U^0: If $xy \in B_0$ or $zw \in B_0$ or $xy \neq zw$, then we have $U^0(x,y,z,w) \in B_0$ and $(xy) \wedge (zw) \in B_0$, for any elements $x, y, z, w \in C \cup B_0 \cup B_1$. On the other hand, if $xy = zw \notin B_0$ it must be that $x = z = \alpha_p$ and $y = w = \beta_{p+1}$ for some $p \in \mathbb{Z}$. In that case, $U^0(x,y,z,w) = xy = zx = (xy) \wedge (zw)$. Thus, $C \cup B_0 \cup B_1$ is closed under U^0 and Φ preserves U^0.

Observe that for $u, v \in C \cup B_0 \cup B_1$, we have $u \prec v$ only when $u = \alpha_p$ and $v = \alpha_{p+1}$ for some $p \in \mathbb{Z}$. In particular,

$$\text{With respect to } \prec, \text{ every element of } C \cup B_0 \cup B_1 \text{ has at most one predecessor and at most one successor.} \tag{$*$}$$

CASE $U^1_{i\gamma\varepsilon}$: In case $F_{i\gamma\varepsilon}(x,y,w) \in B_0$ or $x \not\prec z$, we have $U^1_{i\gamma\varepsilon}(x,y,z,w) \in B_0$. In the alternative case, it follows from the definition of $F_{i\gamma\varepsilon}$ that $x \prec y$. In view of $(*)$ it must be that $y = z$. So $U^1_{i\gamma\varepsilon}(x,y,z,w) = F_{i\gamma\varepsilon}(x,y,w) \in C$. Therefore, the application of $U^1_{i\gamma\varepsilon}$ always results in an element of $C \cup B_0$. Consequently, $C \cup B_0 \cup B_1$ is closed with respect to $U^1_{i\gamma\varepsilon}$ and Φ preserves this operation.

CASE $U^2_{i\gamma\varepsilon}$: This case is like the one above, but it exploits the uniqueness of predecessors instead of successors.

CASE S_0: Since $\mathcal{T}$ does not halt, the set $V_0^{\mathbb{Z}}$ is disjoint from $C \cup B_0 \cup B_1$. It follows that the application of S_0 always results in an element of B_0. Thus S_0 is preserved by Φ and $C \cup B_0 \cup B_1$ is closed with respect to S_0. **It should be noted that this is the sole place in the argument where the fact the $\mathcal{T}$ does not halt comes into play.**

CASE S_1: The set $\{1,3\}^{\mathbb{Z}}$ is disjoint from $C \cup B_0 \cup B_1$. It follows that the application of S_1 always results in an element of B_0. Thus S_1 is preserved by Φ and $C \cup B_0 \cup B_1$ is closed with respect to S_1.

CASE S_2: It follows from $\star$ that the application of S_2 always results in an element of B_0. Thus S_2 is preserved by Φ and $C \cup B_0 \cup B_1$ is closed with respect to S_2.

So $\mathbf{Q}_{\mathbb{Z}}$ belongs to the variety generated by $\mathbf{A}(\mathcal{T})$. □

LESSON 15

WHEN $\mathcal{T}$ HALTS: FINITE SUBDIRECTLY IRREDUCIBLE ALGEBRAS OF SEQUENTIABLE TYPE

Throughout this lecture we assume that $\mathcal{T}$ is a Turing machine that eventually halts when started on the all-0 tape. We denote by $\pi(\mathcal{T})$ the number of squares examined by $\mathcal{T}$ in the course of its computation. Thus $\pi(\mathcal{T})$ is the length of the stretch of tape which comes into use for this computation. Our ambition is to describe all the finite subdirectly irreducible algebras in the variety generated by $\mathbf{A}(\mathcal{T})$, or at any rate to bound their size. From the facts developed in Lectures 1 and 2 we already have a lot of information at our disposal. Once again we take $\mathbf{S}$ to be a finite subdirectly irreducible algebra in the variety and we fix a finite set T, $\mathbf{B}$, and θ, so that

- $\mathbf{B} \subseteq \mathbf{A}(\mathcal{T})^T$
- θ is strictly meet-irreducible in $\operatorname{Con}\mathbf{B}$.
- $\mathbf{S}$ is isomorphic to $\mathbf{B}/\theta$.
- T is as small as possible for representing $\mathbf{S}$ in this way.
- $|T| > 1$ (i.e. $\mathbf{S} \notin HS\mathbf{A}(\mathcal{T})$).

Among other things, we know that $(x \wedge y) \vee (x \wedge z)$ is not a polynomial of $\mathbf{B}$ (Fact 5). We also have an element $p \in B$ so that $(p, 0)$ is critical over θ. In Lecture 2 the analysis revealed that all the elements of S, except 0, arose from a unique longest factorization of p using the product $\cdot$. We want, loosely speaking, to do the same thing now; but the machine operations I and $F_{i\gamma\varepsilon}$ have to be considered along with $\cdot$. We will change the definition of B_1. Thus, the facts that grew out of our analysis of the old version of B_1 must be re-examined. Also, Fact 9 was proved using an analysis by cases, with one case for each basic operation. Now we have more operations. Finally, we have modified all the old operations by extending their domains, (in the case of J, J', and S_2, we have done this by treating the new elements in V like the elements in W). However, in all its essential features the old analysis can be carried forward.

We take B_0 to be the collection of all elements of B which contain at least one 0. In $\mathbf{B}$ the ranges of S_0, S_1, and S_2 lie entirely in B_0. Moreover, V_0^T and $\{1,3\}^T$ are disjoint from B and there are no elements $u, v \in B$ so that $u = \bar{v} \in (V \cup W)^T$. This is just a direct consequence of Fact 5.

Fact 20. Every sequentiable subset of B has fewer than $\pi(\mathcal{T})$ members.

Proof. By the Key Coding Lemma any large enough sequentiable set would allow us, using I and the $F_{i\gamma\varepsilon}$'s, to emulate in **B** the entire halting computation of $\mathcal{T}$, producing an element of V_0^T in B. Then, via S_2, $(x \wedge y) \vee (x \wedge z)$ would be a polynomial of **B**. □

Next we restate a part of Fact 8 in our expanded setting. The only difference is the insertion of V in the statement and the proof.

Fact 21. If $v \in B$ and $p(s) = v(s)$ or $p(s) = \overline{v(s)} \in V \cup W$ for all $s \in T$, then $p = v$. □

The next fact splits our analysis into two cases.

Fact 22. Either $p \in V^T$ or $p \in W^T$.

Proof. First notice that there must be a nonconstant unary polynomial f and $u \in B$ with $f(u) = p$ but $u \neq p$. Otherwise, it follows from Fact 3 that $B - \{p\}$ is a θ-class. This means that our subdirectly irreducible algebra **S** has only two elements, and indeed is isomorphic to a subalgebra of $\mathbf{A}(\mathcal{T})$. This contradicts our assumption that T has at least two elements.

Let λ be a nonconstant unary polynomial of least complexity so that for some $u \in B$ with $u \neq p$ we have $\lambda(u) = p$. Also fix such a u. Now the rest of the argument falls into cases according to the leading operation symbol of λ.

CASE $\wedge$: $\lambda(x) = \mu(x) \wedge r$. Then $p = \mu(u) \wedge r$ Since p is maximal, we conclude that $p = \mu(u)$. This leads to a violation of the minimality of λ.

CASE $\cdot$: The range of λ is included in $B_0 \cup W^T$. This means $p \in W^T$.

CASE I: The range of λ is included in $B_0 \cup V^T$. This means $p \in V^T$.

CASES $F_{i\gamma\varepsilon}$: The range of λ is included in $B_0 \cup V^T$. So $p \in V^T$.

CASES S_i: Impossible: the range of each S_i is included in B_0.

CASES $U^0, U^j_{i\gamma\varepsilon}$: These cases put $p \in W^T$ (for U^0) or $p \in V^T$ (for $U^j_{i\gamma\varepsilon}$'s).

CASE J: $\lambda(x) = J(\mu(x), r, s)$, or $\lambda(x) = J(r, \mu(x), s)$, or $\lambda(x) = J(r, s, \mu(x))$. Under the first alternative, $p = \lambda(u) = J(\mu(u), r, s) \leq \mu(u)$. Then $p = \mu(u) = r$ by Fact 21 and the maximality of p. This violates the minimality of λ. The same reasoning applies to the second alternative. So consider the last alternative. Then $p = J(r, s, \mu(u)) \leq r$. Then $p = r$, and so Fact 21 implies that $p = r = s$. But this means that $\lambda(x) = J(p, p, \mu(x)) = p$, and so λ is constant. This case is impossible.

CASE J': This is like the last case, but easier. □

S is of *sequentiable* type if $p \in W^T$ and of *machine* type otherwise.

Fact 23. Finite subdirectly irreducibles of sequentiable type have fewer than $2\pi(\mathcal{T})$ members.

Proof. We can just follow the old analysis for **A**, paying a modest amount of attention to the additional operations, and observing that a sequentiable set arises in a natural way.

Now $p \in W^T$. Let B_1 be the set of all factors of p with respect to $\cdot$. Now *all* our previously established facts hold, as is evident in all cases except for Fact 9. This fact asserts that, *if $u \in B$ and $\lambda(u) \in B_1$ for some nonconstant translation λ, then $u \in B_1$*. The proof of Fact 9 relied on a case-by-case analysis according to the leading operation symbol. To get a proof for Fact 9 in our expanded similarity type, we have to consider the operations $I, F_{i\gamma\varepsilon}, U^1_{i\gamma\varepsilon}, U^2_{i\gamma\varepsilon}$, and S_0. (Actually, there are also minor changes in the definitions of J, J', and S_2, which merit a small amount of attention not provided here.) All these cases are trivial because $\lambda(u) \notin B_1$ for any u if the leading operation is any of these, since $B_1 \subseteq U^T \cup W^T$.

As in our analysis for $\mathbf{A}$, we have $B_1 = \{a_0, a_1, \ldots, a_{n-1}\} \cup \{b_0, b_1, \ldots, b_n\}$ where $b_k = a_k b_{k+1}$ for all $k < n$ and $b_0 = p$. Also $B - B_1$ is the θ-class of 0, B_1 splits into singletons modulo θ, and $a_k \in U^T$ and $b_k \in W^T$ for all k. It remains to see that $\{a_k : k < n\}$ is a sequentiable set. Since $\pi(\mathcal{T})$ bounds the size of sequentiable sets, we would be finished. We need $a_k \prec a_{k+1}$ for all $k < n-1$. Let $t \in T$, and suppose first that $a_{k+1}(t) = 1$. Then $b_k(t) \in \{r, \overline{r}\}$, so $a_k(t) \in \{1,2\}$. Hence $a_k(t) \prec a_{k+1}(t)$. Next, suppose that $a_{k+1}(t) = 2$. Then $b_k(t) \in \{q, \overline{q}\}$, so $a_k(t) = 3$. Hence, $a_k(t) \prec a_{k+1}(t)$. Finally, suppose $a_{k+1}(t) = 3$. Then $b_k(t) \in \{q, \overline{q}\}$, so $a_k(t) = 3 \prec 3 = a_{k+1}(t)$. thus, $a_k \prec a_{k+1}$ and $\{a_k : k < n\}$ is sequentiable. □

LESSON 16

WHEN $\mathcal{T}$ HALTS: FINITE SUBDIRECTLY IRREDUCIBLE ALGEBRAS OF MACHINE TYPE

We now consider the case when the finite subdirectly irreducible algebra **S** introduced in Lecture 6 is of machine type. So we have $p \in V^T$. In this case, we let B_1 be the smallest subset of B which includes p and which is closed under the inverses of all the machine operations I and $F_{i\gamma\varepsilon}$. Hence,

$$B_1 = \{u : \lambda(u) = p \text{ for some nonconstant translation } \lambda \text{ of } \mathbf{A}(\mathcal{T}) \text{ built only from the machine operations}\}$$

It is easy to see that since $p \in V^T$, then $B_1 \subseteq U^T \cup V^T$. It also follows that if λ is a translation built up from the machine operations, and $\lambda(u) = p$, then all the coefficients of λ also belong to B_1.

Since we have now substantially altered the definition of B_1, we will need to re-examine Facts 8 and 9. Here is the new version of Fact 8. It is an immediate consequence of Fact 21 and Fact 19.

Fact 24. If $u \in B_1$ and $v \in B$ so that for all $s \in T$ either $u(s) = v(s)$ or $u(s) = \overline{v(s)} \in V \cup W$, then $u = v$.

Here is the new version of Fact 9. The statement has not changed, but the proof is different, accommodating the change in the definition of B_1.

Fact 25. If $u \in B$ and $\lambda(u) \in B_1$ for some nonconstant translation λ, then $u \in B_1$.

Proof. The proof is by induction on the complexity of λ. The initial step of the induction is obvious, since the identity function is the only simplest nonconstant translation. For the inductive step we take $\lambda(x) = \nu(\mu(x))$, where $\nu(x)$ is a basic translation and $\mu(x)$ is a translation with smaller complexity than λ. The work breaks down into cases according to the basic operation associated with ν.

CASE $\wedge$: $\lambda(x) = \mu(x) \wedge r$. But every element of B_1 is maximal with respect to the semilattice order. So $\lambda(u) = \mu(u) \in B_1$. Invoking the induction hypothesis for $\mu(x)$, we get $u \in B_1$.

CASE $\cdot$: This cannot happen since then the range λ would be included in $B_0 \cup W^T$, which is disjoint from B_1.

CASES $F_{i\gamma\varepsilon}$: Since $\nu(\mu(u)) = \lambda(u) \in B_1$, it follows from the definition of B_1, that $\mu(u) \in B_1$. Now the induction hypothesis applies.

CASE I: $\lambda(x) = I(\mu(x))$. By the definition of B_1, $\mu(u) \in B_1$. So the induction hypothesis applied.

CASE J: $\nu(x) = J(v, y, z)$, where x is one of v, y, and z, while the remaining two are coefficients. First, suppose x is either v or y. From Fact 24 and the maximality of the members of B_1 it follows that $\mu(u) =$

$\lambda(u) \in B_1$. So the induction hypothesis applies. Now suppose x is z and son v and y are coefficients. In this case, it follows from Fact 24 that $v = y = \lambda(u) \in B_1$. But this means that $\nu(x) = v$ and so λ is constant. That cannot happen.

CASE J': This case is easier than the last one and its discussion is omitted.

CASES S_0, S_1 AND S_2: Too easy—the range of λ would be included in B_0.

CASE U^0: This cannot happen since the range of λ would be included in $B_0 \cup W^T$, which is disjoint from B_1.

CASES $U^j_{i\gamma\varepsilon}$: $\nu(x) = U^j_{i\gamma\varepsilon}(v, y, z, w)$, where exactly one of v, y, z, and w is x and the reamining ones are coefficients, which we will regard as constant functions.

The other case being similar, we suppose that $j = 1$. Evidently, $\lambda(u)$ and $F_{i\gamma\varepsilon}(v(u), y(u), w(u))$ satisfy the hypotheses of Fact 24. So $\lambda(u) = F_{i\gamma\varepsilon}(v(u), y(u), w(u)) = F_{i\gamma\varepsilon}(v(u), z(u), w(u))$ (since also $y(u) = w(u)$) follows from the definition of $U^1_{i\gamma\varepsilon}$. So $v(u), y(u), z(u), w(u) \in B_1$, by the definition of B_1. So $\mu(u) \in B_1$ and the induction hypothesis applies. □

Here is the new version of Fact 10. Again, the statement is the same, but B_1 has a new meaning. The proof it like that for Fact 10, but it uses Fact 25 in place of Fact 9 and Fact 19 in place of Fact 6.

Fact 26. $u/\theta = \{u\}$ for each $u \in B_1$ and $0/\theta = B - B_1$.

Thus to bound the cardinality of **S** we need to bound $|B_1|$. This will be the focus of our efforts in the next lecture. However, here we can remark that in fact a complete analysis of finite subdirectly irreducible algebras of machine type, as well as those of sequentiable type, is at hand. This further analysis would describe the behavior of all the operations. We will not pursue this more detailed analysis, except to point out that all these subdirectly irreducible algebras are flat.

LESSON 17

WHEN $\mathcal{T}$ HALTS: BOUNDING THE SUBDIRECTLY IRREDUCIBLES

In this lecture we will complete our analysis of the subdirectly irreducible algebras generated by $\mathbf{A}(\mathcal{T})$ in the case when $\mathcal{T}$ halts. Fact 23 already provides a bound on the size of the finite subdirectly irreducible algebras of sequentiable type. The last lecture provided a description of the finite subdirectly irreducible algebras of machine type. Our next task is to bound the size of these algebras. So we continue to consider the case when $\mathbf{S}$ is of machine type.

We can suppose that no component of $p \in V^T$ is a barred element. (The basic reason is that the operations $F_{i\gamma\varepsilon}$ do not alter whether a symbol is barred. Hence the distribution of bars in any member of $B_1 \cap V^T$ is the same as the distribution of bars in p.) Now $B_1 \subseteq U^T \cup V^T$. Let $\Omega = B_1 \cap V^T$ and $\Sigma = B_1 \cap U^T$. Look first in more detail at Ω. We define Ω_n by the following recursion.

$$\begin{aligned}\Omega_0 &= \{p\}\\ \Omega_{n+1} &= \Omega_n \cup \{u \in B_1 : F_{i\gamma\varepsilon}(f,g,u) \in \Omega_n \text{ for some } f,g \in B \text{ and some } i,\gamma,\varepsilon\}\end{aligned}$$

Evidently, $\Omega = \bigcup_n \Omega_n$. We will say that $f \in U^T$ **matches** $v \in V^T$ provided for all $t \in T$

$$\begin{aligned}f(t) = 1 &\Leftrightarrow v(t) \text{ is a } r^{\nu}_{i\gamma}\\ f(t) = 2 &\Leftrightarrow v(t) \text{ is an } M^{\gamma}_{i}\\ f(t) = 3 &\Leftrightarrow v(t) \text{ is a } q^{\nu}_{i\gamma}\end{aligned}$$

Observe that every $v \in V^T$ matches exactly one $f \in U^T$. For each natural number n, we let $\Sigma_n = \{f \in \Sigma : f \text{ matches } v \text{ for some } v \in \Omega_n\}$. By referring to the definition of $F_{i\gamma\varepsilon}$, we have that the elements of the two element set $\{f,g\}$ match the elements of the two element set $\{u,v\}$ whenever $F_{i\gamma\varepsilon}(f,g,u) = v \in \Omega$ (the order in which this matching occurs depends on whether the underlying Turing machine instruction is right-moving or left-moving). It follows that $\Sigma = \bigcup_n \Sigma_n$.

Fact 27. Σ is a sequentiable set.

Proof. We argue by induction that Σ_n is sequentiable.

INITIAL STEP: Observe that Σ_0 has only one element. (Σ_0 cannot be empty, since then our subdirectly irreducible $\mathbf{S}$ would be in $HS\mathbf{A}(\mathcal{T})$.) Since $\Sigma_0 \subseteq B_1 \cap U^T$ and B is disjoint for $\{1,2\}^T$, we see that its element has to have H in at least one place. Thus, Σ_0 is a sequentiable set.

INDUCTIVE STEP: Suppose $h \in \Sigma_{n+1} - \Sigma_n$. Pick $u \in \Omega_{n+1} - \Omega_n$ so that h matches u. Further, pick $F_{i\gamma\varepsilon}, f, g$, and v so that $F_{i\gamma\varepsilon}(f,g,u) = v \in \Omega_n$. It does no harm to suppose that we have a left-moving operation. So g matches u and f matches v. It follows that $h = g$, that $f \in \Sigma_n$, and that $f \prec g$. By the inductive hypothesis, we have that Σ_n is sequentiable. Let us display Σ_n as

$$f_a \prec f_{a+1} \prec \ldots f_b$$

In the event that $f = f_b$ we have $\Omega_n \cup \{h\}$ sequentiable as desired. On the other hand, if $f = f_c$ for some $c < b$, then, in view of Fact 24, we know $U^1_{i\gamma\varepsilon}(f,h,f_{c+1},u) = F_{i\gamma\varepsilon}(f,h,u)$. So we would be able to conclude that $h = f_{c+1} \in \Sigma_n$, contrary to our choice of h. Reasoning in the same way, we see that it is not possible that Σ_{n+1} extends Σ_n on the right in any more elaborate way. Indeed, suppose $h' \in \Sigma_{n+1} - \Sigma_n$ and that $F_{i'\gamma'\varepsilon'}(f_b, g', u') = v' \in \Omega_n$, where h' matches u'. We take this operation to be left-moving. Then from $U^1_{i'\gamma'\varepsilon'}(f_b, h', h, u') = F_{i'\gamma'\varepsilon'}(f_b, h', u')$ we are able to conclude that $h = h'$.

Right-moving operations are handled in a way similar to what we just did for left-moving operations, but using $U^2_{i\gamma\varepsilon}$. □

Fact 28. Σ has fewer than $\pi(\mathcal{T})$ elements. □

To obtain a bound on the cardinality of Ω we must recall that the sequentiable set Σ partitions T into $T_L, T_a, \ldots, T_b, T_R$ where $\Sigma = \{f_a, \ldots, f_b\}$.

Fact 29. $u \restriction T_c$ is constant for each $u \in \Omega$ and each $c \in \{a, \ldots, b\}$.

Proof. The proof is accomplished in stages, each stage showing that more elements of Ω are constant on more T_c's until everything is accomplished. This proof needs some preliminary observations.

Suppose that $u \in \Omega_{n+1} - \Omega_n$ with $F_{i\gamma\varepsilon}(f_c, f_{c+1}, u) = v \in \Omega_n$. In this case we will say that u, c and $c+1$ become *active* at stage $n+1$. (We regard p as the only element active at stage 0 and no member of $c \in \{a, \ldots, b\}$ as active at stage 0.) The definition of $F_{i\gamma\varepsilon}$ entails that $u \restriction T_c, u \restriction T_{c+1}, v \restriction T_c$ and $v \restriction T_{c+1}$ are all constant. Moreover, for all d, $u \restriction T_d$ is constant if and only if $v \restriction T_d$ is constant. In checking this, it helps to notice that the relevant subscripts and superscripts can all be determined from $F_{i\gamma\varepsilon}$ and the related Turing machine instruction $[i, \gamma, \delta, M, j]$. Also, if $I(f) = u \in \Omega$, then $u \restriction T_d$ is constant for all d.

Now we argue by induction on n, that every member of Ω_n is constant on T_c for all c that have become active by stage n and that, for all d and all $v, v' \in \Omega_n$, $v \restriction T_d$ is constant if and only if $v' \restriction T_d$ is constant.

The initial step of the induction holds vacuously.

For the inductive step, suppose $u, u' \in \Omega_{n+1} - \Omega_n$ with

$$F_{i\gamma\varepsilon}(f_c, f_{c+1}, u) = v \in \Omega_n \qquad \text{and} \qquad F_{i'\gamma'\varepsilon'}(f_{c'}, f_{c'+1}, u') = v' \in \Omega_n$$

Now our preliminary observations give the conclusions that u and u' are constant on all the d's active by stage n as well as for $c, c', c+1$, and $c'+1$, some of which may have become active for stage $n+1$. Moreover, we also conclude that, for all d, u is constant on T_d if and only if v is constant on T_d if and only if v' is constant of T_d if and only if u' is constant on T_d. In this way, the inductive step is complete. □

Now we just count things to obtain:

Fact 30. Ω has no more than $2^s ms$ elements where $s = |\Sigma|$ and m is the number of nonhalting states of $\mathcal{T}$.

Proof. For each $u \in \Omega$ there are no more than s possibilities for $c \in \{a, \ldots, b\}$ so that $u(t) = M^\gamma_i$, for some i and some γ and all $t \in T_c$. Having fixed one of these possibilities there are m choices for i and two choices

for γ. Now for d with $a \leq d < c$ we must have a ν so that $u(t) = C^{\nu}_{i\gamma}$ for all $t \in T_d$. Thus for each such d there are no more than two possibilities for ν. Likewise, if $c < d \leq b$, then there is some ν so that $u(t) = D^{\nu}_{i\gamma}$ for all $t \in T_d$. Again, for each such d there are no more than two possibilities for ν. Thus, far we have bounded the number of possibilities for u by $2^s ms$, as desired—but we still have to examine what $u(t)$ is like when $t \in T_L \cup T_R$. Suppose $t \in T_L$. Then $f_c(t) = 1$ for all $c \in \{a, \ldots, b\}$. From the definition of the operations $F_{i\gamma\varepsilon}$, it follows that $u(t) = C^{\nu}_{i\gamma}$, where ν is determined by $p(t) = C^{\nu}_{i'\gamma'}$, and i and γ are the same subscripts that occur throughout u. So u is determined on T_L by our previous choices and by the structure of p. Likewise, u is determined on T_R. So the desired bound is established. □

Theorem 17.0.1. *If $\mathcal{T}$ halts, then the cardinality of any subdirectly irreducible member of the variety generated by $\mathbf{A}(\mathcal{T})$ is no greater than the maximum of 2π, $2^{(\pi-1)}m(\pi-1)+\pi$ and $20m+28$, where π is the number of tape squares used by $\mathcal{T}$ in its halting computation and m is the number of nonhalting states of $\mathcal{T}$; moreover, every subdirectly irreducible algebra in the variety is flat.* □

The $20m+28$ that occurs above is just the cardinality of $\mathbf{A}(\mathcal{T})$. It bounds the cardinalities of the subdirectly irreducibles that belong to $HS\mathbf{A}(\mathcal{T})$. The 2π bounds the cardinalities of the subdirectly irreducible algebras of sequentiable type. The $2^{(\pi-1)}m(\pi-1)+\pi$ bounds the cardinalities of the subdirectly irreducible algebras of machine type.

It is clear that much more was accomplished than just establishing the bound on subdirectly irreducible algebras given above. Our analysis is very close to a complete description (given a description of the behavior of $\mathcal{T}$) of all the subdirectly irreducible algebras, even in the case that $\mathcal{T}$ does not halt. The only way in which the hypothesis that $\mathcal{T}$ does not halt entered into consideration of the finite subdirectly irreducible algebras was in bounding their size. The analysis of their structure holds regardless. In the case that $\mathcal{T}$ does not halt, McKenzie describes how to carry this description of the finite subdirectly irreducible algebras up to the infinite subdirectly irreducibles, via an argument relying on Quackenbush's Theorem. His conclusion is that such varieties have residual character ω_1: while they have countably infinite subdirectly irreducible algebras, they have none of any larger cardinality.

Finally, we have in hand all the pieces of McKenzie's first undecidability result about finite algebras:

Theorem 17.0.2. *The set of finite algebras of finite type which generate residually very finite varieties is not recursive. Indeed, that set is recursively inseparable from the set of finite algebras of finite type which generate varieties of residual character ω_1.* □

INDEX