2-Semilattices: Residual Properties and Applications to the Constraint Satisfaction Problem

by

Ian Payne

A thesis
presented to the University of Waterloo
in fulfillment of the
thesis requirement for the degree of
Doctor of Philosophy
in
Pure Mathematics

Waterloo, Ontario, Canada, 2017

Examining Committee Membership

The following served on the Examining Committee for this thesis. The decision of the Examining Committee is by majority vote.

External Examiner	George McNulty Professor University of South Carolina
Supervisor	Ross Willard Professor University of Waterloo
Internal Member	Stanley Burris Distinguished Professor Emeritus Adjunct Professor University of Waterloo
Internal Member	Rahim Moosa Professor University of Waterloo
Internal-external Member	Chris Godsil Professor University of Waterloo

I hereby declare that I am the sole author of this thesis. This is a true copy of the thesis, including any required final revisions, as accepted by my examiners.

I understand that my thesis may be made electronically available to the public.

Abstract

Semilattices are algebras known to have an important connection to partially ordered sets. In particular, if a partially ordered set $(A, \leq)$ has greatest lower bounds, a semilattice $(A; \wedge)$ can be associated to the order where $a \wedge b$ is the greatest lower bound of a and b. In this thesis, we study a class of algebras known as 2-semilattices, which is a generalization of the class of semilattices. Similar to the correspondence between partial orders and semilattices, there is a correspondence between certain digraphs and 2-semilattices. That is, to every 2-semilattice, there is an associated digraph which holds information about the 2-semilattice. Making frequent use of this correspondence, we explore the class of 2-semilattices from three perspectives: (i) Tame Congruence Theory, (ii) the "residual character" of the class of 2-semilattices, and (iii), the constraint satisfaction problem.

Tame Congruence Theory, developed in [29], is a structure theory on finite algebras driven by understanding their prime congruence quotients. The theory assigns to each such quotient a type from **1** to **5**. We show that types **3**, **4**, and **5** can occur in the class of 2-semilattices, but type **4** can not occur in a finite simple 2-semilattice.

Classes of algebras contain "subdirectly irreducible" members which hold information about the class. Specifically, the size of these members has been of interest to many authors. We show for certain subclasses of the class of 2-semilattices that there is no cardinal bound on the size of the irreducible members in that subclass.

The "fixed template constraint satisfaction problem" can be identified with the decision problem hom($\mathbb{A}$) where $\mathbb{A}$ is a fixed finite relational structure. The input to hom($\mathbb{A}$) is a finite structure $\mathbb{B}$ similar to $\mathbb{A}$. The question asked is "does there exist a homomorphism from $\mathbb{B}$ to $\mathbb{A}$?" Feder and Vardi [22] conjectured that for fixed $\mathbb{A}$, this decision problem is either **NP**-complete or solvable in polynomial time. Bulatov [15] confirmed this conjecture in the case that $\mathbb{A}$ is invariant under a 2-semilattice operation. We extend this result.

Acknowledgements

There are many people I would like to thank for many different things. First of all, thanks to my parents, Rita and Scott for nurturing my love of math from when I was very young. Thanks to my siblings Anthony, Cheyanne, Katie, and Natalie, as well as my unwavering best friends Craig, Ditto, and Ryan. Your love and support made my life in Newfoundland as perfect an upbringing as anyone could hope for. Thanks to my uncles, Kirby and Milton, my aunts Nicol and Glennis, and their collective offspring Adam, Bianca, Ella, Evelyn, and Jeffrey. Because of you, I never really left home.

There isn't room on this page to list all of the new friends who made my time in graduate school as enjoyable as it was. In no particular order, thanks to Zack, Pat, Omar, Ale, Jordan, Parisa, Robert, Jess, Seda, Mike (Szestopalow), Matt, Farrah, and Mike (McTavish). Thanks Blake for being a great colleague, roommate, and above all, buddy. Thanks to Michael for the laughs and for sharing and supporting my love of elementary problems. Thanks to Ty for being there from the first day until the last one. The easiest way to clean our office may be with gas and a match. Thanks to Bill for the stories, counseling, and patience. Most of all, thanks to Carrie for tolerating me these last four years.

Thanks to Chris, George, Rahim, and Stan who read my thesis and provided valuable feedback. Finally I would like to thank my supervisor, Ross Willard. I have felt lucky to have your advice and support through this whole journey. Thanks for taking me to all the conferences, helping me when I needed it, not helping me when I didn't need it, and somehow always being in a good mood even at times when you must have wanted to slap me. I sincerely hope that our time working together isn't over.

Dedication

This is dedicated to my parents, Rita and Scott. Now that you have retired, you should have plenty of time to read it.

Table of Contents

List of Figures

Chapter 1

Introduction

One of the broad goals in universal algebra is to understand and classify algebras and classes of algebras called *varieties*. One such coarse classification was given by Hobby and McKenzie in [29]. This six-fold classification of locally finite varieties (defined in Chapter 2) has provided a popular way of restricting interesting problems to a more approachable context. One of these six classes is the class of varieties having the *congruence meet semidistributive* property, which will be defined in Chapter 2.

Perhaps the simplest congruence meet semidistributive variety is the variety of semilattices. A semilattice is a set A equipped with a binary operation $\wedge$ which is commutative, associative, and *idempotent*. By idempotent, we mean $a \wedge a = a$ for every $a \in A$. An important feature of semilattices is their connection with partial orders. Indeed, every semilattice $(A; \wedge)$ carries a natural partial order given by $a \leq b$ if and only if $a \wedge b = a$. As well, every partially ordered set with greatest lower bounds gives rise to a semilattice. In this case, $\wedge$ is the binary operation on the domain of the partial order which selects the greatest lower bound. Semilattices have received attention from researchers in a variety of areas of mathematics. Anderson and Ward in [1] and Rhodes in [50] discuss semilattices with topological motivation. Ellis [21], Papert [45], Nieminen [44], and Freese and Nation [26] have written about semilattices from are algebraic point of view. Jeavons et. al. [31] and Rehof and Torben [49] are papers on constraint satisfaction problems. Semilattices are well-studied and well-understood from many perspectives. For a survey of the theory of semilattices, see Chajda et. al [17].

From the perspectives of this thesis, semilattices are very easily understood. On the purely algebraic side, semilattices are simple in the sense that every semilattice embeds in a direct product of copies of the two-element semilattice. This was essentially shown for finite semilattices in [45] by Papert, but the proof given extends to the infinite case. When an algebra embeds in a product of finite algebras, it is called *residually finite*. Papert's result shows that semilattices are residually finite. Semilattices are also simple from the perspective of the constraint satisfaction or homomorphism problem. For a relational structure $\mathbb{A}$ with domain A, we say that $\mathbb{A}$ is invariant under a semilattice operation $\wedge$ on A if each relation of $\mathbb{A}$ is a subuniverse of some power of the algebra $(A; \wedge)$. If $\mathbb{A}$ is a relational structure which is invariant under a semilattice operation, the problem of determining whether or not a relational structure $\mathbb{B}$ similar to $\mathbb{A}$ has a homomorphism to $\mathbb{A}$ is decidable in polynomial time. This was first shown using arc-consistency by Jeavons et. al. in [31]. For a more in-depth discussion of arc-consistency, see Chapter 11 of Dechter [20].

This thesis studies a class of algebras called 2-semilattices, which generalizes the class of semilattices. The class of 2-semilattices is another concrete example of a congruence meet semidistrbutive variety. In the general theme of understanding congruence meet semidistributive varieties, the class of 2-semilattices is a natural place to look once one has understood the variety of semilattices. Indeed, Bulatov's understanding of the constraint satisfaction problem when $\mathbb{A}$ is invariant under a 2-semilattice operation in [15] was a key insight towards Barto and Kozik's solution to the so-called "bounded width conjecture" in

Figure 1.1: The operation table for $\mathbf{T}_3$

$*$	0	1	2
0	0	1	0
1	1	1	2
2	0	2	2

[3]. This will be discussed in Chapter 6.

A 2-semilattice is a set A equipped with a commutative, idempotent, binary operation $*$ which satisfies $a * (a * b) = a * b$ for every $a, b \in A$. Using the associative and idempotent identities, if $(A; \wedge)$ is a semilattice, then for any $a, b \in A$, we have $a \wedge (a \wedge b) = (a \wedge a) \wedge b = a \wedge b$. This shows that every semilattice is a 2-semilattice. We now give an example of a 2-semilattice which is not a semilattice. The algebra $\mathbf{T}_3$ in Example 1.0.1 will be mentioned several times throughout this thesis.

Example 1.0.1. Let $\mathbf{T}_3 = (T_3; *)$ be such that $T_3 = \{0, 1, 2\}$ and $*$ is defined in Figure 1.1. It can be seen immediately from the table that $*$ is commutative and idempotent and easily checked that $a * (a * b) = a * b$ for every $a, b \in T_3$. Therefore, $\mathbf{T}_3$ is a 2-semilattice. However, $0 * (1 * 2) = 0 * 2 = 0$, but $(0 * 1) * 2 = 1 * 2 = 2$, so $*$ is not associative. Thus, $\mathbf{T}_3$ is not a semilattice.

The algebra $\mathbf{T}_3$ in Example 1.0.1 satisfies the additional property that for every $a, b \in T_3$, $a * b \in \{a, b\}$. A 2-semilattice satisfying this property is often called a *tournament*. Any semilattice whose associated partial order has two incomparable elements is an example of a 2-semilattice which is not a tournament. Tournaments will be discussed briefly in Chapter 3.

The earliest explicit reference to 2-semilattices that I have been able to find is by Quackenbush in [47]. In [47] 2-semilattices are mentioned in the introduction and nowhere else in the paper. However, Quackenbush does point to an earlier paper he coauthored with Ježek [34] where they explore the class of *commutative directoids*. A 2-semilattice $\mathbf{A}$ is a commutative directoid if it satisfies $a * ((a * b) * c) = (a * b) * c$ for every $a, b, c \in A$. The class of commutative directoids lies strictly between the class of semilattices and the class of 2-semilattices. There has been some foundational work done on commutative directoids. They were introduced in [34], further explored by Gardner and Parmenter in [27], and have been mentioned by several other authors such as Chajda et. al. in [17] and Ježek and McNulty in [33].

In Maróti's PhD thesis [40], the main objects of study were tournaments. As mentioned earlier, the class of tournaments is properly contained in the class of 2-semilattices. Given a 2-semilattice $\mathbf{A}$, there is a natural class obtained by closing $\{\mathbf{A}\}$ under products, subalgebras, and homomorphic images. This class it known as the *variety generated by* $\mathbf{A}$. The answer to certain questions about this class boils down to understanding its subdirectly irreducible members. "Subdirectly irreducible" will be defined in Chapter 2. When $\mathbf{A}$ is a semilattice with more than one element in its domain, the generated variety is always the class of all semilattices. In this case, the only subdirectly irreducible member up to isomorphism is the two-element semilattice as shown by Papert in [45]. In [40], Maróti proved that if $\mathbf{A}$ is a tournament then there are only finitely many subdirectly irreducible members up to isomorphism in its generated variety and they are all finite.

Bulatov in [15] studied 2-semilattices from the point of view of the homomorphism problem mentioned at the end of the third paragraph. In particular, he showed that if $\mathbb{A}$ is a finite relational structure which is invariant under a 2-semilattice operation, then the path consistency algorithm correctly answers the question "Given a finite relational structure $\mathbb{B}$ similar to $\mathbb{A}$, is there a homomorphism from $\mathbb{B}$ to $\mathbb{A}$?"

The new results in this thesis up to and including Chapter 5 are purely algebraic.

Chapter 3 is a collection of basic properties of 2-semilattices, most of which were previously known. Chapter 4 is a more in-depth exploration of the structure of 2-semilattices. In particular, we give a structure theorem on minimal congruences in finite 2-semilattices. We then use this theorem to explore the "tame-congruence-theoretic" types, as defined by Hobby and McKenzie in [29], which occur in varieties of 2-semilattices. Tame Congruence Theory will be introduced in Section 4.2. In Chapter 5 we show that if a finite 2-semilattice has certain properties, its generated variety contains arbitrarily large subdirectly irreducible members. Chapter 6 is devoted to extending the result of Bulatov from [15] mentioned in the previous paragraph. We prove that if $\mathbb{A}$ is a finite relational structure which is invariant under a binary operation satisfying certain properties, then the homomorphism problem described in the previous paragraph can be answered in polynomial time. The properties required of the binary operation are strictly weaker than being a 2-semilattice operation.

Chapter 2 is an introduction to the basics of universal algebra that will be used in this thesis. The reader may wish to skip Chapter 2 and refer to it as needed. As mentioned above, Chapter 3 contains elementary facts about 2-semilattices. Results from Chapter 3 will be applied in Chapters 4, 5, and 6. Chapter 6 does not refer to any results or definitions in either Chapter 4 or Chapter 5, but Chapter 5 does use results from Chapter 4. Finally, we note that Section 6.5 does not really seem to naturally fit anywhere in this thesis. The notation and results introduced there will not be needed until Section 6.6, so it seems as appropriate a choice as any to include them just before they are needed.

Before proceeding to Chapter 2, we acknowledge that since November 2016, three proofs of the "dichotomy conjecture", which will be stated in Chapter 6, have been circulated. The first proof was circulated by Feder, Rafiey, and Kinne, the second by Zhuk, and the third by Bulatov. References to their respective write-ups on arXiv.org can be found in [48], [53], and [10], respectively. At the time of the writing of this thesis, all three proofs are unverified and unpublished. If true, these results will imply Theorem 6.6.5 and Corollary 6.6.9, the main results in Section 6.6. It therefore seems appropriate to note that the results in Section 6.6 had all been obtained by the summer of 2015.

Chapter 2

Universal Algebra

This chapter is intended to introduce the algebraic background needed to understand this thesis. For a more in-depth exposition of basic universal algebra, we point the reader to text books of Burris and Sankapanavar [16] and Bergman [4]. Both are well-written introductions to the subject. The reader may wish to skip this section and only consult it as needed.

For a set A and a non-negative integer n, an *n-ary operation* on A is a function, $f : A^n \to A$. The integer n is called the *arity* of f, and we say that f is n-ary. When $n = 1$ or 2, we will call f unary or binary, respectively. When $n = 0$, the set A^0 is identified with $A^{\varnothing}$, the set of functions from $\varnothing$ to A. By convention, there is exactly one such function, the "empty" function. What is important here is that A^0 contains one element, so any function $f : A^0 \to A$ chooses one element from A. For this reason, we think of a zero-ary operation on A as a distinguished element of A.

For a nonempty set I, we call a function $\rho : I \to \mathbb{N}$ a *similarity type*. We will frequently omit the word "similarity" and just say type. For a similarity type ρ, an *algebra of type ρ* is a pair $\mathbf{A} = (A; F)$ where A is a nonempty set called the *universe* of $\mathbf{A}$, and $F = (f_i : i \in I)$

is a set of operations on A indexed by I such that the arity of f_i is $\rho(i)$ for every $i \in I$. We will always denote the universe by a capital letter and the algebra by the corresponding bold letter. From this point on, we will sacrifice some formality in order to improve readability by thinking of a similarity type as a set of symbols with built-in arities. In this sense, if $\mathcal{F}$ is a similarity type, the operations of an algebra $\mathbf{A}$ of type $\mathcal{F}$ are, for each n-ary symbol $f \in \mathcal{F}$, an operation $f^{\mathbf{A}} : A^n \to A$. The superscripts are used as in the previous sentence to distinguish the operations corresponding to the same symbol in different algebras of the same type. For example, we think of the similarity type of a group as a binary operation symbol $\cdot$, a unary operation symbol $^{-1}$, and a zero-ary operation symbol, 1. A group $\mathbf{A}$ has a binary operation, $\cdot^{\mathbf{A}}$, and unary operation $^{-1^{\mathbf{A}}}$, and a distinguished constant $1^{\mathbf{A}}$. The superscripts will be omitted whenever the algebra is clear from context.

An algebra $\mathbf{A}$ is said to be *finite* if A is finite, and $\mathbf{A}$ is said to have *finite type* if the set of symbols in its type is finite. When a type $\mathcal{F}$ contains just one symbol f, we will denote an algebra $\mathbf{A}$ of type $\mathcal{F}$ by $(A; f^{\mathbf{A}})$ or $(A; f)$ rather than $(A; (f^{\mathbf{A}}))$.

Two algebras are *similar* if they have the same similarity type. As described in the previous paragraph, all groups are similar since they all can be thought of as having the symbol set $\mathcal{F} = \{\cdot, ^{-1}, e\}$ where $\cdot$ is binary, $^{-1}$ is unary, and 1 is zero-ary. For an algebra $\mathbf{A} = (A; F)$, the set F is called its set of *basic operations*. The *term operations* of an algebra consist of the basic operations, all projections, and all operations obtained by composition of these. From the perspective of first-order logic, a term operation of $\mathbf{A}$ is the natural interpretation as an operation of some term in the type of $\mathbf{A}$. For example, if $\mathbf{A}$ has type $\mathcal{F}$ and $\mathcal{F}$ contains a binary symbol f and a three-ary symbol g, then $\mathcal{F}$ has a term $fx_4gx_1x_3fx_5gx_2x_2x_3$. This gives rise to a five-ary term operation $h^{\mathbf{A}}$, defined by $h^{\mathbf{A}}(x_1, x_2, x_3, x_4, x_5) = f^{\mathbf{A}}(x_4, g^{\mathbf{A}}(x_1, x_3, f^{\mathbf{A}}(x_5, g^{\mathbf{A}}(x_2, x_2, x_3))))$ for every algebra $\mathbf{A}$ of type $\mathcal{F}$.

For an algebra of type $\mathcal{F}$, a *subuniverse* of $\mathbf{A}$ is a subset, $B \subseteq A$ which is closed under all operations of $\mathbf{A}$. A *subalgebra* of $\mathbf{A}$ is an algebra $\mathbf{B} = (B; G)$ similar to $\mathbf{A}$ where B is a nonempty subuniverse of $\mathbf{A}$ and $f^{\mathbf{B}} = f^{\mathbf{A}} \restriction_B$ for each symbol $f \in \mathcal{F}$. That is, $G = (f^{\mathbf{B}} : f \in \mathcal{F})$ where $f^{\mathbf{B}}$ is as defined in the previous sentence. We write $B \leq \mathbf{A}$ if B is a subuniverse of $\mathbf{A}$, and $\mathbf{B} \leq \mathbf{A}$ if $\mathbf{B}$ is a subalgebra of $\mathbf{A}$. We now give two definitions regarding subuniverses and subalgebras that will be referred to later.

Definition 2.0.1. For an algebra $\mathbf{A}$ and a subset $X \subseteq A$, the *subuniverse generated by* X is the intersection of all subuniverses of $\mathbf{A}$ which contain X. The subuniverse generated by X is denoted by $\mathrm{Sg}_{\mathbf{A}}(X)$. We say that an algebra $\mathbf{A}$ is *finitely generated* if there is a finite set $X \subseteq A$ such that $A = \mathrm{Sg}_{\mathbf{A}}(X)$.

It can be shown for any algebra $\mathbf{A}$ and $X \subseteq A$ that

$$\mathrm{Sg}_{\mathbf{A}}(X) = \{t(a_1, \ldots, a_n) : t \text{ is an } n\text{-ary term operation of } \mathbf{A} \text{ and } a_1, \ldots, a_n \in X\}.$$

See Theorem 4.32 in [4] for a proof of this fact.

Definition 2.0.2.

1. A subuniverse $B \leq \mathbf{A}$ is *absorbing* in $\mathbf{A}$ if $\mathbf{A}$ has an n-ary term operation t with $n \geq 2$ such that $t(a_1, a_2, \ldots, a_n) \in B$ whenever at least $n-1$ of $a_1, \ldots, a_n$ are in B. In this case, we write $B \lhd \mathbf{A}$.

2. An absorbing subuniverse $B \lhd \mathbf{A}$ is *proper* if it is not equal to A or $\varnothing$.

3. An algebra $\mathbf{A}$ is *absorption free* if it has no proper absorbing subuniverses.

If $\mathbf{A}$ and $\mathbf{B}$ are similar algebras, a *homomorphism* from $\mathbf{A}$ to $\mathbf{B}$ is a function $h : A \to B$ such that for every n-ary operation symbol f of arity n and any $a_1, \ldots, a_n \in A$,

$$h(f^{\mathbf{A}}(a_1, \ldots, a_n)) = f^{\mathbf{B}}(h(a_1), \ldots, h(a_n)).$$

In the case that h is surjective, $\mathbf{B}$ is called a *homomorphic image* of $\mathbf{A}$. As usual in algebra, a homomorphism which is a bijection is called an isomorphism. In this case, we say that $\mathbf{A}$ and $\mathbf{B}$ are isomorphic and write $\mathbf{A} \cong \mathbf{B}$.

For a family $(\mathbf{A}_u : u \in U)$ of similar algebras, the *direct product* or *product* $\mathbf{A} = \prod_{u \in U} \mathbf{A}_u$ has universe $\prod_{u \in U} A_u$, which is formally the set

$$\left\{ \sigma : U \to \bigcup_{u \in U} A_u : \sigma(u) \in A_u \text{ for each } u \in U \right\}.$$

For each n-ary operation symbol f in the type of the $\mathbf{A}_u$, the operation $f^{\mathbf{A}}$ is defined by $f^{\mathbf{A}}(\sigma_1, \ldots, \sigma_n)(u) = f^{\mathbf{A}_u}(\sigma_1(u), \ldots, \sigma_n(u))$. When U is finite, that is, $|U| = n$ for some integer n, the product can be visualized as the usual cartesian product consisting of n-tuples with coordinate-wise operations.

A *congruence* on an algebra $\mathbf{A}$ is an equivalence relation on its domain so that the operations of $\mathbf{A}$ act in a well-defined way on equivalence classes. This means that an equivalence relation θ is a congruence on $\mathbf{A}$ if for any n-ary operation f of $\mathbf{A}$ and any $(a_1, b_1), \ldots, (a_n, b_n) \in \theta$, we have that $(f(a_1, \ldots, a_n), f(b_1, \ldots, b_n)) \in \theta$ as well. We will often use the notation $a \stackrel{\theta}{\equiv} b$ to denote $(a, b) \in \theta$. Congruences will usually be denoted by Greek letters, with a notable exception in the next sentence. Every algebra $\mathbf{A}$ has two congruences, $0_A = \{(a, a) : a \in A\}$ and $1_A = A \times A$. Unless $|A| = 1$, these congruences are different. We say that an algebra $\mathbf{A}$ is *simple* when it has exactly two congruences. This implies that a simple algebra has at least two elements in its universe. Some authors, such as Burris and Sankappanavar in [16] also consider one-element algebras to be simple. We will stick to the convention that simple algebras have at least two elements and say that an algebra is *trivial* if and only if it has a one-element universe. If θ is a congruence on $\mathbf{A}$, the quotient $\mathbf{A}/\theta$ is the algebra similar to $\mathbf{A}$ defined as follows: The universe, denoted by A/θ, is the set of θ-classes. For each $a \in A$, we denote by a/θ the

θ-class containing a. For each n-ary operation symbol f, the operation $f^{\mathbf{A}/\theta}$ is defined by $f^{\mathbf{A}/\theta}(a_1/\theta, \ldots, a_n/\theta) = f^{\mathbf{A}}(a_1, \ldots, a_n)/\theta$. That this definition gives rise to well-defined operations is precisely the condition that distinguishes congruences among equivalence relations. For a homomorphism $h : \mathbf{A} \to \mathbf{B}$, the *kernel* of h is defined by

$$\ker(h) = \{(a_1, a_2) \in A^2 : h(a_1) = h(a_2)\}.$$

The kernel of h is an equivalence relation, and for any $(a_1, b_1), \ldots, (a_n, b_n) \in \ker(h)$ and n-ary operation $f^{\mathbf{A}}$ of $\mathbf{A}$, we have

$$\begin{aligned} h(f^{\mathbf{A}}(a_1, \ldots, a_n)) &= f^{\mathbf{B}}(h(a_1), \ldots, h(a_n)) \\ &= f^{\mathbf{B}}(h(b_1), \ldots, h(b_n)) \\ &= h(f^{\mathbf{A}}(b_1, \ldots, b_n)), \end{aligned}$$

which means $(f^{\mathbf{A}}(a_1, \ldots, a_n), f^{\mathbf{A}}(b_1, \ldots, b_n)) \in \ker(h)$. The first and third equalities in the above calculation are because h is a homomorphism, and the second is because $(a_i, b_i) \in \ker(h)$ for each i. This shows $\ker(h)$ is a congruence on $\mathbf{A}$. Furthermore, if h is a surjective homomorphism, $\mathbf{A}/\ker(h) \cong \mathbf{B}$. For a congruence θ on $\mathbf{A}$, θ is the kernel of the homomorphism $h : \mathbf{A} \to \mathbf{A} / \theta$ which sends $a \mapsto a/\theta$. This gives a correspondence between quotients and homomorphic images which is an extension of the same phenomenon in the setting of groups and rings. The reader may wish to verify that the congruences of a group are precisely the equivalence relations on its domain whose set of equivalence classes is the set of cosets of a fixed normal subgroup.

We now include a brief discussion of lattices. We will not frequently appeal to lattice theory, but the language of lattices will be used when discussing congruences. A lattice can be thought of as a structure with an order or as an algebra with two binary operations. Specifically, one definition of a lattice is a partial order in which every pair of elements has a least upper bound and a greatest lower bound. Another definition of a lattice is an algebra whose similarity type consists of two binary symbols, $\wedge$ and $\vee$, which satisfy some identities including commutativity and associativity. For a complete list of the defining identities of lattices, see Definition 1.7 in [4]. One of the ideas in lattice theory that will be important for us is that there is a correspondence between these two definitions. A lattice in the order theoretic sense can be assigned binary operations $\wedge$ and $\vee$ where $x \wedge y$ is the greatest lower bound of x and y, and $x \vee y$ is the least upper bound of x and y. A lattice in the algebraic sense can be assigned a partial order by $x \leq y$ if and only if $x \wedge y = x$ if and only if $x \vee y = y$. These processes are inverses of each-other. For this reason, we will think of a lattice as a structure with both a partial order $\leq$, and two operations $\wedge$ and $\vee$ which are compatible in the sense that $a \leq b$ if and only if $a \wedge b = a$ if and only if $a \vee b = b$.

Our main use of the language of lattices will be to discuss congruences. We now define congruence lattices, which will be mentioned frequently.

Definition 2.0.3. For an algebra $\mathbf{A}$, we denote by $\mathrm{Con}(\mathbf{A})$ its set of congruences. For congruences α and β on an algebra, we define $\alpha \leq \beta$ if and only if $\alpha \subseteq \beta$, $\alpha \wedge \beta = \alpha \cap \beta$ and $\alpha \vee \beta$ to be the transitive closure of $\alpha \cup \beta$.

It can be shown that $\leq$, $\wedge$ and $\vee$ give a lattice structure on $\mathrm{Con}(\mathbf{A})$ and satisfy all of the conditions in the last sentence of the previous paragraph. In particular, it can be shown that $\alpha \wedge \beta$ and $\alpha \vee \beta$ are congruences on $\mathbf{A}$. Therefore, we will often call $\mathrm{Con}(\mathbf{A})$ the *congruence lattice* of the algebra $\mathbf{A}$.

Definition 2.0.4. For an algebra $\mathbf{A}$ and congruences $\alpha, \beta \in \mathrm{Con}(\mathbf{A})$, we say $\alpha \prec \beta$ if the following conditions are satisfied:

1. $\alpha < \beta$, and
2. $\alpha \leq \gamma \leq \beta$ implies $\gamma \in \{\alpha, \beta\}$.

When $\alpha \prec \beta$, we use the terminology "α is a lower cover of β" and "β is an upper cover of α".

Definition 2.0.5. Let $\mathbf{A}$ be an algebra and $a, b \in A$. The *congruence generated by* (a, b), denoted $\mathrm{Cg}_{\mathbf{A}}(a, b)$ is the intersection of all congruences on $\mathbf{A}$ which contain (a, b).

If $\mathbf{A}$ is an algebra and $\alpha \leq \beta$ are congruences on $\mathbf{A}$, then the relation β/α on A/α defined by

$$\beta/\alpha = \{(u/\alpha, v/\alpha) : (u, v) \in \beta\}$$

is a congruence on the algebra $\mathbf{A}/\alpha$. Furthermore, we have the following generalization of the Correspondence Theorem for groups. For a proof of Theorem 2.0.6, see Theorem 3.6 from [4].

Theorem 2.0.6 (Correspondence Theorem)**.** *Suppose* $\mathbf{A}$ *is an algebra and* $\alpha \in \mathrm{Con}(\mathbf{A})$*. The set of congruences on* $\mathbf{A}/\alpha$ *is exactly the set* $\{\beta/\alpha : \alpha \leq \beta \in \mathrm{Con}(\mathbf{A})\}$*. Furthermore,* $\beta/\alpha = \gamma/\alpha$ *if and only if* $\beta = \gamma$*, and* $\beta/\alpha < \gamma/\alpha$ *if and only if* $\beta < \gamma$*.*

The Correspondence Theorem implies that a congruence $\alpha \in \mathrm{Con}(\mathbf{A})$ is maximal if and only if $\mathbf{A}/\alpha$ is a simple algebra. We will use this fact frequently.

We say that $\mathbf{R}$ is a *subdirect product* of $(\mathbf{A}_u : u \in U)$ if $\mathbf{R} \leq \prod_{u \in U} \mathbf{A}_u$ and $\mathrm{pr}_u(R) = A_u$ for every $u \in U$. In this case, we write $\mathbf{R} \leq_{\mathrm{sd}} \prod_{u \in U} \mathbf{A}_u$. When we are only referring to the universe of $\mathbf{R}$, we will write $R \leq_{\mathrm{sd}} \prod_{u \in U} \mathbf{A}_u$. If $R \leq_{\mathrm{sd}} \mathbf{A}_1 \times \mathbf{A}_2 \times \mathbf{A}_3$, we will occasionally use the fact that $\mathbf{R}$ is isomorphic to a subalgebra of $\mathrm{pr}_{1,2}(\mathbf{R}) \times \mathbf{A}_3$ where (a_1, a_2, a_3) is identified with $((a_1, a_2), a_3)$. In fact, the image of this isomorphism is subdirect in $\mathrm{pr}_{1,2}(\mathbf{R}) \times \mathbf{A}_3$. We will occasionally abuse notation and write $R \leq_{\mathrm{sd}} \mathrm{pr}_{1,2}(\mathbf{R}) \times \mathbf{A}_3$ in this and similar situations. If $h : \mathbf{B} \hookrightarrow \prod_{u \in U} \mathbf{A}_u$ is an embedding with $h(B) \leq_{\mathrm{sd}} \prod_{u \in U} \mathbf{A}_u$, h is called a *subdirect embedding.*

Definition 2.0.7. An algebra $\mathbf{B}$ is called *subdirectly irreducible* if for any subdirect embedding $h : \mathbf{B} \hookrightarrow \prod_{u\in U} \mathbf{A}_u$, there is some $u \in U$ with the property that $\mathrm{pr}_u \circ h : \mathbf{B} \to \mathbf{A}_u$ is an isomorphism.

As is pointed out in [4], the property of an algebra being subdirectly irreducible is completely encoded in its lattice of congruences. Fix an algebra $\mathbf{A}$ and a congruence $\theta \in \mathrm{Con}(\mathbf{A})$. We say that θ is *completely meet irreducible* if whenever $\bigwedge_{\alpha\in C} \alpha = \theta$ for $C \subseteq \mathrm{Con}(\mathbf{A})$, it must be that $\theta \in C$. The connection between subdirect irreducibility and the congruence lattice is that $\mathbf{A}/\theta$ is subdirectly irreducible if and only if θ is completely meet irreducible. In particular, $\mathbf{A}$ is subdirectly irreducible if and only if 0_A is completely meet irreducible. The statement "θ is completely meet irreducible" is equivalent to the existence of a congruence μ satisfying (i) $\theta < \mu$, and (ii), if $\theta < \alpha$ then $\mu \leq \alpha$. For a proof of this and more discussion on subdirectly irreducible algebras, see Section 3.3 of [4]. When θ is a completely meet irreducible congruence on an algebra $\mathbf{A}$, the congruence μ defined by $\mu = \bigwedge_{\theta<\psi} \psi$ is the unique upper cover of θ. See Definition 2.0.4 for the definition of "upper cover". Putting the facts in the previous few sentences together, we get that an algebra is subdirectly irreducible if and only if the intersection of its non-zero congruences is non-zero. For any algebra $\mathbf{A}$, let Θ be the set of its completely meet irreducible congruences. It can be shown using Zorn's Lemma that $\bigwedge \Theta = 0_A$. Because of this, the homomorphism $h : \mathbf{A} \hookrightarrow \prod_{\theta\in\Theta} \mathbf{A}/\theta$ is a subdirect embedding, where each $\mathbf{A}/\theta$ is subdirectly irreducible. This is a sketch of a proof that every algebra is isomorphic to a subdirect product of subdirectly irreducible algebras. This result is due to Birkhoff [8]. As with simple algebras, there is some inconsistency in the literature about whether or not one-element algebras are subdirectly irreducible. For example, Burris and Sankappanavar in [16] and Gardner and Parmenter in [27] consider one-element algebras as subdirectly irreducible. We will stick with the convention that one-element algebras are not subdirectly irreducible.

An *identity* is a universally quantified equation of terms. For example, consider an algebra $\mathbf{A}$ whose similarity type yields terms s and t. We say that $\mathbf{A}$ *satisfies* $s \approx t$ and write $\mathbf{A} \models s \approx t$ if $s^{\mathbf{A}} = t^{\mathbf{A}}$. This may seem nonsensical as $s^{\mathbf{A}}$ and $t^{\mathbf{A}}$ need not even be functions of the same variables. For example, perhaps $\mathbf{A}$ has a binary basic operation symbol f, $s = fx_3fx_1x_2$, and $t = fx_3fx_5x_4$. In this case, $s^{\mathbf{A}} = t^{\mathbf{A}}$ means for any assignment $(a_1, a_2, a_3, a_4, a_5) \in A^5$ of $(x_1, x_2, x_3, x_4, x_5)$, we have that $s^{\mathbf{A}}(a_1, a_2, a_3) = t^{\mathbf{A}}(a_3, a_4, a_5)$. In the previous sentence, we consider s and t as functions of $x_1, \ldots, x_5$. That is, $s = s(x_1, \ldots, x_5)$ and $t = t(x_1, \ldots, x_5)$. The variables that s *mentions syntactically* are x_1, x_2, and x_3, and s does not mention x_4 or x_5 syntactically. Similarly, t mentions x_3, x_4, and x_5 syntactically, but it does not mention x_1 or x_2 syntactically.

Returning to our discussion on identities, we say that a class $\mathcal{V}$ of similar algebras satisfies $s \approx t$ and write $\mathcal{V} \models s \approx t$ if every algebra in $\mathcal{V}$ satisfies $s \approx t$. This thesis will almost exclusively deal with algebras which are *idempotent*. This was defined for binary operations in the introduction, but here we give a more general definition.

Definition 2.0.8.

1. An operation t on a set A is *idempotent* if it satisfies $t(a, a, \ldots, a) = a$ for every $a \in A$.

2. An algebra is *idempotent* if all of its basic operations are idempotent.

3. A class of algebras is *idempotent* if each algebra it contains is idempotent.

We note that for any algebra $\mathbf{A}$, the statements "all basic operations of $\mathbf{A}$ are idempotent" is equivalent to the statement "all term operations of $\mathbf{A}$ are idempotent". One of many useful properties of idempotent algebras is stated in the next proposition.

Proposition 2.0.9. *Let $\mathbf{A}$ be an idempotent algebra and $\theta \in \mathrm{Con}(\mathbf{A})$. Every equivalence class of θ is a subuniverse of $\mathbf{A}$.*

The class of groups are a familiar example of a class which satisfies identities. In fact, the class of all groups is precisely the class of algebras with similarity type given by the symbol set $\{\cdot, ^{-1}, e\}$ having arities $2, 1$, and 0, respectively, that satisfy the familiar group identities. Such a defining set of identities is often called an *axiomatization*. A class of similar algebras which has an axiomatization is called a *variety*. Birkhoff in [7] showed that a class of similar algebras is a variety if and only if it is closed under the formation of products, subalgebras, and homomorphic images. We note that an axiomatization of a variety need not be finite or unique.

For a class $\mathcal{K}$ of similar algebras, the variety generated by $\mathcal{K}$ is defined to be the intersection of all varieties containing $\mathcal{K}$. In [7], Birkhoff proved that the variety generated by a set $\mathcal{K}$ of similar algebras is exactly the class of homomorphic images of subalgebras of products of members of $\mathcal{K}$. For this reason, we make the following definition:

Definition 2.0.10. Let $\mathcal{K}$ be a class of algebras. The variety generated by $\mathcal{K}$, denoted $\mathbf{HSP}(\mathcal{K})$ is the class of all algebras which are a homomorphic image of a subalgebra of a product of members of $\mathcal{K}$. When $\mathcal{K}$ consists of just one algebra $\mathbf{A}$, we abbreviate $\mathbf{HSP}(\{\mathbf{A}\})$ by $\mathbf{HSP}(\mathbf{A})$.

For varieties $\mathcal{V}$ and $\mathcal{W}$ of the same type, we denote by $\mathcal{V} \vee \mathcal{W}$ the variety generated by their union. We finish off with a definition concerning varieties.

Definition 2.0.11. A variety is called *locally finite* when each finitely generated algebra it contains is finite.

It is well known that every variety generated by finitely many finite algebras is locally finite. For a proof of this, see Theorem 3.49 in [4].

Chapter 3

Basic Properties of 2-Semilattices

The goal of this chapter is to formally define 2-semilattices and establish some of their important properties that will be used later. In Section 3.1, we associate to each 2-semilattice a digraph and prove Lemma 3.1.5 which is a collection of properties of this digraph. Section 3.2 establishes two properties of the variety of 2-semilattices: it is not locally finite, and it is meet semidistributive. “Locally finite” was defined in Definition 2.0.11, and meet semidistributivity is a property of congruence lattices which will be defined formally in Definition 3.2.2. The rest of the chapter uses the Absorption Theorem of Barto and Kozik to establish facts about subdirect products of finite 2-semilattices. Section 3.4 is a collection of results that were known to Bulatov in [15]. We include proofs here because many of Bulatov's proofs were either quite terse or omitted altogether. Most of these results will not be needed until Chapter 6

Definition 3.0.1. Let A be a set.

1. A 2-*semilattice operation* on A is a binary operation $* : A^2 \to A$ satisfying

    (a) $x * x \approx x$,

(b) $x * y \approx y * x$, and

(c) $x * (x * y) \approx x * y$.

2. An algebra is called a 2-*semilattice* if its only basic operation is a 2-semilattice operation.

3. A 2-semilattice $\mathbf{A}$ is called a *tournament* if for any $a, b \in A$, $a * b \in \{a, b\}$.

4. A 2-semilattice is called a *commutative directoid* if it satisfies $x*((x*y)*z) \approx (x*y)*z$.

5. The variety of 2-semilattices will be denoted by $\mathcal{S}$.

The symbol $*$ will usually be omitted and we will simply denote the operation by concatenation. Since 2-semilattice operations are not associative, care must be taken to include parentheses when composing a 2-semilattice operation with itself. To avoid some clumsiness, we adopt the convention that association takes place on the left when parentheses are omitted. For example, xyz will always mean $(xy)z$, and $xyzw$ will always mean $((xy)z)w$.

Given any algebra $\mathbf{A}$ and $a, b \in A$, recall from Definition 2.0.5 that $\mathrm{Cg}_{\mathbf{A}}(a, b)$ is the congruence on $\mathbf{A}$ generated by the pair (a, b). Proposition 3.0.2 is about congruences of 2-semilattices and is stated without proof. For proofs of more general versions of Proposition 3.0.2 parts (1) and (2), see Theorems 4.16 and 4.17 from [4]. Proposition 3.0.2 (2) is traditionally credited to Maltsev since it is immediately implied by results from [38].

Proposition 3.0.2. *Let* $\mathbf{A}$ *be a 2-semilattice,* θ *be an equivalence relation on* A*, and* $a, b \in A$*.*

1. *The relation* θ *is a congruence on* $\mathbf{A}$ *if and only if for every* $(u, v) \in \theta$ *and* $c \in A$*, the pair* $(cu, cv) \in \theta$*.*

2. *The congruence* $\mathrm{Cg}_{\mathbf{A}}(a,b)$ *is the transitive closure of*

$$\{(t(a,a_1,\ldots,a_n),t(b,a_1,\ldots,a_n)) : t \textit{ is a term operation of } \mathbf{A} \textit{ and } a_1,\ldots,a_n \in A\}.$$

It is worth mentioning that Proposition 3.0.2 (1) is true for any algebra whose only basic operation is binary, and Proposition 3.0.2 (2) is true of any algebra. The following is a result of Papert from [45] on semilattices. This is not necessarily the best place for it, but it will be used later and seems to fit here as well as anywhere.

Proposition 3.0.3. *The only subdirectly irreducible semilattice is the two element semilattice.*

3.1 The Digraph

We now define the digraph mentioned in the beginning of this chapter as well as some associated vocabulary.

Definition 3.1.1. Let $\mathbf{A}$ be a 2-semilattice. We define a digraph relation on A by $a \xrightarrow{\mathbf{A}} b$ if and only if $ab = b$.

1. A subset $X \subseteq A$ is *strongly connected* if for any $x, y \in X$ there are $z_1, \ldots, z_n \in X$ such that
$$x \xrightarrow{\mathbf{A}} z_1 \xrightarrow{\mathbf{A}} \cdots \xrightarrow{\mathbf{A}} z_n \xrightarrow{\mathbf{A}} y.$$

2. A subset $X \subseteq A$ is a *strongly connected component* of $\mathbf{A}$ if it is strongly connected and $X \subsetneq Y \subseteq A$ implies Y is not strongly connected.

3. We say that $\mathbf{A}$ is *strongly connected* if and only if A is a strongly connected component of $\mathbf{A}$.

4. A subset $X \subseteq A$ is *acyclic* if $x_1, \dots, x_n \in X$ with

$$x_1 \xrightarrow{\mathbf{A}} \cdots \xrightarrow{\mathbf{A}} x_n \xrightarrow{\mathbf{A}} x_1$$

implies $x_1 = x_2 = \cdots = x_n$.

5. We say $\mathbf{A}$ is *acyclic* if A is acyclic.

We will sometimes write $\longrightarrow$ rather than $\xrightarrow{\mathbf{A}}$ when it can be done without a loss of clarity. Notice that in the case that $\mathbf{A}$ is a semilattice, this digraph relation is exactly the associated partial order with the arrow pointing to the smaller element. The terminology in Definition 3.1.1 (1)-(3) is standard graph theoretic terminology for the digraph $(A, \xrightarrow{\mathbf{A}})$ imposed on the algebra $\mathbf{A}$. Our notation agrees with Bulatov's in [15], but Maróti defines $a \longrightarrow b$ if and only if $ab = a$ in [40]. This difference in notation is superficial. As we will see, the structure of $(A, \xrightarrow{\mathbf{A}})$ holds a lot of information about $\mathbf{A}$. In general, however, the operation of a 2-semilattice can not be recovered from its digraph alone. The following example gives two non-isomorphic 2-semilattices which have identical digraphs.

Example 3.1.2. We define 2-semilattices $\mathbf{A}_1$ and $\mathbf{A}_2$ as follows. The universes, A_1 and A_2 are both equal to $\{0, 1, 2, 3\}$. The operations, $*^{\mathbf{A}_1}$ and $*^{\mathbf{A}_2}$ are defined by the tables in Figure 3.1. It can be seen from the table that both $*^{\mathbf{A}_1}$ and $*^{\mathbf{A}_2}$ are commutative and idempotent. That they satisfy $x(xy) = xy$ can be easily checked. The associated digraph for both $\mathbf{A}_1$ and $\mathbf{A}_2$ is in Figure 3.2. As will be the case for all digraphs of 2-semilattices in this thesis, we omit all loops from the picture. We note that $\mathbf{A}_2$ is a commutative directoid referred to as a "fork" by Gardner and Parmenter in [27] and Ježek and McNulty in [34].

Notice that $\mathbf{A}_1$ is a semilattice, but $1 *^{\mathbf{A}_2} (2 *^{\mathbf{A}_2} 3) = 0$ while $(1 *^{\mathbf{A}_2} 2) *^{\mathbf{A}_2} 3 = 1$. Therefore, $\mathbf{A}_2$ is not associative, so $\mathbf{A}_1 \ncong \mathbf{A}_2$.

Figure 3.1: The operation tables for $\mathbf{A}_1$ and $\mathbf{A}_2$

$*^{\mathbf{A}_1}$	0	1	2	3
0	0	0	0	0
1	0	1	1	1
2	0	1	2	1
3	0	1	1	3

$*^{\mathbf{A}_2}$	0	1	2	3
0	0	0	0	0
1	0	1	1	1
2	0	1	2	0
3	0	1	0	3

.

Figure 3.2: The digraph of $\mathbf{A}_1$ and $\mathbf{A}_2$

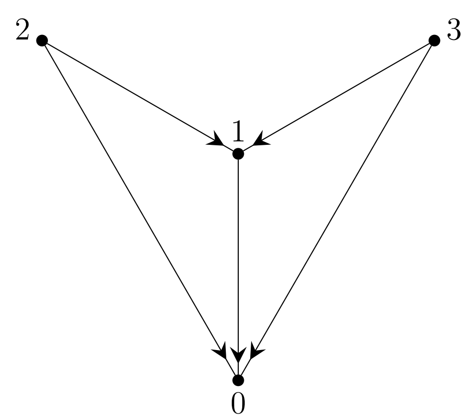

Suppose $\mathbf{A}$ is a tournament. Since $ab \in \{a, b\}$ for all $a, b \in A$, either $a \longrightarrow b$ or $b \longrightarrow a$, so the product ab can be determined from the direction of the arrow between a and b. In other words, if $\mathbf{A}$ is a tournament, then its operation is completely defined by its digraph. This is consistent with Example 3.1.2 since neither $\mathbf{A}_1$ nor $\mathbf{A}_2$ is a tournament as $2 *^{\mathbf{A}_i} 3 \notin \{2, 3\}$ for $i = 1, 2$.

The structure of strongly connected components of digraphs associated to 2-semilattices will be of importance throughout the thesis. The next definition and lemma are about strongly connected components.

Definition 3.1.3. Let $\mathbf{A}$ be a 2-semilattice and $\xrightarrow{\mathbf{A}}$ be the digraph relation from Definition 3.1.1.

1. Denote by $\overset{\mathbf{A}}{\sim}$ the equivalence relation on A whose classes are the strongly connected components with respect to $\xrightarrow{\mathbf{A}}$. We omit the over-set symbol and denote this relation by $\sim$ whenever it can be done without a loss of clarity.

2. Define a digraph relation $\xrightarrow{A/\sim}$ on $A/{\sim}$ by $a/{\sim} \xrightarrow{A/\sim} b/{\sim}$ if and only if there are $a' \in a/{\sim}$ and $b' \in b/{\sim}$ with $a' \xrightarrow{\mathbf{A}} b'$.

3. Write $a/{\sim} \succeq b/{\sim}$ if there is a directed walk from $a/{\sim}$ to $b/{\sim}$ in the digraph $(A, \xrightarrow{A/\sim})$.

Note that while $\sim$ is an equivalence relation on the domain of $\mathbf{A}$, it is not necessarily a congruence on $\mathbf{A}$. The relation $\xrightarrow{A/\sim}$ is simply a digraph relation on the set $A/{\sim}$.

Lemma 3.1.4. *Let $\mathbf{A}$ be a 2-semilattice and $\overset{\mathbf{A}}{\sim}$ and $\succeq$ be as in Definition 3.1.3.*

1. *For classes $a/{\sim}$ and $b/{\sim}$, $a/{\sim} \succeq b/{\sim}$ if and only if there is a directed walk in $(A, \xrightarrow{\mathbf{A}})$ from a' to b' for every $a' \in a/{\sim}$ and $b' \in b/{\sim}$.*

2. *The digraph $(A/{\sim}, \xrightarrow{A/\sim})$ is acyclic.*

3. *The relation* $\succeq$ *is a partial order.*

Proof.

1. Suppose $a/{\sim} \xrightarrow{A/\!\sim} b/{\sim}$. By definition, this means there are $a' \in a/{\sim}, b' \in b/{\sim}$ with $a' \xrightarrow{\mathbf{A}} b'$. Now let $a'' \in a/{\sim}$ and $b'' \in b/{\sim}$ be arbitrary. since $a/{\sim}$ and $b/{\sim}$ are both strongly connected, there are directed walks from a'' to a' and from b' to b''. This gives a directed walk from a'' to b''. If $a/{\sim} \succeq b/{\sim}$, there is, by definition, a directed walk in $(A/{\sim}, \xrightarrow{A/\!\sim})$ from $a/{\sim}$ to $b/{\sim}$. Repeatedly applying the argument for when $a/{\sim} \xrightarrow{A/\!\sim} b/{\sim}$ and concatenating the directed walks obtained shows that there is a directed walk from every element in $a/{\sim}$ to every element in $b/{\sim}$. Conversely, if there is a directed walk from every element in $a/{\sim}$ to every element in $b/{\sim}$, then there is a directed walk from a to b in $(A, \xrightarrow{\mathbf{A}})$. Let
$$a \xrightarrow{\mathbf{A}} a_1 \xrightarrow{\mathbf{A}} \cdots \xrightarrow{\mathbf{A}} a_n \xrightarrow{\mathbf{A}} b$$
be such a walk. Then
$$a/{\sim} \xrightarrow{A/\!\sim} a_1/{\sim} \xrightarrow{A/\!\sim} \cdots \xrightarrow{A/\!\sim} a_n/{\sim} \xrightarrow{\mathbf{A}} b/{\sim}$$
which is a directed walk from $a/{\sim}$ to $b/{\sim}$ in $(A/{\sim}, \xrightarrow{A/\!\sim})$, so $a/{\sim} \succeq b/{\sim}$.

2. Suppose
$$a_1/{\sim} \xrightarrow{A/\!\sim} a_2/{\sim} \xrightarrow{A/\!\sim} \cdots \xrightarrow{A/\!\sim} a_n/{\sim} \xrightarrow{A/\!\sim} a_1/{\sim}\ .$$
By part (1), this means, for each $i, j \leq n$, there is a directed walk from every $a \in a_i/{\sim}$ to every $b \in a_j/{\sim}$. Therefore,
$$\bigcup_{i=1}^{n} a_i/{\sim}$$
is strongly connected, so $a_1/{\sim} = a_2/{\sim} = \cdots = a_n/{\sim}$ because each $a_i/{\sim}$ is a strongly connected component.

3. Since $a \xrightarrow{\mathbf{A}} a$ for every $a \in A$, we have $a/\!\sim \xrightarrow{A/\sim} a/\!\sim$ and hence $a/\!\sim \succeq a/\!\sim$ for every $a \in A$. This shows that $\succeq$ is reflexive, it is antisymmetric by part (2). To see that $\succeq$ is transitive, suppose $a/\!\sim \succeq b/\!\sim \succeq c/\!\sim$. Let $a' \in a/\!\sim$ and $c' \in c/\!\sim$ be arbitrary. For any $b' \in b/\!\sim$, there are directed walks from a' to b' and from b' to c' in $(A, \xrightarrow{\mathbf{A}})$ by part (1). Concatenating these walks gives a walk from a' to c'. Applying part (1) again shows that $a/\!\sim \succeq c/\!\sim$. Therefore, $\succeq$ is a partial order.

$\square$

When $a/\!\sim \succeq b/\!\sim$, we view $b/\!\sim$ as being lower than $a/\!\sim$ in the partial order. When we refer to minimal components, we mean with respect to this order. Lemma 3.1.5 is a collection of properties of and relating to the digraph structure of a 2-semilattice. Many of the properties appear in [15] either implicitly or explicitly. Maróti also proved parts of Lemma 3.1.5 in [40] in the context of tournaments. We remind the reader that absorbing subuniverses were defined in Definition 2.0.2.

Lemma 3.1.5. *Let $\mathbf{A} \in \mathcal{S}$ be finite. The following hold for the digraph $(A, \xrightarrow{\mathbf{A}})$.*

1. *For any $a, b \in A$, $a \longrightarrow a$, $a \longrightarrow ab$ and $b \longrightarrow ab$.*

2. *$\mathbf{A}$ has a unique minimal strongly connected component denoted A' with respect to the partial order $\succeq$ from Definition 3.1.3. This component has the property that for any $b \in A$ there is $a \in A'$ such that $b \longrightarrow a$.*

3. *With A' as in (2), $a \in A'$ if and only if for every $b \in A$ there is a directed walk from b to a.*

4. *A' is an absorbing subuniverse of $\mathbf{A}$ with respect to $*$.*

5. *If $a, b \in A$ with $a \longrightarrow b$, then $(\{a, b\}; *)$ is a semilattice with smallest element b.*

6. *If* $\mathbf{B} \leq \mathbf{A}$*, then* $\xrightarrow{\mathbf{B}}$ *is* $(\xrightarrow{\mathbf{A}}) \cap B^2$*. That is, the digraph on* $\mathbf{B}$ *is the subdigraph induced from* $(A, \xrightarrow{\mathbf{A}})$ *on the domain of* $\mathbf{B}$*.*

7. *If* $\alpha \in \mathrm{Con}(\mathbf{A})$ *then* $a/\alpha \xrightarrow{\mathbf{A}/\alpha} b/\alpha$ *if and only if there are* $a' \in a/\alpha$ *and* $b' \in b/\alpha$ *with* $a' \xrightarrow{\mathbf{A}} b'$*.*

8. *Let* α *be a congruence on* $\mathbf{A}$*. If* $(A, \xrightarrow{\mathbf{A}})$ *is strongly connected, then* $(A/\alpha, \xrightarrow{\mathbf{A}/\alpha})$ *is strongly connected.*

Proof.

1. $a \longrightarrow a$ because $aa = a$, $a \longrightarrow ab$ since $a(ab) = ab$, and $b \longrightarrow ab$ because $b \longrightarrow ba$ and $*$ is commutative. We are using the identity $x * (x * y) \approx x * y$ from the definition of a 2-semilattice operation.

2. The partial order $\succeq$ is acyclic by Lemma 3.1.4 (2), so it has minimal elements by finiteness. By minimality, any such component U has the property that if $u \in U$ and $u \longrightarrow v$, then $v \in U$. Suppose U and V are minimal components. Fix $u \in U$ and $v \in V$. From (1), $u \longrightarrow uv$ and $v \longrightarrow uv$. From the previous remark, we have $uv \in U \cap V$. Since U and V are classes of an equivalence relation, we get $U = V$. We have shown that A' exists and is unique. Now pick $b \in A$ and $a' \in A'$, then set $a = ba'$. By the previous remark, $a \in A'$, and $b \longrightarrow a$ by (1).

3. If $a \in A'$ and $b \in A$, there is $c \in A'$ such that $b \longrightarrow c$ by (2). Since A' is strongly connected, there is a directed walk from c to a, so there is a directed walk from b to a. Conversely, suppose $a \in A$ and that there is a directed walk from b to a for every $b \in A$. In particular, for any $b \in A'$ there is a directed walk from b to a. By the minimality of A', every vertex in any such walk must be in A', so a is in A'.

4. Suppose $\{a, b\} \cap A' \neq \varnothing$. Then $a \longrightarrow ab$ and $b \longrightarrow ab$ by (1). Since either $a \in A'$ or $b \in A'$ there is some $c \in A'$ such that $c \longrightarrow ab$. This means $ab \in A'$ by minimality. Note that this proves A' is a subuniverse of $\mathbf{A}$ and that it is an absorbing subuniverse with respect to $t(x, y) = x * y$.

5. This is true since $aa = a$ and $ab = ba = bb = b$.

6. The notation $a \xrightarrow{\mathbf{B}} b$ means $a, b \in B$ and $ab = b$.

7. Suppose $a/\alpha \xrightarrow{\mathbf{A}/\alpha} b/\alpha$. By definition, this means $(a/\alpha)(b/\alpha) = b/\alpha$. Because of how quotient algebras are defined, $(ab)/\alpha = b/\alpha$, so there is $b' \in b/\alpha$ so that $ab = b'$. This means $ab' = a(ab) = ab = b'$, so $a \xrightarrow{\mathbf{A}} b'$. Taking $a = a'$ shows one direction of the implication. Now suppose there are $a' \in a/\alpha$ and $b' \in b/\alpha$ with $a' \longrightarrow b'$. Then

$$\begin{aligned} (a/\alpha)(b/\alpha) &= (ab)/\alpha \\ &= (a'b')/\alpha \\ &= b'/\alpha \\ &= b/\alpha, \end{aligned}$$

so $a/\alpha \xrightarrow{\mathbf{A}/\alpha} b/\alpha$.

8. Suppose $a, b \in A$. Since $(A, \xrightarrow{\mathbf{A}})$ is strongly connected, there are $a_1, \ldots, a_n$ such that

$$a = a_1 \xrightarrow{\mathbf{A}} a_2 \xrightarrow{\mathbf{A}} \cdots \xrightarrow{\mathbf{A}} a_n = b.$$

By (7),

$$a/\alpha = a_1/\alpha \xrightarrow{\mathbf{A}/\alpha} a_2/\alpha \xrightarrow{\mathbf{A}/\alpha} \cdots \xrightarrow{\mathbf{A}/\alpha} a_n/\alpha = b/\alpha$$

is a directed walk from a/α to b/α in $(A/\alpha, \xrightarrow{\mathbf{A}/\alpha})$.

$\square$

Figure 3.3: The digraph of $\mathbf{E}$

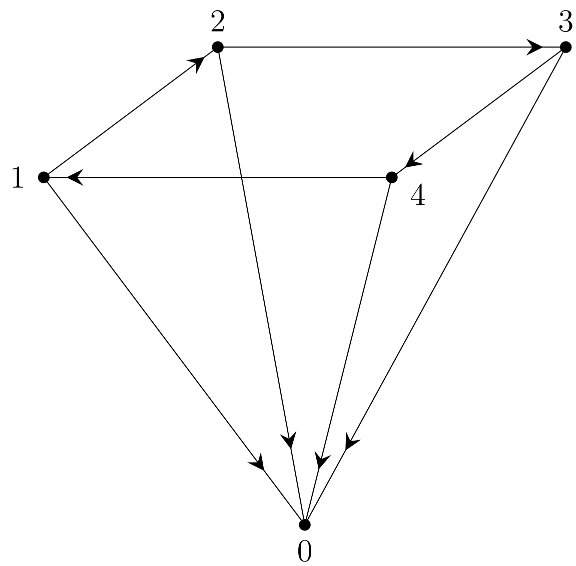

With the exception of the smallest strongly connected component, the strongly connected components of a 2-semilattice $\mathbf{A}$ are not, in general, subuniverses of $\mathbf{A}$. For an example of this, see Example 3.1.6. By Proposition 2.0.9, the equivalence classes of a congruence on a 2-semilattice $\mathbf{A}$ are all subuniverses of $\mathbf{A}$. Thus, Example 3.1.6 also shows, as mentioned earlier, that $\overset{\mathbf{A}}{\sim}$ will not be a congruence on $\mathbf{A}$ in general. In the case that $\mathbf{A}$ is a tournament, $\overset{\mathbf{A}}{\sim}$ is a congruence and the quotient by $\overset{\mathbf{A}}{\sim}$ is a semilattice. For a proof of this, see Lemma 4.2 in Maróti's doctoral thesis [40].

Example 3.1.6. We define an algebra $\mathbf{E}$ with universe $E = \{0, 1, 2, 3, 4\}$, $1 * 3 = 2 * 4 = 0$, and all other products can be deduced from its associated digraph in Figure 3.3. Note that the strongly connected components of $(E, \longrightarrow)$ are $\{0\}$ and $\{1, 2, 3, 4\}$, the latter of which is not closed under $*$.

There are several standard "products" in graph theory. For clarity on what we mean by a product of digraphs, we include the following definition:

Definition 3.1.7. Let $(D_i : i \in I)$ be a family of digraphs. By $\prod_{i\in I} D_i$ we mean the digraph whose domain is the Cartesian product of the domains, and $\mathbf{v} \longrightarrow \mathbf{w}$ if and only if $\mathbf{v}(i) \longrightarrow \mathbf{w}(i)$ for each i.

This is the usual product of relational structures from a model theoretic point of view, and it is the tensor product from a graph theoretic point of view. For us, the product of graphs will always be the product in Definition 3.1.7.

Lemma 3.1.8. *Suppose $(\mathbf{A}_i : i \in I)$ is a family of 2-semilattices and let $\mathbf{A} = \prod_{i\in I} \mathbf{A}_i$.*

1. *The graph $(A, \xrightarrow{\mathbf{A}})$ equals $\prod_{i\in I}(A_i, \xrightarrow{\mathbf{A}_i})$.*
2. *If each $\mathbf{A}_i$ is strongly connected and there is some $n \in \omega$ such that $|A_i| \leq n$ for each $i \in I$, then $\mathbf{A}$ is strongly connected.*

Proof.

1. The two digraphs in question both have domain equal to $\prod_{i\in I} A_i$, so it is enough to show that they have the same arrows. Fix $\mathbf{a}, \mathbf{b} \in A$. By Definition 3.1.1, $\mathbf{a} \xrightarrow{\mathbf{A}} \mathbf{b}$ if and only if $\mathbf{ab} = \mathbf{b}$. Because of the way products of algebras are defined, this is equivalent to $\mathbf{a}(i)\mathbf{b}(i) = \mathbf{b}(i)$ for each $i \in I$, which is equivalent to $\mathbf{a}(i) \xrightarrow{\mathbf{A}_i} \mathbf{b}(i)$ for each i.
2. Suppose $f, g \in \mathbf{A}$. Since the digraph $(A_i, \xrightarrow{\mathbf{A}_i})$ is reflexive for each i and $|A_i| \leq n$, there is a directed walk,

$$f(i) \xrightarrow{\mathbf{A}_i} a_1^i \xrightarrow{\mathbf{A}_i} a_2^i \xrightarrow{\mathbf{A}_i} \cdots \xrightarrow{\mathbf{A}_i} a_{n-1}^i \xrightarrow{\mathbf{A}_i} g(i)$$

from $f(i)$ to $g(i)$ in $(A_i, \xrightarrow{\mathbf{A}_i})$ of length exactly n. For each $j = 1, \ldots, n-1$, if we define $h_j \in A$ by $h_j(i) = a^i_j$, we have that

$$f \xrightarrow{\mathbf{A}} h_1 \xrightarrow{\mathbf{A}} \cdots \xrightarrow{\mathbf{A}} h_n \xrightarrow{\mathbf{A}} g,$$

so $\mathbf{A}$ is strongly connected.

$\square$

In particular, Lemma 3.1.8 shows that any power of a finite strongly connected 2-semilattice is strongly connected.

3.2 Local Finiteness and Meet Semidistributivity

This thesis will be mainly concerned with varieties generated by a finite 2-semilattice. It was noted that such varieties are locally finite after Definition 2.0.11. We now show that the variety of all 2-semilattices is not locally finite. The algebra $\mathbf{A}$ in Example 3.2.1 is a finitely-generated 2-semilattice which is infinite. Figure 3.4 contains a partial picture of the digraph of $\mathbf{A}$. One way to visualize the digraph is in "layers". The elements of $\mathbf{A}$ are denoted by the nonnegative integers, and the layers are $\{0, 1, 2\}, \{3, 4, 5\}, \{6, 7, 8\}, \ldots$. There is an arrow from n to $n+1$ for every n, so we think of the directed path $0 \longrightarrow 1 \longrightarrow 2 \longrightarrow \cdots$ as spiraling to the right through the layers. The arrows in this path are drawn thicker than the others. Additionally, there is an arrow from i to j whenever $i < j$ and either $i \equiv j \pmod 3$ or $i + 1 \equiv j \pmod 3$. We have omitted from the picture in Figure 3.4 all such arrows with $i \equiv j \pmod 3$, but included some of those with $i + 1 \equiv j \pmod 3$. There is also a loop on each vertex which has been omitted from the picture.

Figure 3.4: A finitely-generated infinite 2-semilattice

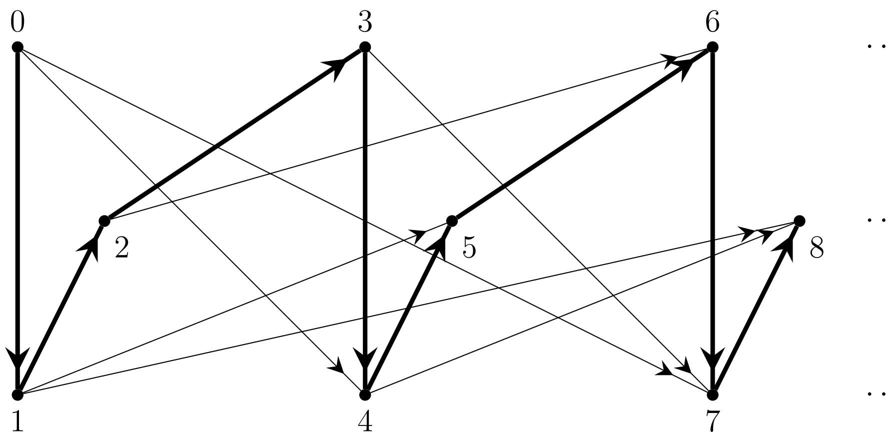

Example 3.2.1. Let $A = \omega$, the natural numbers and define $f : A^2 \to \{0, 1, 2\}$ by $f(x, y) = (x \bmod 3) * (y \bmod 3)$ where $*$ is the 2-semilattice operation of $\mathbf{T}_3$ from Definition 1.0.1. Now define a binary operation $\cdot : A^2 \to A$ as follows:

$$x \cdot y = \begin{cases} \max\{x, y\} & \text{if } x \equiv y \pmod 3 \\ \text{z} & \begin{array}{l} \text{where } z \geq x, y \text{ is minimal such that} \\ z \equiv f(x, y) \pmod 3 \text{, otherwise.} \end{array} \end{cases}$$

Let $\mathbf{A}$ be the algebra $(A, \cdot)$. We will show that $\mathbf{A}$ is a finitely generated 2-semilattice. Since $A = \omega$ is infinite, this will show that the variety of all 2-semilattices is not locally finite. Since f is commutative, it follows that $\cdot$ is commutative. As well, the definition of the operation $\cdot$ says that $x \cdot x = \max\{x, x\} = x$, so $\cdot$ is idempotent. We now fix $x, y \in A$. If $x \cdot y = x$, then $x \cdot (x \cdot y) = x \cdot x = x = x \cdot y$. If $x \cdot y = y$, then $x \cdot (x \cdot y) = x \cdot y$. Therefore, the only way $\mathbf{A}$ can fail identity (3) in Definition 3.0.1 (1) is if $x \cdot y \notin \{x, y\}$. This means $x \not\equiv y \pmod 3$, so $x \cdot y = z$ where $z \geq x, y$ is the smallest for which $z \equiv f(x, y)$ $\pmod 3$, and since $z \notin \{x, y\}$, $z > x, y$. Because of how $*^{\mathbf{T}_3}$ is defined, we have that

$f(x, y) \equiv x \pmod 3$ or $f(x, y) \equiv y \pmod 3$. Suppose the first occurs and that $x = n$ for some $n \geq 0$. Then $z = n + 3k$ for some $k \geq 1$, but since $f(x, y) \equiv x \pmod 3$, we have that $y \equiv n - 1 \pmod 3$. By the minimality of z, it must be that $y = z - 1$. We can now calculate $x \cdot (x \cdot y) = n \cdot (n + 3k) = n + 3k = z = x \cdot y$. If $f(x, y) \equiv y \pmod 3$, then by the same reasoning as the previous case, we have $y = n, z = n + 3k$ for $k \geq 1$, and $x = n + 3k - 1$. In this case, $x \cdot (x \cdot y) = (n + 3k - 1) \cdot (n + 3k)$. Since $n + 3k - 1 \not\equiv n + 3k \pmod 3$, the definition of $\cdot$ implies this product is equal to $n + 3k$ which is $z = x \cdot y$.

To see that $\mathbf{A}$ is finitely generated, fix n and suppose $n \equiv 0 \pmod 3$. Then $n + 2 \equiv 2 \pmod 3$, so $n \cdot (n + 2) \equiv 0 * 2 = 0 \pmod 3$. The smallest integer which is $0 \pmod 3$ and at least as large as n and $n + 2$ is $n + 3$. Therefore, $n \cdot (n + 2) = n + 3$. A similar argument holds when $n \equiv 1, 2 \pmod 3$. This means $n+3$ is in the subalgebra generated by $\{n, n+1, n+2\}$ for any n. Setting $n = 0$, this inductively shows that A is the subuniverse generated by $\{0, 1, 2\}$. Thus, $\mathbf{A}$ is finitely generated. This algebra will be used to provide another counterexample in Section 5.2

We now define congruence meet semidistributivity and prove that the variety of 2-semilattices is congruence meet semidistributive.

Definition 3.2.2. Let $\mathbf{A}$ be an algebra. We say that $\mathbf{A}$ is *congruence meet semidistributive* if for any congruences α, β, γ, and δ of $\mathbf{A}$, if $\alpha \wedge \beta = \alpha \wedge \gamma = \delta$, then $\alpha \wedge (\beta \vee \gamma) = \delta$. A variety is *congruence meet semidistributive* if every algebra it contains is congruence meet semidistributive.

We will often drop the word "congruence" and simply say that an algebra or a variety is meet semidistributive.

Theorem 8.1 from [35] is a ten-fold characterization given by Kearnes and Kiss of meet semidistributivity of a variety. Much of the theorem was shown to be true in the locally finite case by Hobby and McKenzie in Theorem 9.10 of [29]. To state either of these two theorems would be a substantial digression from the material in this thesis. Instead, we note that as an immediate consequence of Theorem 8.1 from [35], to show that $\mathcal{S}$ is meet semidistributive, it suffices to prove that a nontrivial module over a unital ring cannot have a term operation which is a 2-semilattice operation. Put more precisely, it suffices to prove that a nontrivial left module over a unital ring can not have a commutative, idempotent, binary term operation t satisfying $t(x, t(x, y)) \approx t(x, y)$. Recall from Chapter 2 that a nontrivial module is a module with more than one element.

Proposition 3.2.3. *The variety $\mathcal{S}$ is congruence meet semidistributive.*

Proof. Fix a unital ring $\mathbf{R}$ and a left $\mathbf{R}$-module $\mathbf{M}$. From a universal algebraic point of view, the operations of $\mathbf{M}$ are a binary operation $+$, a unary operation $-$, a constant 0, and for each $r \in R$, a unary operation $t_r : M \to M$ given by $t_r(m) = rm$. The only operations that can be constructed by composing these are of the form $t(x_1, \ldots, x_n) = r_1x_1 + \cdots r_nx_n$ for some $r_1, \ldots, r_n \in R$ and variables $x_1, \ldots, x_n$. This means that all binary term operations of $\mathbf{M}$ are of the form $t(x, y) = ax + by$ for some $a, b \in R$. We now suppose $\mathbf{M}$ has a binary term operation $t(x, y)$ given by $t(x, y) = ax + by$ for some $a, b \in R$ which satisfies the conditions outlined before the statement of the Proposition. In other words, there are $a, b \in R$ such that for all $m, n \in M$ the following hold:

1. $am + bm = m$
2. $am + bn = an + bm$
3. $am + b(am + bn) = am + bn$.

Notice that condition 3 is equivalent to $b(am + bn) = bn$ since am can be cancelled from both sides of the equation. Setting $n = 0$ and using (2), we get $am = bm$ for all m.

Applying this to the simplified version of equation 3 gives $a(am + bn) = bn$, and again setting $n = 0$ gives $a^2 m = 0$ for all $m \in M$. Applying $am = bm$ to equation 1, we get $am + am = m$ and multiplying both sides by a gives $a^2m + a^2m = am$. The left side is 0, so $am = 0$ for all $m \in M$. By what we showed earlier, this means $bm = 0$ for all $m \in M$ as well. Applying this to equation 1, we get $m = 0$ for all $m \in M$. Therefore, $\mathbf{M}$ is trivial and $\mathcal{S}$ is meet semidistributive. $\square$

3.3 Absorption

We will make use of the so-called Absorption Theorem of Barto and Kozik, which appeared as Theorem 2.3 in [2]. To apply the Absorption Theorem, we usually use Theorem 3.3.2. Before stating Lemma 3.3.1 which will be used in the proof of Theorem 3.3.2, we point the reader to the discussion on identities in Section 2 for an explanation of what "syntactically mentions" means.

Lemma 3.3.1. *Suppose* $\mathbf{A}$ *is a finite* 2*-semilattice,* $a, b \in A$*,* $a \longrightarrow b$*, and* $t(x_1, \cdots, x_n)$ *is a term which syntactically mentions* x_1*. Then* $t^{\mathbf{A}}(b, a, \ldots, a) = b$*.*

Proof. Suppose t is a counterexample which is minimal with respect to the number of times $*$ occurs in t. If $*$ does not occur in t, then t must be x_1, so $t(b, a, \ldots, a) = b$. Since t is a counterexample, this can't happen, so there are terms r and s so that $t(x_1, x_1, \ldots, x_n) = r(x_1, x_2, \ldots, x_n) * s(x_1, x_2, \ldots, x_n)$. Since $\{a, b\}$ is a subuniverse of $\mathbf{A}$ by Lemma 3.1.5 (5), each of $r(b, a, \ldots, a)$ and $s(b, a, \ldots, a)$ is either a or b. We are also assuming $t(b, a, \ldots, a) = a$, so $r(b, a, \ldots, a) = s(b, a, \ldots, a) = a$ because $ab = ba = bb = b$. Furthermore, one of r and s must mention x_1 syntactically since t does. This contradicts the minimality of t. $\square$

Theorem 3.3.2. *Let* $\mathbf{A} \in \mathcal{S}$ *be finite and strongly connected. Then* $\mathbf{A}$ *has no proper absorbing subuniverse.*

Proof. First, we note that if $B \lhd \mathbf{A}$ is proper, then there is some term operation t witnessing the absorption that depends on all of its variables. Indeed, if u is an n-ary term operation witnessing $B \lhd \mathbf{A}$ that depends on variables $x_1, \ldots, x_k$ with $k \leq n$, then $t(x_1, \ldots, x_k)$ defined by $u(x_1, \ldots, x_{k-1}, x_k, x_k, \ldots, x_k)$ depends on all of its variables and witnesses that $B \lhd \mathbf{A}$. Suppose B is a proper absorbing subuniverse of $\mathbf{A}$ with respect to some n-ary term operation t which depends on all of its variables. Since B is proper and $(A, \longrightarrow)$ is strongly connected, there are $b \in B$ and $c \in A - B$ with $b \longrightarrow c$. Since t depends on x_1, any term which defines t must mention x_1 syntactically. By Lemma 3.3.1, $t(c, b, \ldots, b) = c$, which contradicts $B \lhd \mathbf{A}$. □

Using what has already been done in this section, we will apply the Absorption Theorem to 2-semilattices in the next section. Before stating the Absorption Theorem, we collect some definitions.

Definition 3.3.3. Let $\mathbf{A}$ and $\mathbf{B}$ be finite algebras with $R \leq_{\text{sd}} \mathbf{A} \times \mathbf{B}$. Let $\mathbb{G}_R$ be the bipartite simple graph with vertex set $A \cup B$ and an edge connecting a and b exactly when $(a, b) \in R$. We say that R is *linked* if and only if $\mathbb{G}_R$ is connected.

We now define Taylor operations. These operations are named after Walter Taylor, who first defined them in [52]. The defining identities of a Taylor operation cannot be satisfied by any projection in a nontrivial algebra. Hence, an algebra having a Taylor term operation in a sense indicates a "non-triviality" of its term operations.

Definition 3.3.4. Let A be a set and $n \geq 2$. A *Taylor operation* of arity n is an idempotent operation $t : A^n \to A$ which, for each $i \leq n$, satisfies an identity $t(x_1^i, x_2^i, \ldots, x_n^i) \approx t(y_1^i, y_2^i, \ldots, y_n^i)$ where $x_k^i, y_k^i \in \{x, y\}$ for $1 \leq k \leq n$ and $x_i^i \neq y_i^i$.

Definition 3.3.5.

1. An idempotent algebra $\mathbf{A}$ is said to be *Taylor* if it has a Taylor term operation.

2. A variety $\mathcal{V}$ is said to be *Taylor* if there is a term t in its type and a set of identities so that $t^{\mathbf{A}}$ is a Taylor operation with respect to those identities for each algebra $\mathbf{A} \in \mathcal{V}$.

For example, the variety of groups is Taylor since the operation given by $t(x, y, z) = xy^{-1}z$ is a Taylor operation. To see this, note that the variety of groups satisfies $t(x, y, y) \approx t(y, y, x) \approx t(x, x, x) \approx x$. This shows that in the variety of groups t is idempotent, $t(x, y, y) \approx t(y, y, x)$ is an identity mentioning two variables with the variables unequal in the first and third positions, and $t(x, y, y) \approx t(x, x, x)$ is an identity mentioning two variables with the second variables unequal.

Proposition 3.3.6. *The variety $\mathcal{S}$ is Taylor.*

Proof. The basic operation $*$ is idempotent and satisfies $x * y \approx y * x$ which is an identity where the first variables are different, and the second variables are different. Indeed, any binary, idempotent, commutative operation is a Taylor operation. $\square$

Now we state the Absorption Theorem of Barto and Kozik, which appears as Theorem 2.3 in [2]. We remind the reader that an algebra is absorption free if it has no proper absorbing subuniverse.

Theorem 3.3.7 (Absorption Theorem)**.** *Suppose $\mathbf{A}$ and $\mathbf{B}$ are finite, absorption free, and in an idempotent variety with a Taylor term. If $R \leq_{\mathrm{sd}} \mathbf{A} \times \mathbf{B}$ is linked, then $R = A \times B$.*

3.4 Subdirect Products of Strongly Connected 2-Semilattices

We will now study subdirect products of 2-semilattices. All results in this section were known to Bulatov in [15]. Bulatov's proofs relied on careful analysis of the nature of strongly connected 2-semilattices. We present shorter proofs which rely on the Absorption Theorem, which was unknown at the time Bulatov wrote [15]. Recall from Chapter 2 that an algebra $\mathbf{A}$ is simple if and only if it has exactly two congruences: $0_A = \{(a,a) : a \in A\}$ and $1_A = A \times A$.

Lemma 3.4.1. *Let $\mathbf{A}, \mathbf{B} \in \mathcal{S}$ be finite and strongly connected with $\mathbf{B}$ simple. If $R \leq_{\text{sd}} \mathbf{A} \times \mathbf{B}$ then either $R = A \times B$ or there is a surjective homomorphism $\varphi : \mathbf{A} \to \mathbf{B}$ such that $R = \{(a, \varphi(a)) : a \in A\}$. In other words, R is the graph of φ.*

Proof. Assume $R \leq_{\text{sd}} \mathbf{A} \times \mathbf{B}$. If R is the graph of a function, it is necessarily the graph of a homomorphism. Therefore, if we assume R is not the graph of a homomorphism, subdirectness of R guarantees that there is some $a \in A$ and distinct $b_1, b_2 \in B$ with both (a, b_1) and $(a, b_2) \in R$. The relation

$$\tau = \{(c,d) \in B^2 : \text{ there is } a \in A \text{ such that } (a,c),(a,d) \in R\}$$

is a symmetric and reflexive subuniverse of $\mathbf{B}^2$, so its transitive closure, α, is a congruence on $\mathbf{B}$. This follows from Proposition 3.0.2 (1). We also have that $(b_1, b_2) \in \tau \subseteq \alpha$, so $\alpha = B \times B$ because $b_1 \neq b_2$ and $\mathbf{B}$ is simple. This shows that R is linked. We also have that $\mathbf{A}$ and $\mathbf{B}$ are absorption free by Theorem 3.3.2, and that $\mathcal{S}$ is Taylor by Proposition 3.3.6. The conditions of Theorem 3.3.7 are satisfied, so $R = A \times B$. $\square$

Lemma 3.4.2. *Let $\mathbf{A}, \mathbf{B} \in \mathcal{S}$ be finite, strongly connected, and simple with $R \leq_{\text{sd}} \mathbf{A} \times \mathbf{B}$. Either $R = A \times B$ or R is the graph of a bijection.*

Proof. Suppose $R \leq_{\text{sd}} \mathbf{A} \times \mathbf{B}$ and that $R \neq A \times B$. Since $\mathbf{B}$ is simple, Lemma 3.4.1 implies R is the graph of a surjective homomorphism from $\mathbf{A}$ to $\mathbf{B}$. Using that $\mathbf{A}$ is simple and applying Lemma 3.4.1 again, we get that this homomorphism is injective, as well. Therefore, R is the graph of a bijection from A to B. $\square$

Recall that for a 2-semilattice $\mathbf{D}$, the smallest strongly connected component of $(D, \longrightarrow)$ is denoted by D'.

Lemma 3.4.3. *Suppose* $\mathbf{A}$ *and* $\mathbf{B}$ *are finite* 2*-semilattices.*

1. *If* $\mathbf{R} \leq_{\text{sd}} \mathbf{A} \times \mathbf{B}$*, then* $\mathbf{R}' \leq_{\text{sd}} \mathbf{A}' \times \mathbf{B}'$.
2. $(\mathbf{A} \times \mathbf{B})' = \mathbf{A}' \times \mathbf{B}'$.

Proof.

1. First, suppose $(a, b) \in R'$ and choose any $u \in A$. Since $\mathbf{R} \leq_{\text{sd}} \mathbf{A} \times \mathbf{B}$, there is $v \in B$ with $(u, v) \in R$. By Lemma 3.1.5 (3), there is a directed walk from (u, v) to (a, b) in R. Restricting to the first coordinate gives a directed walk from u to a in A. Since u was chosen arbitrarily in A, Lemma 3.1.5 (3) implies $a \in A'$. A similar argument shows that $b \in B'$. We have that $R' \leq \mathbf{A}' \times \mathbf{B}'$. Now choose $a \in A'$ and $(u, v) \in R'$. Using Lemma 3.1.5 (3), find a directed walk,

$$u = u_1 \longrightarrow u_2 \longrightarrow \cdots \longrightarrow u_n = a,$$

from u to a. Using that $R \leq_{\text{sd}} \mathbf{A} \times \mathbf{B}$, extend $u_1, \ldots, u_n$ to pairs (u_i, v_i) in R, with $v_1 = v$. Consider the element

$$(x, y) = (u_1, v_1)(u_2, v_2) \cdots (u_{n-1}, v_{n-1})(u_n, v_n)$$

of R and recall that association takes place on the left. By construction, we have $x = u_n = a$. Each (u_i, v_i) was chosen to be in R and $(u_1, v_1) \in R'$, so by repeated application of Lemma 3.1.5 (4), $(a, y) \in R'$. This proves that $A' \subseteq \mathrm{pr}_1(R')$. Similarly, we can prove that $B' \subseteq \mathrm{pr}_2(R')$. The opposite inclusions follow from the first part of the argument, so $\mathrm{pr}_1(R') = A'$ and $\mathrm{pr}_2(R') = B'$. By definition, $R' \leq_{\mathrm{sd}} \mathbf{A}' \times \mathbf{B}'$.

2. Since $\mathbf{A} \times \mathbf{B} \leq_{\mathrm{sd}} \mathbf{A} \times \mathbf{B}$, we have that $(\mathbf{A} \times \mathbf{B})' \leq_{\mathrm{sd}} \mathbf{A}' \times \mathbf{B}'$ by (1), so it is enough to show that $A' \times B' \subseteq (A \times B)'$. To see this, choose $(a, b) \in A' \times B'$ and $(c, d) \in A \times B$. By Lemma 3.1.5 (3), it is enough to show there is a directed walk from (c, d) to (a, b) in $\mathbf{A} \times \mathbf{B}$. Since $a \in A'$, by Lemma 3.1.5 (3), there is a directed walk, $c = u_1 \longrightarrow \cdots \longrightarrow u_n = a$, so $(c, d) = (u_1, d) \longrightarrow \cdots \longrightarrow (u_n, d) = (a, d)$ is a directed walk in $\mathbf{A} \times \mathbf{B}$ from (c, d) to (a, d). By similar reasoning, there is a directed walk from (a, d) to (a, b) in $\mathbf{A} \times \mathbf{B}$. Concatenating these walks completes the proof.

□

Lemma 3.4.4 will be used in several places, including the proof of Lemma 3.4.5 which follows it. It is essentially the same as Lemma 3.7 from [15], but the proof is different.

Lemma 3.4.4. *Suppose $n > 1$ and $\mathbf{A}_1, \ldots, \mathbf{A}_n$ are 2-semilattices which are simple and strongly connected. Let $\mathbf{T} \leq_{\mathrm{sd}} \mathbf{A}_1 \times \cdots \times \mathbf{A}_n$ and $\mathrm{pr}_{i,j}(T) = A_i \times A_j$ for all $i \neq j$. Then $T = A_1 \times \cdots \times A_n$.*

Before proceeding with the proof, we remind the reader that for an algebra $\mathbf{A}$ and a congruence α on $\mathbf{A}$, the algebra $\mathbf{A}/\alpha$ is simple if and only if α is maximal in $\mathrm{Con}(\mathbf{A})$. This follows from the Correspondence Theorem, which we stated as Theorem 2.0.6.

Proof. Let $\mathbf{T} \leq \mathbf{A}_1 \times \cdots \times \mathbf{A}_n$ be a counterexample which is minimal with respect to n. Note that $n > 2$ since the result is immediate if $n = 2$. We will denote by $\mathbf{B}_i$ the algebra

$\mathbf{A}_1 \times \cdots \times \mathbf{A}_{i-1} \times \mathbf{A}_{i+1} \times \cdots \times \mathbf{A}_n$. For each i, we can identify $\mathbf{T}$ as a subdirect product of $\mathrm{pr}_{1,2,\ldots,i-1,i+1,\ldots,n}(\mathbf{T})$ and $\mathbf{A}_i$. The algebra $\mathrm{pr}_{1,2,\ldots,i-1,i+1,\ldots,n}(\mathbf{T}) \leq_{\mathrm{sd}} \mathbf{B}_i$ and satisfies the conditions of the lemma, so by the minimality of $\mathbf{T}$, it is equal to $\mathbf{B}_i$. By Lemma 3.1.8 (2), $\mathbf{B}_i$ is strongly connected, so $\mathrm{pr}_{1,\ldots,i-1,i+1,\ldots,n}(\mathbf{T})$ is strongly connected. Since $\mathbf{A}_i$ is simple and $\mathbf{T}$ is a counterexample to the lemma, we can apply Lemma 3.4.1 to get that, for each i, T is the graph of a surjective homomorphism from $\mathbf{B}_i$ to $\mathbf{A}_i$. Let α_j be the kernel of the j^{th} projection from $\mathbf{T}$ to $\mathbf{A}_j$. That is, α_j is the congruence on $\mathbf{T}$ which identifies elements of T which agree in the j^{th} position. For now, fix i and suppose $(\mathbf{u}, \mathbf{v}) \in \bigwedge_{j \neq i} \alpha_j$. This means $\mathbf{u}$ and $\mathbf{v}$ agree everywhere except possibly the i^{th} position. Since $\mathbf{T}$ is the graph of a function from $\mathbf{B}_i$ to $\mathbf{A}_i$, this means $\mathbf{u}$ and $\mathbf{v}$ agree in the i^{th} position, as well. In other words, $\bigwedge_{j \neq i} \alpha_j \leq \alpha_i$. This implies

$$\left(\bigwedge_{j \neq i} \alpha_j\right) \leq \left(\bigwedge_{j \neq i} \alpha_j\right) \wedge \alpha_i.$$

Since α_j is the congruence which collapses elements which agree in the j^{th} coordinate, $\bigwedge_{j=1}^{n} \alpha_j = 0_T$. This simplifies to

$$\bigwedge_{j \neq i} \alpha_j = 0_T.$$

This shows that $\bigwedge_{j \neq i} \alpha_j = 0_A$ for each i.

Suppose $|A_i| = 1$ for some i. Since $\mathbf{T} \leq_{\mathrm{sd}} \prod_{i=1}^{n} \mathbf{A}_i$, in this case, $\mathbf{T} = \prod_{i=1}^{n} \mathbf{A}_i$ if and only if $\mathrm{pr}_{1,2,\ldots,i-1,i+1,\ldots,n}(\mathbf{T}) = \mathbf{B}_i$. Because we are assuming $T \neq \prod_{i=1}^{n} A_i$, we have $|A_i| > 1$ for each i. If $\alpha_i = \alpha_j$ for some $i \neq j$, then $\mathrm{pr}_{i,j}(T)$ is the graph of a bijection. We are also assuming $\mathrm{pr}_{i,j}(T) = A_i \times A_j$, and the only way for $A_i \times A_j$ to be the graph of a bijection is if $|A_i| = |A_j| = 1$. We already showed that this is not the case, so we conclude that the α_i are distinct.

For a subset $X \subseteq \{1, 2, \ldots, n\}$, let $\theta(X) = \bigwedge_{i \in X} \alpha_i$. Suppose $|X| = |Y| = |X \cap Y| + 1$ with $\theta(X) = \theta(Y)$. If we set $U = X \cap Y$, there must be $i \neq j$ so that $X = U \cup \{i\}$

and $Y = U \cup \{j\}$, which means $\theta(X) = \theta(U) \wedge \alpha_i$ and $\theta(Y) = \theta(U) \wedge \alpha_j$. By meet semidistributivity, $\theta(X) = \theta(Y) = \theta(U) \wedge (\alpha_i \vee \alpha_j)$. By subdirectness, $\mathbf{T}/\alpha_i \cong \mathbf{A}_i$ and $\mathbf{T}/\alpha_j \cong \mathbf{A}_j$, so α_i and α_j are maximal because $\mathbf{A}_i$ and $\mathbf{A}_j$ are simple. By the previous paragraph, $\alpha_i \neq \alpha_j$, so $\alpha_i \vee \alpha_j = 1_T$ and it follows that $\theta(X) = \theta(Y) = \theta(U)$. Applying this, we get that if $|X| = |Y| = |X \cap Y| + 1$ and $\theta(X) = \theta(Y) = 0_T$, then $\theta(X \cap Y) = 0_T$. The first part of the proof showed that $\theta(\{1, 2, \ldots, n\} - \{i\}) = 0_T$ for all i, so by repeatedly applying the fact in the previous sentence, we get that $\theta(X) = 0_T$ whenever $|X| \geq 1$, which means $\alpha_1 = \cdots = \alpha_n = 0_A$. This forces $\mathrm{pr}_{i,j}(T)$ to be the graph of a bijection between $\mathbf{A}_i$ and $\mathbf{A}_j$ for each i, j. As argued earlier, this means $|A_i| = |A_j| = 1$. It was established in the previous paragraph that this is not the case. We have a contradiction, so the proof is complete.

□

The next lemma is a special case of Lemma 3.8 from [15]. The proof given is an adaptation of Bulatov's proof. The only place that Lemma 3.4.5 will be used is in the proof of Lemma 6.3.4, but we include it here because it fits well in this section.

Lemma 3.4.5 (Lemma 3.8 from [15]). *Let $\mathbf{A}_1, \mathbf{A}_2, \mathbf{A}_3 \in \mathcal{S}$ be finite and strongly connected and $T \leq_{\mathrm{sd}} \mathbf{A}_1 \times \mathbf{A}_2 \times \mathbf{A}_3$ satisfy the following:*

1. *$\mathbf{A}_3$ is simple,*
2. *$\mathrm{pr}_{1,2}(T)$ is strongly connected,*
3. *$\mathrm{pr}_{i,3}(T) = A_i \times A_3$ for $i = 1, 2$.*

Then $T = \mathrm{pr}_{1,2}(T) \times A_3$.

Proof. For now, suppose $\mathbf{A}_1$ and $\mathbf{A}_2$ are both simple. Since $T \leq_{\mathrm{sd}} \mathbf{A}_1 \times \mathbf{A}_2 \times \mathbf{A}_3$, we get that $\mathrm{pr}_{1,2}(T) \leq_{\mathrm{sd}} \mathbf{A}_1 \times \mathbf{A}_2$. By Lemma 3.4.2, either $\mathrm{pr}_{1,2}(T) = A_1 \times A_2$, or it is the graph of a bijection. If $\mathrm{pr}_{1,2}(T) = A_1 \times A_2$, then $T = \mathbf{A}_1 \times \mathbf{A}_2 \times \mathbf{A}_3$ by Lemma 3.4.4 and the result holds. Now suppose $\mathrm{pr}_{1,2}(T)$ is the graph of a bijection. Choose $(a,b) \in \mathrm{pr}_{1,2}(T)$ and $c \in A_3$. Since $\mathrm{pr}_{1,3}(T) = A_1 \times A_3$, there is a $b' \in A_2$ with $(a,b',c) \in T$, but this means $(a,b') \in \mathrm{pr}_{1,2}(T)$, so $b' = b$ and $(a,b,c) \in T$. This shows that $T = \mathrm{pr}_{1,2}(T) \times A_3$.

We now assume $\mathbf{A}_1, \mathbf{A}_2, \mathbf{A}_3$ and T form a counterexample minimal with respect to $|A_1| + |A_2| + |A_3|$. By the first paragraph, a counterexample can not have both $\mathbf{A}_1$ and $\mathbf{A}_2$ simple so we assume without loss of generality that $\mathbf{A}_1$ is not simple. Let α be a maximal congruence of $\mathbf{A}_1$ and define $T_\alpha \leq_{\mathrm{sd}} \mathbf{A}_1/\alpha \times \mathbf{A}_2 \times \mathbf{A}_3$ by

$$T_\alpha = \{(a_1/\alpha, a_2, a_3) : (a_1, a_2, a_3) \in T\}.$$

The algebra $\mathbf{T}_\alpha$ inherits the conditions of the lemma, but since $\mathbf{A}_1$ is not simple, $\alpha \neq 0_A$, and $|A_1/\alpha| < |A_1|$. By the minimality of $(T, \mathbf{A}_1, \mathbf{A}_2, \mathbf{A}_3)$, $T_\alpha = \mathrm{pr}_{1,2}(T_\alpha) \times A_3$.

Choose an α-class C, let $S = \{(a,b,c) \in T : a \in C\}$ and define $D = \mathrm{pr}_2(S)$. Since $\mathrm{pr}_{1,3}(T) = A_1 \times A_3$, for any $c \in A_3$ and $a \in C$, there is $b \in A_2$ such that $(a,b,c) \in T$, but $a \in C$ which means $(a,b,c) \in S$. This shows $\mathrm{pr}_3(S) = A_3$ and hence $S \leq_{\mathrm{sd}} \mathrm{pr}_{1,2}(S) \times \mathbf{A}_3$. By Lemma 3.4.3 (1), $S' \leq_{\mathrm{sd}} \mathrm{pr}_{1,2}(S)' \times A_3$ since $A_3' = A_3$, which implies $\mathrm{pr}_{1,2}(S') = \mathrm{pr}_{1,2}(S)'$. Because of how C and D are defined, we have $\mathrm{pr}_{1,2}(S) \leq_{\mathrm{sd}} \mathbf{C} \times \mathbf{D}$. Applying Lemma 3.4.3 gives $\mathrm{pr}_{1,2}(S') = \mathrm{pr}_{1,2}(S)' \leq_{\mathrm{sd}} \mathbf{C}' \times \mathbf{D}'$. This implies $S' \leq_{\mathrm{sd}} \mathbf{C}' \times \mathbf{D}' \times \mathbf{A}_3$, and since $\mathrm{pr}_{1,2}(S') = \mathrm{pr}_{1,2}(S)'$ and the latter is strongly connected by definition, we have that $\mathrm{pr}_{1,2}(S')$ is strongly connected. If we can show that $\mathrm{pr}_{1,3}(S') = C' \times A_3$ and $\mathrm{pr}_{2,3}(S') = D' \times A_3$, we will have that $\mathbf{S}' \leq_{\mathrm{sd}} \mathbf{C}' \times \mathbf{D}' \times \mathbf{A}_3$ satisfies the conditions of the lemma.

To see that $\mathrm{pr}_{1,3}(S') = C' \times A_3$, suppose $a \in C'$ and $c \in \mathbf{A}_3$. Since $S' \leq_{\mathrm{sd}} \mathbf{C}' \times \mathbf{D}' \times \mathbf{A}_3$, there is a triple $(a, b, c^\times) \in S'$. Since $\mathrm{pr}_{1,3}(T) = A_1 \times A_3$, there is $(a, b^\times, c) \in T$ and in fact, since $a \in C' \subseteq C$, $(a, b^\times, c) \in S$. Since $\mathbf{A}_3$ is strongly connected, there is a walk, $c^\times \longrightarrow c_1 \longrightarrow \dots \longrightarrow c_n \longrightarrow c$, and again since $\mathrm{pr}_{1,3}(T) = A_1 \times A_3$, there are $b_1, \dots, b_n \in A_2$ with $(a, b, c^\times), (a, b_i, c_i), (a, b^\times, c) \in T$ for all $1 \leq i \leq n$. Since $a \in C$, each of these triples is in S. Multiplying, we get, $(aa \cdots a, bb_1b_2 \cdots b_nb^\times, c^\times c_1c_2 \cdots c_nc) = (a, \bar{b}, c)$ for some $\bar{b} \in A_2$. Since $(a, b, c^\times) \in S'$ and each triple is in S, we get that $(a, \bar{b}, c) \in S'$, as well. Therefore, $\mathrm{pr}_{1,3}(S') = C' \times A_3$. To show that $\mathrm{pr}_{2,3}(S') = D' \times A_3$, we start by choosing $b \in D$ and $c \in A_3$ arbitrarily. By definition of D, $(C, b) \in \mathrm{pr}_{1,2}(T_\alpha)$. From earlier, we have that $T_\alpha = \mathrm{pr}_{1,2}(T_\alpha) \times A_3$, so $(C, b, c) \in T_\alpha$, which means there is some $a \in C$ such that $(a, b, c) \in T$, and since $a \in C$, we actually have $(a, b, c) \in S$. This shows that $\mathrm{pr}_{2,3}(S) = D \times A_3$. By Lemma 3.4.3 (1), $S' \leq_{\mathrm{sd}} \mathbf{C}' \times \mathrm{pr}_{2,3}(S)'$, which shows that $\mathrm{pr}_{2,3}(S') = \mathrm{pr}_{2,3}(S)'$. W have just shown that $\mathrm{pr}_{2,3}(S) = D \times A_3$, so

$$\mathrm{pr}_{2,3}(S') = \mathrm{pr}_{2,3}(S)' = (D \times A_3)' = D' \times A_3$$

by Lemma 3.4.3 (2).

Since $\mathbf{T}$ was a minimal counterexample to the lemma with respect to $|A_1| + |A_2| + |A_3|$ and α is a proper congruence on $\mathbf{A}_1$, we have $|C'| \leq |C| < |A_1|$, so the conclusion of the lemma applies to $S' \leq_{\mathrm{sd}} \mathbf{C}' \times \mathbf{D}' \times \mathbf{A}_3$, which means $S' = \mathrm{pr}_{1,2}(S') \times A_3$. If A_3 were a singleton, we would already have $T = \mathrm{pr}_{1,2}(T) \times A_3$, which we are assuming is not the case. Therefore, assume there are distinct $c_1, c_2 \in A_3$, and choose any $(a, b) \in \mathrm{pr}_{1,2}(S')$. Then $(a, b, c_1), (a, b, c_2) \in S' \subseteq T$. This means $T \leq_{\mathrm{sd}} \mathrm{pr}_{1,2}(T) \times \mathbf{A}_3$ is not the graph of a surjective homomorphism, so by Lemma 3.4.1, $T = \mathrm{pr}_{1,2}(T) \times \mathbf{A}_3$.

□

Chapter 4

Congruence Lattices of Finite 2-Semilattices

In this chapter, we explore some of the properties of congruences of 2-semilattices. Section 4.1 is devoted to Theorem 4.1.1 and its proof. Theorem 4.1.1 is about the structure of minimal congruences in 2-semilattices. The rest of the chapter is devoted to the exploration of how "Tame Congruence Theory", developed by Hobby and McKenzie in [29], fits into the context of 2-semilattices. Tame Congruence Theory is a broad structure theory for locally finite varieties which is based on an understanding of minimal congruences of finite algebras. We give a brief introduction to Tame Congruence Theory in the beginning of Section 4.2. Sections 4.2 and 4.3 apply Theorem 4.1.1 and Tame Congruence Theory to 2-semilattices.

4.1 Connectivity and Minimal Congruences

For a 2-semilattice $\mathbf{A}$, recall from Definition 3.1.3 that $\overset{\mathbf{A}}{\sim}$ is the equivalence relation on A whose equivalence classes are the strongly connected components of the digraph $(A, \longrightarrow)$. In this section, when we say "component", we mean "strongly connected component". As well, we will write $\sim$ rather than $\overset{\mathbf{A}}{\sim}$ whenever the algebra $\mathbf{A}$ is clear from context. The goal of this section is to prove the following theorem:

Theorem 4.1.1. *Let θ be a congruence of a finite 2-semilattice $\mathbf{A}$ which is minimal with respect to the order in $\mathrm{Con}(\mathbf{A})$. Exactly one of the following two conditions holds:*

1. *There are two distinct components, X and Y of $\mathbf{A}$ and a function $f : X \to Y$ so that*

$$\theta = \{(u, v) \in (X \cup Y)^2 : f(u) = f(v) \text{ or } f(u) = v \text{ or } f(v) = u\} \cup 0_A,$$

or

2. *There is a component X of $(A, \longrightarrow)$ such that $\theta \subseteq X^2 \cup 0_A$.*

This structural theorem has several interesting consequences. For example, we will show in Theorem 4.3.2 that finite simple 2-semilattices always have tame congruence theoretic type **3** or **5** (see Section 4.2). Before moving on, we give a definition and lemma about a type of congruence that always occurs in a 2-semilattice. Recall from Definition 3.1.3 (2) and Lemma 3.1.4 that $\succeq$ is a partial order on components where $X \succeq Y$ if and only if there is a directed walk from x to y for all $x \in X$ and $y \in Y$.

Definition 4.1.2. Let $\mathbf{A}$ be a finite 2-semilattice and suppose X is a component of $\mathbf{A}$. Define $\psi_X \subseteq A^2$ by

$$\psi_X = \{(a, b) : X \succeq a/{\sim} \text{ and } X \succeq b/{\sim}\} \cup 0_A.$$

Lemma 4.1.3. *Let* $\mathbf{A}$ *be a finite* 2*-semilattice and suppose* X *is a component of* $\mathbf{A}$*. The relation* ψ_X *from Definition 4.1.2 is a congruence on* $\mathbf{A}$*.*

Proof. Because $a \longrightarrow ab$ for any $a, b \in A$, the set $B = \bigcup_{X \succeq Y} Y$ is an absorbing subuniverse of $\mathbf{A}$ with respect to the operation $*$. Since $\psi_X = B^2 \cup 0_A$, it follows that ψ is a congruence. $\square$

Definition 4.1.4. Let $\mathbf{A}$ be a 2-semilattice, $\theta \in \mathrm{Con}(\mathbf{A})$, and B be a component. We say B is θ*-nontrivial* if there is some $b \in B$ such that $b/\theta \neq \{b\}$. Naturally, we say that B is θ*-trivial* otherwise.

Note that the condition $\theta \cap B^2 \nsubseteq 0_A$ is not implied by θ-nontriviality. It is possible that elements of B witnessing θ-nontriviality are only θ-related to elements not in B.

Suppose $\mathbf{A}$ is a finite 2-semilattice and $\theta \in \mathrm{Con}(\mathbf{A})$ is not 0_A. Then there must be some θ-nontrivial component. The restriction of $\succeq$ to θ-nontrivial components is therefore a non-empty, finite partial order, so it must have a minimal element. Such a component B has the property that if $b \in B$, $b \longrightarrow a$, and $a/\theta \neq \{a\}$, then $a \in B$. As suggested by Theorem 4.1.1, if θ is a minimal congruence, there is a unique minimal θ-nontrivial component in this sense and there are at most two θ-nontrivial components in total. This will be shown along the way to proving Theorem 4.1.1.

Lemma 4.1.5. *Let* $\mathbf{A}$ *be a finite* 2*-semilattice,* $\theta \in \mathrm{Con}(\mathbf{A})$ *be minimal, and* B *be a minimal* θ*-nontrivial component. If* $\theta \cap B^2 \nsubseteq 0_A$*, then* $\theta \subseteq B^2 \cup 0_A$*.*

Proof. Choose $a, b \in B$ such that $a \neq b$ and $(a, b) \in \theta$. Since θ is minimal, $\theta = \mathrm{Cg}_{\mathbf{A}}(a, b)$. We also have $(a, b) \in \psi_B$ which is a congruence by Lemma 4.1.3. This means $\theta \leq \psi_B$, so every θ-nontrivial component C is such that $B \succeq C$. By the minimality of B, every component C with $B \succ C$ is θ-trivial. Therefore, if $a/\theta \neq \{a\}$, then $a \in B$. This is the same as $\theta \subseteq B^2 \cup 0_A$. $\square$

Lemma 4.1.5 shows that Theorem 4.1.1 (2) holds when $\theta \cap B^2 \nsubseteq 0_A$. To finish the proof of Theorem 4.1.1, we will show that Theorem 4.1.1 (1) holds when $\theta \cap B^2 \subseteq 0_A$. This is the content of Lemma 4.1.6.

Lemma 4.1.6. *Suppose* $\mathbf{A}$ *is a finite* 2*-semilattice,* $\theta \in \operatorname{Con}(\mathbf{A})$ *is minimal, and* B *is a minimal* θ*-nontrivial component. Further suppose* $\theta \cap B^2 \subseteq 0_A$*. There is a component* $C \neq B$ *and a function* $f : C \to B$ *so that* $c \xrightarrow{\mathbf{A}} f(c)$ *for each* $c \in C$ *and*

$$\theta = \{(u, v) \in (X \cup Y)^2 : f(u) = f(v) \text{ or } f(u) = v \text{ or } f(v) = u\} \cup 0_A.$$

A verbal description of θ is that the nontrivial θ-blocks are precisely $\{f(a)\} \cup f^{-1}(f(a))$ for each $a \in C$.

Proof. Since $\theta \cap B^2 \subseteq 0_A$ and B is θ-nontrivial, there are $a \in A - B$ and $b \in B$ with $(a, b) \in \theta$. We now let $\mathcal{X}$ be the set of components C satisfying

1. $C \neq B$
2. There are $c \in C$ and $b \in B$ with $(c, b) \in \theta$.

By the first sentence, $a/{\sim} \in \mathcal{X}$, so $\mathcal{X}$ is nonempty. Thus, we can choose a component $C \in \mathcal{X}$ which is minimal with respect to $\succeq$. Note that conditions (1) and (2) guarantee that C is θ-nontrivial. Suppose $c_1, c_2 \in C$ and $b' \in B$ with $(c_1, b') \in \theta$ and $c_1 \longrightarrow c_2$. Then $(c_1c_2, b'c_2) = (c_2, b'c_2) \in \theta$. If $c_2 = b'c_2$ then $b' \longrightarrow c_2$, which means $B \longrightarrow C$. This is impossible by the minimality of B. Therefore, $c_2 \neq b'c_2$ and $b'c_2/{\sim}$ is θ-nontrivial. Since $b' \longrightarrow b'c_2$, we have $B \longrightarrow b'c_2/{\sim}$, so $B = b'c_2/{\sim}$ by the minimality of B. This means $b'c_2 \in B$. If we take $b = b'c_2$, we have shown that there is $b \in B$ with $c_2 \longrightarrow b$ and $(c_2, b) \in \theta$. Since C is strongly connected, this argument can be repeated to show that every $c \in C$ has this property. Furthermore, if $c \in C$ and there are $b_1, b_2 \in B$ with

$(c, b_i) \in \theta$ for each i, then $(b_1, b_2) \in \theta$ since θ is an equivalence relation. This means $b_1 = b_2$ since $\theta \cap B^2 \subseteq 0_A$. We have shown that $\theta \cap (C \times B)$ is the graph of a function $f : C \to B$ that satisfies $c \longrightarrow f(c)$ for each $c \in C$.

To complete the proof, we set

$$\sigma = \{(u, v) : f(u) = f(v) \text{ or } f(u) = v \text{ or } f(v) = u\} \cup 0_A$$

and show that $\sigma = \theta$. Because of how σ is defined, we know that $0_A \neq \sigma \subseteq \theta$. Therefore, to prove that $\sigma = \theta$, it is enough to show that σ is a congruence since θ is assumed to be a minimal congruence. The relation σ is defined in such a way that it is an equivalence relation, so by Proposition 3.0.2 (1), it is enough to show for any $(u, v) \in \sigma$ and $a \in A$ that $(au, av) \in \sigma$. We assume u, v, and a are as stated and note that $au = av$ implies $(au, av) \in \sigma$, so we additionally assume $au \neq av$. We will use that $(au, av) \in \theta$, which is true because $(u, v) \in \theta$. Suppose $f(u) = v$. In this case, $(u, v) \in C \times B$ and since $au \neq av$, $av/\!\sim$ is θ-nontrivial. This puts $av \in B$ because $v \longrightarrow av$ and B is a minimal θ-nontrivial component. Since $\theta \cap B^2 \subseteq 0_A$ and $au \neq av$, we have $au/\!\sim \in \mathcal{X}$, but $u \longrightarrow au$ and C is minimal in $\mathcal{X}$, so $au \in C$. It follows that $f(au) = av$ and $(au, av) \in \sigma$. The case when $f(v) = u$ is handled in the same way. Now assume $f(u) = f(v)$ and set $w = f(u) = f(v)$. By definition, we also have $(u, w), (w, v) \in \sigma \subseteq \theta$, so $(au, aw), (aw, av) \in \theta$ as well. Since $au \neq av$, we assume without loss of generality that $au \neq aw$. By the previous case, this means $au \in C$ and $aw \in B$ so $f(au) = aw$. If $aw = av$, we have $f(au) = av$ which puts $(au, av) \in \sigma$. Otherwise, the same reasoning leads to $f(av) = aw$, so $f(au) = f(av)$ and $(au, av) \in \sigma$. This completes the proof that σ is a congruence, which shows that $\sigma = \theta$ and θ has the claimed structure. □

4.2 Tame Congruence Theory in 2-Semilattices

In [29], Hobby and McKenzie developed what is now known as Tame Congruence Theory, which will hence-forth be abbreviated by TCT. In some sense, TCT is a way of relating local and global behaviour in a locally finite variety. We now give a brief overview of TCT. An in-depth and systematic development of the subject can be found in [29].

For an algebra $\mathbf{A}$, an n-ary *polynomial operation* p on $\mathbf{A}$ is an n-ary operation defined by $p(x_1, \ldots, x_n) = t(x_1, \ldots, x_n, a_1, \ldots, a_m)$ where t is an $m+n$-ary term operation, $m, n \geq 0$, and $a_1, \ldots, a_m \in A$. If we take $m = 0$, p is a term operation, so every term operation is a polynomial operation. For two congruences α and β, the pair $\langle \alpha, \beta \rangle$ is called a *congruence quotient* if $\alpha < \beta$. The quotient is called *prime* if $\alpha \prec \beta$. Recall from Definition 2.0.4 that $\alpha \prec \beta$ means $\alpha < \beta$ and there are no congruences strictly between α and β.

Given a finite algebra $\mathbf{A}$ and a prime congruence quotient $\langle \alpha, \beta \rangle$, the authors of [29] consider the range $p(A)$ of unary polynomials p satisfying $\{(p(a), p(b)) : (a, b) \in \beta\} \nsubseteq \alpha$. The set $\{(p(a), p(b)) : (a, b) \in \beta\}$ will be abbreviated by $p(\beta)$. Since A is finite, there are only finitely many possible ranges, so there must be such a polynomial p with $p(A) = U$ minimal with respect to inclusion. Such a U is called $\langle \alpha, \beta \rangle$-*minimal*. The authors then restrict their attention to the algebra $\mathbf{A}{\restriction_U}$ whose domain is U and whose basic operations are the restrictions to U of all polynomials of $\mathbf{A}$ under which U is closed. The relations $\alpha{\restriction_U} = \alpha \cap U^2$ and $\beta{\restriction_U} = \beta \cap U^2$ are congruences of this algebra, and $\langle \alpha{\restriction_U}, \beta{\restriction_U} \rangle$ is a prime congruence quotient of $\mathbf{A}{\restriction_U}$. The authors then choose a $\beta{\restriction_U}$-class N which is not an $\alpha{\restriction_U}$ class and with the same notion of restriction, consider the algebra $(\mathbf{A}{\restriction_U}){\restriction_N} / (\alpha{\restriction_N})$. This algebra is essentially one of the following:

type **1**	A set whose operations are a group of permutations on that set,
type **2**	A vector space over a finite field,
type **3**	A two element boolean algebra,
type **4**	A two element lattice,
type **5**	A two element semilattice.

Remarkably, the type of this algebra depends only on the congruence quotient $\langle\alpha,\beta\rangle$, and not on the choices of U and N which suggests that the algebra $(\mathbf{A}\restriction_U)\restriction_N /(\alpha\restriction_N)$ should carry information about $\langle\alpha,\beta\rangle$. Since the type depends only on $\langle\alpha,\beta\rangle$, there is a well defined function typ from prime congruence quotients of finite algebras to the set $\{\mathbf{1},\mathbf{2},\mathbf{3},\mathbf{4},\mathbf{5}\}$. The authors of [29] adapt the definition of typ to take a variety as input by defining

$$\operatorname{typ}(\mathcal{V}) = \{\operatorname{typ}(\langle\alpha,\beta\rangle) : \langle\alpha,\beta\rangle \text{ is a prime quotient of finite } \mathbf{A}\in\mathcal{V}\}.$$

The set $\operatorname{typ}(\mathcal{V})$ is called the *type set* of $\mathcal{V}$. The type set of a locally finite variety $\mathcal{V}$ holds information about identities which are satisfied by $\mathcal{V}$, as well as structural properties of the congruence lattices of its members. For example, it follows from results in [29] and [52] that a locally finite variety is Taylor (Definition 3.3.4) if and only if its type set does not include **1**. The theorem of most interest for us is Theorem 9.10 from [29] which says, among other things, that a locally finite variety is meet semidistributive if and only if its type set omits **1** and **2**.

By Proposition 3.2.3, the variety of 2-semilattices is meet semidistributive, so any locally finite variety of 2-semilattices is meet semidistributive. By Theorem 9.10 from [29], a locally finite variety of 2-semilattices can only have prime congruence quotients of types **3**, **4**, or **5**. In Example 4.3.5 we will see that each of these can occur. The remainder of this chapter is an exploration of the TCT types that occur in varieties of 2-semilattices. By Corollary 5.3 in [29], the type of a prime quotient $\langle\alpha,\beta\rangle$ of $\mathbf{A}$ is the same as the type of $\langle 0_{A/\alpha},\beta/\alpha\rangle$ in $\mathbf{A}/\alpha$. It follows from the Correspondence Theorem (Theorem 2.0.6) that $\langle\alpha,\beta\rangle$ is a prime congruence quotient in $\mathbf{A}$ if and only if $\langle 0_{A/\alpha},\beta/\alpha\rangle$ is prime in $\mathbf{A}/\alpha$. Hence, to understand the TCT types in a locally finite variety, it is sufficient to study prime quotients of the form $\langle 0_A,\alpha\rangle$.

Definition 4.2.1. Let $\mathbf{A}$ be a finite algebra.

1. A *tolerance* τ on $\mathbf{A}$ is a symmetric and reflexive subuniverse of $\mathbf{A}^2$.
2. Let $\langle 0_A,\theta\rangle$ be a prime congruence quotient of $\mathbf{A}$. The *basic tolerance* τ_θ for $\langle 0_A,\theta\rangle$ is the intersection of all tolerances τ satisfying $0_A \neq \tau \subseteq \theta$.

The next theorem is a special case of Theorem 5.27 from [29]. In its statement, the notation $(\theta\times\theta)\restriction_R$ represents

$$\big\{\big((a,c),(b,d)\big) : (a,b),(c,d)\in\theta\big\}\cap R^2,$$

which is a congruence on $\mathbf{R}$.

Theorem 4.2.2 (Hobby, McKenzie)**.** *Let $\langle 0_A,\theta\rangle$ be a prime congruence quotient of a finite 2-semilattice $\mathbf{A}$ and let τ_θ be its basic tolerance. Let $\mathbf{R}$ be the subalgebra of $\mathbf{A}^2$ with universe τ_θ. If there are more than four congruences, $\psi\in\mathrm{Con}(\mathbf{R})$ satisfying $0_R\le\psi\le(\theta\times\theta)\restriction_R$, then $\langle 0_A,\theta\rangle$ has TCT type* $\mathbf{5}$*.*

Theorem 4.2.3. *Let θ be a minimal congruence of a finite 2-semilattice $\mathbf{A}$. If Theorem 4.1.1 (1) holds, then $\langle 0_A,\theta\rangle$ has TCT type* $\mathbf{5}$*. If Theorem 4.1.1 (2) holds, the TCT type of $\langle 0_A,\theta\rangle$ can be* $\mathbf{3}$*,* $\mathbf{4}$ *or* $\mathbf{5}$*.*

Proof. The second assertion is demonstrated by the algebras in Example 4.3.5, so we will only prove the first here. Assume θ is minimal and satisfies Theorem 4.1.1 (1), and that X, Y, and f are as in the statement of Theorem 4.1.1 (1). That is, X and Y are distinct components of $(A, \longrightarrow)$, $f : X \to Y$, and

$$\theta = \{(u, v) \in (X \cup Y)^2 : f(u) = f(v) \text{ or } f(u) = v \text{ or } f(v) = u\} \cup 0_A.$$

Suppose τ_θ is the basic tolerance of θ and $x_1 \neq x_2$ are elements of X such that $(x_1, x_2) \in \tau_\theta$. Then $(x_1, x_2) \in \theta$, so there is $y \in Y$ such that $f(x_1) = f(x_2) = y$. This $y \in Y$ has the property that $|y/\theta| \neq 1$. Since θ is minimal, $\theta = \mathrm{Cg}(x_1, x_2)$. By Theorem 3.0.2 (2), there is an $n+1$-ary term operation t of $\mathbf{A}$, elements $a_1, \ldots, a_n \in A$, and $z \neq y$ such that $(t(x_1, a_1, \ldots, a_n), t(x_2, a_1, \ldots, a_n)) = (y, z)$ or $(t(x_1, a_1, \ldots, a_n), t(x_2, a_1, \ldots, a_n))) = (z, y)$. Either way, this gives a $z \in A$ so that $y \neq z$ and $(y, z) \in \tau_\theta$. Because of the structure of θ, we now can conclude that τ_θ contains a pair $(x, y) \in X \times Y$ such that $f(x) = y$. Now suppose $x' \in X$ is such that $x \longrightarrow x'$. Then $(xx', yx') = (x', yx') \in \tau_\theta \subseteq \theta$. It is not possible that $x' = yx'$ because $x' \in X$ and $y \longrightarrow yx'$, so $yx' \notin X$. This means $yx' \in Y$, and it follows that $f(x') = yx'$ and $(x', f(x')) \in \tau_\theta$. Because X is strongly connected, this reasoning can be applied to get $(x, f(x)) \in \tau_\theta$ for all $x \in X$. Since τ_θ is symmetric, we also have $(f(x), x) \in \tau_\theta$ for all $x \in X$. We have shown that

$$\sigma = \{(x, f(x)) : x \in X\} \cup \{(f(x), x) : x \in X\} \cup 0_A \subseteq \tau_\theta,$$

so if we can show that σ is a tolerance, we will have that $\sigma = \tau_\theta$ by the minimality of τ_θ. We already know that $\sigma \subseteq \theta$. In fact, σ is exactly $[\theta - \theta \cap (X \times X)] \cup 0_A$. Assume $(a, b), (c, d) \in \sigma$. Then $(a, b), (c, d) \in \theta$, so $(ac, bd) \in \theta$. The only way this pair can fail to be in σ is if $ac \neq bd$ and $ac, bd \in X$. If this is the case, then either $a \neq b$ or $c \neq d$. Without loss of generality, assume $a \neq b$. Since $(a, b) \in \sigma$, either a or b is in Y, but this means either ac or bd is not in X. We conclude that σ is a subuniverse of $\mathbf{A}^2$, which shows that it is a tolerance because it is symmetric and reflexive by construction. Therefore, $\sigma = \tau_\theta$.

Now that we have our hands on τ_θ, we can show that the main hypothesis of Theorem 4.2.3 is satisfied. Let $\mathbf{R} \leq \mathbf{A}^2$ have universe τ_θ. We extend f to a function $f : X \cup Y \to Y$ by $f(b) = b$ for $b \in Y$. Note that if $(a, b) \in \theta \cap (X \cup Y)^2$, then $f(a) = f(b)$. Define $\psi \subseteq R^2$ to be the relation

$$\psi = \big\{\big((a, b), (c, d)\big) \in R^2 : f(a) = b = d = f(c)\big\} \cup 0_R.$$

We will show that ψ is a congruence on $\mathbf{R}$. First of all ψ is reflexive and symmetric by definition. For transitivity, assume $\big((a, b), (c, d)\big), \big((c, d), (e, g)\big) \in \psi$. If $(a, b) = (c, d)$ or $(c, d) = (e, g)$, then $\big((a, b), (e, g)\big) \in \psi$, so we assume $(a, b) \neq (c, d)$ and $(c, d) \neq (e, g)$. This implies $\big((a, b), (c, d)\big), \big((c, d), (e, g)\big) \notin 0_R$, so $b = d$ and $d = g$. Therefore, $a \neq c$ and $c \neq e$. We also have that $f(a) = f(c) = b = d$ and $f(c) = f(e) = d = g$, so $f(a) = f(e) = b = d$. Now that we have shown ψ is an equivalence relation, we can verify that ψ is a congruence. By Proposition 3.0.2 (1), it suffices to show for every $(u, v) \in R$ and $\big((a, b), (c, d)\big) \in \psi$ that $\big((ua, vb), (uc, vd)\big) \in \psi$. If $(a, b) = (c, d)$ then $\big((ua, vb), (uc, vd)\big) \in 0_R \subseteq \psi$, so we assume $(a, b) \neq (c, d)$. This means $b = d$, $f(a) = f(c) = b = d$, and $\{a, c\} \cap X \neq \varnothing$. The last of these three assertions is because if $a, c \in Y$, then $a = f(a) = f(c) = c$. Since $f : X \cup Y \to Y$, we have $b, d \in Y$. Assume, without loss of generality, that $a \in X$. Since $(u, v), (a, b), (c, d) \in \theta$, we have that $(ua, vb), (uc, vd) \in \theta$. If $vb \notin Y$, then $ua = vb$ because Y is a minimal θ-nontrivial component, $b \in Y$, and $b \longrightarrow vb$. Since $b = d$, we have, $vb = vd$, which means $uc = vd$ by the same reasoning. In the case that $vb \notin Y$, we have shown $\big((ua, vb), (uc, vd)\big) \in 0_R \subseteq \psi$. Otherwise, $vb = vd \in Y$, which means, by the structure of θ, that $ua, uc \in X \cup Y$, and $f(ua) = vb = vd = f(uc)$, so $\big((ua, vb), (uc, vd)\big) \in \psi$.

Similar to $(\theta \times \theta)\restriction_R$, define $(0_A \times \theta)\restriction_R$ by

$$\big\{\big((a, b), (c, d)\big) : a = c \text{ and } (b, d) \in \theta\big\} \cap R^2$$

and define $(\theta \times 0_A)\restriction_R$ similarly. Both of these relations are congruences on $\mathbf{R}$. Notice that the congruences $0_R, (\theta \times \theta)\restriction_R, \psi, (\theta \times 0_A)\restriction_R$, and $(0_A \times \theta)\restriction_R$ are all between 0_A and $(\theta \times \theta)\restriction_R$.

Figure 4.1: Distinguishing the congruences in the proof of Theorem 4.2.3

	0_R	$(\theta \times \theta)\restriction_R$	ψ	$(\theta \times 0_A)\restriction_R$	$(0_A \times \theta)\restriction_R$
$\big((b,a),(a,a)\big)$	no	yes	no	yes	no
$\big((a,b),(b,b)\big)$	no	yes	yes	yes	no
$\big((a,b),(a,a)\big)$	no	yes	no	no	yes

To complete the proof, we will show that they are all distinct. To see this, choose $a \in X$ and $b \in Y$ such that $(a,b) \in \theta$ and consider the table in Figure 4.1. The rows are indexed by pairs in R^2 and the columns are indexed by the five congruences above. A cell contains a "yes" if the pair is in the congruence, and "no" if the pair is not in the congruence. Notice that for each pair of congruences, there is a row where one of the congruences has "yes" in its column and the other congruence has "no" in its column. The hypotheses of Theorem 4.2.2 hold, so $\langle 0_A, \theta \rangle$ has TCT type **5**.

□

4.3 Maximal Congruences

The following Theorem is a special case of Theorem 5.26 in [29].

Theorem 4.3.1 (Hobby, McKenzie)**.** *Let α be a minimal congruence on a finite algebra $\mathbf{A}$. Then $\langle 0_A, \alpha \rangle$ has type **4** or **5** if and only if there is a partial order $\eta \leq \mathbf{A}^2$ with the property that α is the transitive closure of $\eta \cup \{(a,b) : (b,a) \in \eta\}$.*

We will use Theorem 4.3.1 to give a tame-congruence-theoretic dichotomy on finite simple 2-semilattices.

Theorem 4.3.2. *Let* $\mathbf{A}$ *be a finite simple* 2*-semilattice. If* $\mathbf{A}$ *is strongly connected, then* $\langle 0_A, 1_A\rangle$ *has TCT type* **3**. *Otherwise, it has TCT type* **5**.

Proof. If $\mathbf{A}$ is not strongly connected, then it has at least two components. However, simplicity and Lemma 4.1.3 together imply that $\mathbf{A}$ has at most two strongly connected components. Since 1_A is minimal, it must satisfy Theorem 4.1.1 (1), so it has TCT type **5** by Theorem 4.2.3.

We now assume $\mathbf{A}$ is simple and strongly connected, and that $\langle 0_A, 1_A\rangle$ has type **4** or **5**. By Theorem 4.3.1, there is a partial order $\eta \leq \mathbf{A}^2$ such that the transitive closure of $\eta \cup \{(a,b) : (b,a) \in \eta\}$ is 1_A. Since η is reflexive, $\eta \leq_{\text{sd}} \mathbf{A} \times \mathbf{A}$. Let $u, v \in A$. Since $(u,v) \in 1_A$, we have that (u,v) is in the transitive closure of $\eta \cup \{(a,b) : (b,a) \in \eta\}$. This means there are $w_1, w_2, \ldots, w_n$ such that

$$(u, w_1), (w_2, w_1), (w_2, w_3), (w_4, w_3), \ldots, (w_{n-1}, w_n), (v, w_n) \in \eta.$$

If we think of R as a bipartite graph as in Definition 3.3.3, this shows that there is a walk from every vertex on the left to every other vertex on the left. Since $(v,v) \in \eta$ by reflexivity, we conclude that $\eta \leq_{\text{sd}} \mathbf{A} \times \mathbf{A}$ is linked. Since $\mathbf{A}$ is strongly connected, it is absorption free by Theorem 3.3.2. By the Absorption Theorem, $\eta = A \times A$, which is not antisymmetric unless $|A| = 1$. If $|A| = 1$, then $|A|$ is not simple, so we have contradicted the fact that η is a partial order. This shows that if $\mathbf{A}$ is simple and strongly connected, then $\langle 0_A, 1_A\rangle$ can not have TCT type **4** or **5**. By Theorem 9.10 from [29], $\langle 0_A, 1_A\rangle$ can not have type **1** or **2**, so it must have type **3**. □

By the correspondence theorem and the remark before Definition 4.2.1, Theorem 4.3.2 implies that if $\mathbf{A}$ is a finite 2-semilattice and $\alpha \in \mathrm{Con}(\mathbf{A})$ is maximal, then $\langle \alpha, 1_A\rangle$ has type **3** exactly when the quotient is strongly connected, and type **5** otherwise. A two-element 2-semilattice is a semilattice, and a semilattice is simple if and only if it has two elements by Proposition 3.0.3. Therefore, among maximal congruences α on $\mathbf{A}$ with $\langle\alpha, 1_A\rangle$ of type **5**, the quotient $\mathbf{A}/\alpha$ is a semilattice if and only if α has two equivalence classes. In Chapter 5.2 we will show that the appearance of a type **5** maximal congruence with more than two equivalence classes implies the generated variety has no cardinal bound on the size of its subdirectly irreducible members.

Proposition 4.3.3. *Let* $\mathbf{A}$ *be a finite* 2*-semilattice and* θ *be the smallest congruence such that* $\mathbf{A}/\theta$ *is a semilattice. Set* $\Phi = \{\varphi \in \mathrm{Con}(\mathbf{A}) : |A/\varphi| = 2\}$*. Then* $\theta = \bigwedge \Phi$*.*

Proof. In the paragraph following Definition 2.0.7, it was mentioned that for any algebra $\mathbf{A}$,

$$0_A = \bigwedge\{\alpha : \mathbf{A}/\alpha \text{ is subdirectly irreducible}\}.$$

By the correspondence theorem, θ is equal to the meet of all congruences, $\alpha \in \mathrm{Con}(\mathbf{A})$ with the property that $\mathbf{A}/\alpha$ is a subdirectly irreducible semilattice. It follows from Proposition 3.0.3 and the fact that a two-element 2-semilattice is a semilattice that α has this property if and only if $|A/\alpha| = 2$. Therefore, $\theta = \bigwedge \Phi$. $\square$

We can also prove a somewhat similar result for the meet of maximal congruences α with the TCT type of $\langle\alpha, 1_A\rangle$ equal to **3**.

Proposition 4.3.4. *Suppose* $\mathbf{A}$ *is a finite* 2*-semilattice and let* S *be the set of maximal congruences* φ *on* $\mathbf{A}$ *such that* $\langle\varphi, 1_A\rangle$ *has type* **3***. Set* $\Phi = \bigwedge S$*. Then*

$$\mathbf{A}/\Phi \cong \prod_{\varphi \in S} \mathbf{A}/\varphi.$$

Proof. Since $\mathbf{A}$ is finite, so is S, so let $\alpha_1, \ldots, \alpha_n$ be an enumeration of the members of S. The homomorphism $h : \mathbf{A} \to \mathbf{A}/\alpha_1 \times \cdots \times \mathbf{A}/\alpha_n$ given by $a \mapsto (a/\alpha_1, \ldots, a/\alpha_n)$ has kernel Φ. This means $\mathbf{A}/\Phi$ embeds subdirectly in $\mathbf{A}/\alpha_1 \times \cdots \times \mathbf{A}/\alpha_n$. If $\mathbf{B}$ is the image

Figure 4.2: The operation table of **A** from Example 4.3.5

$*$	0	1	2	3
0	0	1	2	0
1	1	1	2	1
2	2	2	2	3
3	0	1	3	3

of h, we have $\mathbf{B} \leq_{\text{sd}} \mathbf{A}/\alpha_1 \times \cdots \times \mathbf{A}/\alpha_n$. For each i, $\langle \alpha_i, 1_A \rangle$ has type **3**, so $\mathbf{A}/\alpha_i$ is strongly connected by Theorem 4.3.2. Also, each α_i is maximal so $\mathbf{A}/\alpha_i$ is simple. By Lemma 3.4.4, the result will follow if we can show that $\text{pr}_{i,j} \circ h(A) = \mathbf{A}/\alpha_i \times \mathbf{A}/\alpha_j$ for each i and j. Since α_i and α_j are different and both maximal, there is a pair (a, b) such that $(a, b) \in \alpha_i$ but $(a, b) \notin \alpha_j$. Applying $\text{pr}_{i,j} \circ h$ gives $\text{pr}_{i,j} \circ h(a) = (a/\alpha_i, a/\alpha_j)$ and $\text{pr}_{i,j} \circ h(b) = (b/\alpha_i, b/\alpha_j)$. These pairs are equal in the first coordinate, and unequal in the second. This shows $\text{pr}_{i,j} \circ h(A) \leq_{\text{sd}} \mathbf{A}/\alpha_i \times \mathbf{A}/\alpha_j$ is not the graph of a bijection, so it must be $\mathbf{A}/\alpha_i \times \mathbf{A}/\alpha_j$ by Lemma 3.4.2. As mentioned earlier, the result now follows by Lemma 3.4.4. □

We will finish off this chapter with Example 4.3.5. It finishes the proof of Theorem 4.2.3. That is, Example 4.3.5 presents minimal congruences satisfying Theorem 4.1.1 (2) whose corresponding quotients have each of the types **3**, **4**, and **5**. The Universal Algebra Calculator [23] was used to compute congruences on the algebras presented in Example 4.3.5.

Example 4.3.5. The directed three cycle $\mathbf{T}_3$ from Definition 1.0.1 is simple and strongly connected, so the congruence quotient $\langle 0_T, 1_T \rangle$ has TCT type **3** by Theorem 4.3.2. Since 1_A is minimal in this case, this gives an example of a type **3** minimal congruence which satisfies Theorem 4.1.1 (2).

The algebra $\mathbf{A} = (\{0, 1, 2, 3\}, *)$, where $*$ is defined in Figure 4.2, has a minimal congruence θ which only collapses 0 with 1. Figure 4.3 is a picture of the digraph of $\mathbf{A}$, and it

Figure 4.3: The graph of $\mathbf{B}$ from Example 4.3.5

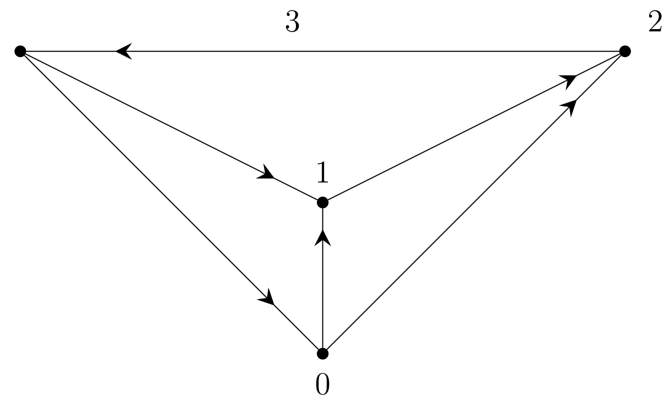

is not hard to see from this picture that $\mathbf{A}/\theta \cong \mathbf{T}_3$. The algebra $\mathbf{A}$ is strongly connected, so θ trivially satisfies Theorem 4.1.1 (2). Since θ only collapses two elements, the only nontrivial tolerance it contains is itself. This means θ is its own basic tolerance. If $\mathbf{R}$ is the subalgebra of $\mathbf{A}^2$ whose universe is θ, one can easily check that the partitions below all correspond to equivalence relations which are congruences:

$$\Big|(0,0)\Big|(0,1)\,(1,1)\Big|(1,0)\Big|(2,2)\Big|(3,3)\Big|$$

$$\Big|(0,0)\Big|(0,1)\Big|(1,0)\,(1,1)\Big|(2,2)\Big|(3,3)\Big|$$

$$\Big|(0,0)\Big|(0,1)\,(1,0)\,(1,1)\Big|(2,2)\Big|(3,3)\Big|.$$

Furthermore, these congruences are all strictly between 0_R and $(\theta \times \theta)\restriction_R$, so $\langle 0_A, \theta\rangle$ has type **5** by Theorem 4.2.3.

Congruence quotients of type **4** seem somewhat more elusive in the variety of 2-semilattices. The following 10 element 2-semilattice is the smallest the author has found with a type **4** congruence quotient. The domain is $B = \{0, 1, 2, \ldots, 9\}$ and $*$ is defined in Figure 4.4

Figure 4.4: The operation table of **B** from Example 4.3.5

$*$	0	1	2	3	4	5	6	7	8	9
0	0	0	2	3	0	0	2	9	0	9
1	0	1	2	3	1	1	3	9	1	9
2	2	2	2	3	4	5	2	4	9	9
3	3	3	3	3	4	5	3	5	9	9
4	0	1	4	4	4	4	9	4	0	9
5	0	1	5	5	4	5	9	5	1	9
6	2	3	2	3	9	9	6	9	9	9
7	9	9	4	5	4	5	9	7	9	9
8	0	1	9	9	0	1	9	9	8	9
9	9	9	9	9	9	9	9	9	9	9

Of course, this is an eleven by eleven array that is rather offensive to the eye, so Figure 4.5 is a partial picture of the digraph of **B** and is intended to explain the operation and how the type **4** congruence quotient arises.

The picture omits the vertex 9, the arrow from each other vertex to 9, and all loops. When constructing this algebra, the plan was to have a unary polynomial $p(x)$ which would send $(0, 1)$ to $(2, 3)$, a unary polynomial $q(x)$ to send $(2, 3)$ to $(4, 5)$, and a unary polynomial $r(x)$ to send $(4, 5)$ back to $(0, 1)$. Introducing the element 6 and setting $p(x) = 6*x$ achieved this. Similarly, $q(x)$ is defined to be $7*x$ and $r(x)$ is $8*x$. Each of $\{0, 1\}$, $\{2, 3\}$, and $\{4, 5\}$ are two element subuniverses which are semilattices. The key here is that the subalgebra with universe $\{0, 1\}$ is a meet semilattice (the operation chooses the minimum), and the subalgebra with universe $\{2, 3\}$ is a join semilattice (the operation chooses the maximum). The element 9 is introduced as a default for all remaining undefined products. The UACalc was used to verify that the partition

$$\left|0\,1\middle|2\,3\middle|4\,5\middle|6\middle|7\middle|8\middle|9\right|$$

is a congruence on **B**. We will refer to this congruence as θ. It is a minimal congruence since collapsing any of its nontrivial blocks generates all of θ using the polynomials p, q, and r above. Notice that $0, 1, 2, 3, 4$, and 5 are mutually reachable from one another via directed walks. Therefore, they are all in the same component of $\mathbf{B}$. This means there is only one θ-nontrivial class in the sense of Definition 4.1.4. We have shown that θ satisfies Theorem 4.1.1 (2).

Now let $\mu = \{(0,1), (2,3), (4,5)\} \cup 0_B$. This is a partial order which is a subuniverse of $\mathbf{B}^2$. The relation $\mu \cup \{(a,b) : (b,a) \in \mu\}$ is θ, so its transitive closure is certainly θ. By Theorem 4.3.1, $\langle 0_B, \theta \rangle$ has type either **4** or **5**. Following the discussion at the beginning of Section 4.2, define a unary polynomial on $\mathbf{B}$ by $q(x) = 8*(1*x)$. Notice that $8*(1*0) = 0$ and $8*(1*1) = 1$, so $q(\theta) \nsubseteq 0_A$. It is also easy to check that $q(B) = U = \{0, 1, 9\}$. We now argue that U is a $\langle 0_A, \theta \rangle$-minimal set. Notice that any unary polynomial f has $f(9) = 9$, so any minimal set contains 9. A constant unary polynomial will collapse θ to 0_A, so a minimal set must contain at least two elements. Therefore, if $\{0, 1, 9\}$ is not minimal, then either $\{0, 9\}$ or $\{1, 9\}$ is. However, a polynomial with such a range necessarily collapses θ to 0_A. We conclude that U is a $\langle 0_A, \theta \rangle$-minimal set. The only nontrivial $\theta\restriction_U$ class is $N = \{0, 1\}$, and we know that $\mathbf{B}\restriction_N$ is either a semilattice or a lattice. The set $N = \{0, 1\}$ is closed under $*$ and has $0*0 = 0$ and $0*1 = 1*0 = 1*1 = 1$. Also, the binary polynomial given by $t(x, y) = 8*(7*((6*x)*(6*y)))$ has $t(1, 1) = 1$ and $t(0, 1) = t(1, 0) = t(0, 0) = 0$. Defining $x \wedge y = x * y$ and $x \vee y = t(x, y)$, we have shown that $\mathbf{A}\restriction_N$ has at least the meet and join operations of a lattice, so it cannot be a semilattice and hence, $\langle 0_A, \theta \rangle$ must have type **4**.

Figure 4.5: The graph of **B** from Example 4.3.5

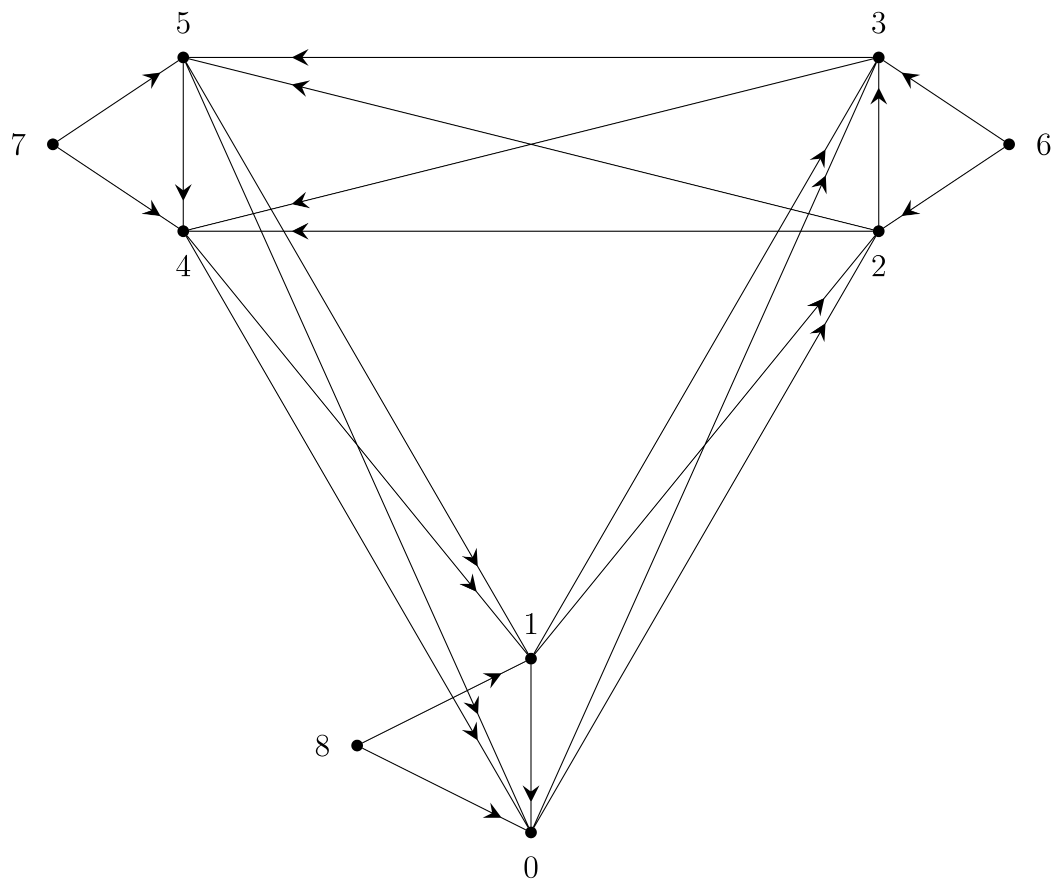

Chapter 5

Residually Large Varieties

5.1 Background

Recall from Definition 2.0.7 that an algebra $\mathbf{A}$ is called subdirectly irreducible if whenever $h : \mathbf{A} \hookrightarrow \prod_{i \in I} \mathbf{A}_i$ is a subdirect embedding, $\mathrm{pr}_i \circ h : \mathbf{A} \to \mathbf{A}_i$ is an isomorphism for some i. This really means that the algebra cannot be represented as a subdirect product in a nontrivial way, as the name "subdirectly irreducible" suggests. In the discussion following Definition 2.0.7, we also mentioned Birkhoff's theorem from [8] that every algebra embeds subdirectly in a product of subdirectly irreducible algebras. In fact, Birkhoff showed that every algebra embeds subdirectly in a product of its subdirectly irreducible quotients. A weaker fact is that if $\mathcal{V}$ is a variety and $\mathcal{V}_{\mathrm{si}}$ is the class[1] of subdirectly irreducible members in $\mathcal{V}$, then every $\mathbf{A} \in \mathcal{V}$ embeds subdirectly in a product of algebras from $\mathcal{V}_{\mathrm{si}}$. It follows that $\mathcal{V}_{\mathrm{si}}$ generates the variety $\mathcal{V}$. Therefore, much can be learned about a variety $\mathcal{V}$ by studying $\mathcal{V}_{\mathrm{si}}$. We now give a definition which is standard in the literature:

[1]There are plenty of examples where $\mathcal{V}_{\mathrm{si}}$ is not a set.

Definition 5.1.1. Let $\mathcal{V}$ be a variety.

1. If there is a cardinal κ such that $|\mathbf{A}| < \kappa$ for every $\mathbf{A} \in \mathcal{V}_{\mathrm{si}}$, then $\mathcal{V}$ is called *residually small*.

2. If $\mathcal{V}$ is not residually small, it is called *residually large*.

3. If $\mathcal{V}$ satisfies (1) with $\kappa = \omega$, we say that $\mathcal{V}$ is *residually finite*.

4. If $\mathcal{V}$ satisfies (1) with $\kappa = n$ for some $n \in \omega$, then we say $\mathcal{V}$ is *residually very finite*.

For an arbitrary variety $\mathcal{V}$, completely understanding $\mathcal{V}_{\mathrm{si}}$ is an unreasonable expectation. One way of restricting the question of understanding $\mathcal{V}_{\mathrm{si}}$ given $\mathcal{V}$ is to assume $\mathcal{V} = \mathbf{HSP}(\mathbf{A})$ for some finite algebra, $\mathbf{A}$. In particular, the question "given a finite algebra $\mathbf{A}$, which parts of Definition 5.1.1 does $\mathbf{HSP}(\mathbf{A})$ satisfy?" has received attention in the literature. It was shown by Papert in [45] and stated in Proposition 3.0.3 that the only subdirectly irreducible semilattice has two elements. It follows from Birkhoff's theorem that $\mathbf{HSP}(\mathbf{A})$ contains a subdirectly irreducible algebra if $\mathbf{A}$ is nontrivial. Therefore, if $\mathbf{A}$ is a semilattice, then $\mathbf{HSP}(\mathbf{A})$ has exactly one subdirectly irreducible member up to isomorphism, so $\mathbf{HSP}(\mathbf{A})$ is residually very finite. In fact, it follows from Birkhoff's theorem that $\mathbf{HSP}(\mathbf{A})$ is the variety of all semilattices whenever $\mathbf{A}$ is a nontrivial semilattice. Maróti was able to show in [40] that if $\mathbf{A}$ is a finite tournament, then $\mathbf{HSP}(\mathbf{A})$ is residually very finite. By a result of Freese and McKenzie in [24], if $\mathbf{A}$ is a finite group, then $\mathbf{HSP}(\mathbf{A})$ is residually very finite if all of the Sylow subgroups of $\mathbf{A}$ are abelian, and residually large otherwise.

McKenzie showed in [42] that there is no algorithm which takes a finite algebra $\mathbf{A}$ as input and decides which parts of Definition 5.1.1 $\mathbf{HSP}(\mathbf{A})$ satisfies. However, as mentioned in the previous paragraph, there is sometimes an algorithm when the question is localized. For example, the question "Given a finite semilattice $\mathbf{A}$, which parts of Definition 5.1.1

does $\mathbf{HSP}(\mathbf{A})$ satisfy?" always has the answer (1), (3), and (4). For the same question with "semilattice" replaced by "group", one algorithm which gives a correct answer is to compute the Sylow subgroups of $\mathbf{A}$ and check to see if they are abelian. If they are all abelian, the answer is (1), (3), and (4), and otherwise, the answer is (2).

By a Theorem of McKenzie in [43], which was re-phrased by Davey, Pitkethly, and Willard in [19], there is a characterization of when a finite 2-semilattice $\mathbf{A}$ has the property that $\mathbf{HSP}(\mathbf{A})$ is residually large. The characterization says that $\mathbf{HSP}(\mathbf{A})$ is residually large if and only if there is a finite algebra $\mathbf{B}$ which embeds in some power of $\mathbf{A}$ and satisfies a list of properties. However, this characterization is difficult to use. In fact, if $\mathbf{A}$ is a finite 2-semilattice, it is not known if determining whether or not $\mathbf{HSP}(\mathbf{A})$ contains such an algebra is decidable. The results in this chapter give some decidable conditions on a 2-semilattice $\mathbf{A}$ which are sufficient to imply $\mathbf{HSP}(\mathbf{A})$ is residually large.

5.2 Acyclic and Weakly Acyclic 2-Semilattices

We start this section by defining what it means for a 2-semilattice to be "weakly acyclic" and "acyclic". In Section 5.3, we use the facts collected in this section to prove that if $\mathbf{A}$ is a weakly acyclic 2-semilattice which is not a semilattice, then it generates a residually large variety.

Definition 5.2.1. Let $\mathbf{A}$ be a 2-semilattice and $\mathcal{V}$ be a variety of 2-semilattices.

1. $\mathbf{A}$ is *acyclic* if $a_1 \longrightarrow a_2 \longrightarrow a_3 \longrightarrow \cdots \longrightarrow a_n \longrightarrow a_1$ implies $a_1 = a_2 = \cdots = a_n$.
2. $\mathbf{A}$ is *weakly* acyclic if all of its strongly connected subuniverses are singletons.
3. $\mathcal{V}$ is *acyclic* if all of its members are acyclic.

4. $\mathcal{V}$ is *weakly acyclic* if it has no nontrivial finite strongly connected members.

We now prove that if $\mathbf{A}$ is finite and acyclic, then $\mathbf{HSP}(\mathbf{A})$ is acyclic. We will later prove the analogous result where "acyclic" is replaced with "weakly acyclic".

Theorem 5.2.2. *Acyclicity is closed under taking subalgebras and products. Furthermore, if $\mathbf{A}$ is acyclic $\mathbf{HSP}(\mathbf{A})$ is locally finite, then all of its quotients are acyclic.*

Proof. To see that acyclicity is closed under taking subalgebras, suppose $\mathbf{A}$ is acyclic and $\mathbf{B} \leq \mathbf{A}$. If

$$a_1 \longrightarrow \cdots \longrightarrow a_n \longrightarrow a_1$$

in $\mathbf{B} \leq \mathbf{A}$, then

$$a_1 \longrightarrow \cdots \longrightarrow a_n \longrightarrow a_1$$

in $\mathbf{A}$. Therefore, $a_1 = \cdots = a_n$ since $\mathbf{A}$ is acyclic, which shows that $\mathbf{B}$ is also acyclic. Suppose $\prod_{i \in I} \mathbf{A}_i$ is a product of acyclic 2-semilattices and let $f_1, \ldots, f_n \in \prod_{i \in I} A_i$ satisfy

$$f_1 \longrightarrow f_2 \longrightarrow \cdots f_n \longrightarrow f_1.$$

Then for each i,

$$f_1(i) \longrightarrow f_2(i) \longrightarrow \cdots \longrightarrow f_n(i) \longrightarrow f_1(i),$$

so

$$f_1(i) = f_2(i) = \cdots = f_n(i)$$

for each i because $\mathbf{A}_i$ is acyclic for each i. Since $i \in I$ was arbitrary, $f_1 = f_2 = \cdots = f_n$. Now suppose $\mathbf{A}$ is an acyclic 2-semilattice with $\mathbf{HSP}(\mathbf{A})$ locally finite. Choose $\theta \in \mathrm{Con}(\mathbf{A})$ and $a_1, \cdots, a_n \in A$ with $a_1/\theta \longrightarrow a_2/\theta \longrightarrow \cdots \longrightarrow a_n/\theta \longrightarrow a_1/\theta$. It follows for $B = \mathrm{Sg}_{\mathbf{A}}(a_1, \ldots, a_n)$ that

$$a_1/\theta{\restriction}_B \longrightarrow a_2/\theta{\restriction}_B \longrightarrow \cdots \longrightarrow a_n/\theta{\restriction}_B \longrightarrow a_1/\theta{\restriction}_B,$$

and if

$$a_1/\theta\restriction_B = a_2/\theta\restriction_B = \cdots = a_n/\theta\restriction_B,$$

then

$$a_1/\theta = a_2/\theta = \cdots = a_n/\theta.$$

Since $\mathbf{HSP}(\mathbf{A})$ is locally finite, $\mathbf{B}$ is finite, so it suffices to prove the result assuming $\mathbf{A}$ is finite. From now on we will assume $\mathbf{A}$ is finite, $\theta \in \mathrm{Con}(\mathbf{A})$, and that

$$a_1/\theta \longrightarrow \cdots \longrightarrow a_n/\theta \longrightarrow a_1/\theta.$$

For $1 \leq i \leq n$ and $j \geq 1$, define a_i^j as follows: $a_1^1 = a_1$, and for $i + j \geq 2$, $a_i^j = a_{i-1}^j * a_i$ if $i \geq 2$, and $a_i^j = a_n^{j-1} * a_1$ if $i = 1$. The a_i^j are defined so that

$$a_1^1 \longrightarrow a_2^1 \longrightarrow \cdots \longrightarrow a_n^1 \longrightarrow a_1^2 \longrightarrow \cdots \longrightarrow a_n^2 \longrightarrow a_1^3 \longrightarrow \cdots \longrightarrow a_n^3 \longrightarrow \cdots,$$

and furthermore, $a_i \overset{\theta}{\equiv} a_i^j$ for each i and for all $j \geq 1$. Since $\mathbf{A}$ is finite, there must some $1 \leq i \leq n$ and $j < k$ such that $a_i^j = a_i^k$, which means

$$a_i^j \longrightarrow \cdots \longrightarrow a_n^j \longrightarrow a_1^{j+1} \longrightarrow \cdots \longrightarrow a_n^{j+1} \longrightarrow a_1^{j+2} \longrightarrow \cdots \longrightarrow a_{i-1}^k \longrightarrow a_i^k.$$

This takes place in $\mathbf{A}$, so by acyclicity, all of the elements in the above chain are equal. We noted that $a_i \overset{\theta}{\equiv} a_i^j$ for all i, j, so $a_i/\theta = a_{i+1}/\theta = \cdots = a_n/\theta = a_1/\theta = \cdots = a_{i-1}/\theta$. This shows $\mathbf{A}/\theta$ is acyclic. □

Recall the algebra in Example 3.2.1. Since for any $a, b \in A$ it is true that $ab \geq a, b$ where $\geq$ is the usual order on the natural numbers, the algebra $\mathbf{A}$ is acyclic. However, it is constructed in such a way that the relation $\theta = \{(a, b) : a \equiv b \pmod 3\}$ is a congruence on $\mathbf{A}$ with $\mathbf{A}/\theta \cong \mathbf{T}_3$ where $\mathbf{T}_3$ is the directed 3-cycle from Definition 1.0.1. This shows that $\mathbf{HSP}(\mathbf{A})$ being locally finite is needed in the statement of Theorem 5.2.2.

Since the variety generated by a finite algebra is locally finite, Theorem 5.2.2 shows that if $\mathbf{A}$ is a finite acyclic 2-semilattice, then every member of $\mathbf{HSP}(\mathbf{A})$ is acyclic. The next proposition asserts that $\mathbf{5}$ is the only TCT type in such a variety's type set.

Theorem 5.2.3. *Let $\mathbf{A}$ be a finite acyclic 2-semilattice. Then every prime congruence quotient of $\mathbf{A}$ has type $\mathbf{5}$.*

Proof. Since acyclicity is closed under taking quotients by Theorem 5.2.2, it suffices to show that $\langle 0_A, \alpha \rangle$ has type $\mathbf{5}$ for every minimal $\alpha \in \mathrm{Con}(\mathbf{A})$. Since every component of $\mathbf{A}$ has one element, no minimal congruence can satisfy Theorem 4.1.1 (2), so α must satisfy Theorem 4.1.1 (1). By Theorem 4.2.3, the quotient has type $\mathbf{5}$. □

As mentioned earlier, we will now prove that if $\mathbf{A}$ is a finite weakly acyclic 2-semilattice, then $\mathbf{HSP}(\mathbf{A})$ is a weakly acyclic variety. We break the proof into three lemmas.

Lemma 5.2.4. *If $(\mathbf{A}_i : i \in I)$ is a family of weakly acyclic 2-semilattices, then the product $\mathbf{A} = \prod_{i \in I} \mathbf{A}_i$ is weakly acyclic.*

Proof. Suppose $\mathbf{B} \leq \mathbf{A}$ is strongly connected and has more than one element. We will show that for some $i \in I$, $\mathbf{A}_i$ has a strongly connected subalgebra with more than one element. Since $\mathbf{B}$ has more than one element, there is an $i \in I$ such that $\mathrm{pr}_i(B)$ has more than one element. If $\mathbf{B}_i$ is the algebra with this universe, then $\mathbf{B}_i \leq \mathbf{A}_i$. For distinct $b, c \in \mathbf{B}_i$, there must be $f, g \in \mathbf{B}$ with $f(i) = b$ and $g(i) = c$. Since $\mathbf{B}$ is strongly connected, there is a directed walk, $f = f_1 \longrightarrow f_2 \longrightarrow \cdots \longrightarrow f_n = g$ in $\mathbf{B}$. Therefore, $b = \mathrm{pr}_i(f_1) \longrightarrow \mathrm{pr}_i(f_2) \longrightarrow \cdots \longrightarrow \mathrm{pr}_i(f_n) = c$ is a directed walk from b to c in $\mathbf{B}_i$, which shows that $\mathbf{B}_i \leq \mathbf{A}_i$ is strongly connected. We have shown that if $\mathbf{A}$ has a nontrivial strongly connected subalgebra, then so does $\mathbf{A}_i$ for some $i \in I$. □

Similar to the acyclic case, we can also show that weak acyclicity is closed under taking quotients in the case that the algebra generates a locally finite variety. We will use the following Lemma.

Lemma 5.2.5. *Let* $\mathbf{A}$ *be a 2-semilattice and* $a_1, \ldots, a_n \in A$. *Suppose* $b = t(a_1, \ldots, a_n)$ *for some term operation,* $t^{\mathbf{A}}$ *of* $\mathbf{A}$. *For some* i, *there is a directed walk from* a_i *to* b *in* $\mathrm{Sg}_{\mathbf{A}}(a_1, \ldots, a_n)$.

Proof. We will use induction on how many times the operation of the 2-semilattice is mentioned in the term t. If $*$ is not mentioned in t, then $t = x_i$ for some i, so $b = a_i$ for some i. Now suppose t mentions the operation $n+1$ times for some $n \geq 0$. Then $t(x_1, \ldots, x_n) = r(x_1, \ldots, x_n) * s(x_1, \ldots, x_n)$ for some terms, r and s. This means $b = r(a_1, \ldots, a_n) * s(a_1, \ldots, a_n)$, so $r(a_1, \ldots, a_n) \longrightarrow b$. By the inductive hypothesis, one of the a_i has a directed walk to $r(a_1, \ldots, a_n)$ in $\mathrm{Sg}_{\mathbf{A}}(a_1, \ldots, a_n)$. Concatenating the walks gives the desired result. □

Lemma 5.2.6. *Let* $\mathbf{A}$ *be weakly acyclic with* $\mathbf{HSP}(\mathbf{A})$ *locally finite. If* θ *is a congruence of* $\mathbf{A}$ *with finitely many equivalence classes, then* $\mathbf{A}/\theta$ *is also weakly acyclic.*

Proof. Let us assume that $\mathbf{A}/\theta$ is finite and $\mathbf{B} \leq \mathbf{A}/\theta$ is strongly connected. This means $B = \{b_i/\theta\}_{i=1}^{n}$ for some $b_1, \ldots, b_n$ with $n \geq 3$. We now set $C = \mathrm{Sg}_{\mathbf{A}}(b_1, \ldots, b_n)$. Since $\mathbf{HSP}(\mathbf{A})$ is locally finite, $\mathbf{C}$ is finite. We also have that $\mathbf{C}/\theta{\restriction_C} \cong \mathbf{B}$. Since $\mathbf{C} \leq \mathbf{A}$, $\mathbf{C}$ is also weakly acyclic. We have produced a finite weakly acyclic algebra with a nontrivial strongly connected quotient. We will prove that this cannot happen.

From now on, we assume that $\mathbf{A}$ is finite and weakly acyclic, and that $\mathbf{A}/\theta$ is strongly connected. For each $a \in A$, the set a/θ is a subuniverse of $\mathbf{A}$. By Lemma 3.1.5 (4), the smallest strongly connected component of $(a/\theta, \longrightarrow)$ is also a subuniverse of $\mathbf{A}$. By weak acyclicity, the smallest strongly connected component of $(a/\theta, \longrightarrow)$ must be a singleton.

Again, let $b_1/\theta, \ldots, b_n/\theta$ be the θ-classes and for each i, let $b_i^\times$ be the unique element in $(b_i/\theta)'$, which satisfies $a \longrightarrow b_i^\times$ for every $a \in b_i/\theta$ by Lemma 3.1.5 (2). We will show that $D = \mathrm{Sg}_{\mathbf{A}}(b_1^\times, \ldots, b_n^\times)$ is strongly connected, which contradicts weak acyclicity of $\mathbf{A}$. Suppose $b_i/\theta \longrightarrow b_j/\theta$. Then $b_i^\times b_j^\times \in b_j/\theta$, but $b_j^\times \longrightarrow b_i^\times b_j^\times$, so $b_i^\times b_j^\times = b_j^\times$ by the choice of $b_j^\times$. This means $b_i^\times \longrightarrow b_j^\times$. Since $\mathbf{A}/\theta$ is strongly connected, we have shown that the set $\{b_1^\times, \ldots, b_n^\times\}$ is strongly connected. Now let $c_1, c_2 \in D$ and suppose $c_1 \in b_i^\times/\theta$, which means $c_1 \longrightarrow b_i^\times$. By the previous remark, there is a directed walk from c_1 to $b_j^\times$ for all j. By the remark following Definition 2.0.1, there is a term operation of $\mathbf{A}$ with $c_2 = t(b_1^\times, \ldots, b_n^\times)$. By Lemma 5.2.5, for some j there is a directed walk from some $b_j^\times$ to c_2 that mentions only elements of D. Therefore, there is a directed walk from c_1 to $b_j^\times$ for every j, and a directed walk from $b_j^\times$ to c_2 for some j. Concatenating these walks gives a directed walk from c_1 to c_2. Since $c_1, c_2 \in D$ were arbitrary, we have shown that $\mathbf{D}$ is strongly connected. □

Theorem 5.2.7. *Let* $\mathbf{A}$ *be a finite weakly acyclic 2-semilattice. Then* $\mathbf{HSP}(\mathbf{A})$ *is a weakly acyclic variety.*

Proof. Let $\mathbf{C} \in \mathbf{HSP}(\mathbf{A})$ be finite and nontrivial. By Birkhoff's Theorem, there is a set I, an algebra $\mathbf{B}$, and a congruence $\psi \in \mathrm{Con}(\mathbf{B})$ such that $\mathbf{B} \leq \mathbf{A}^I$ and $\mathbf{C} \cong \mathbf{B}/\psi$. By Lemma 5.2.4, $\mathbf{A}^I$ is weakly acyclic, so $\mathbf{B}$ is weakly acyclic. Since we are assuming $\mathbf{C}$ is finite, it is weakly acyclic by Lemma 5.2.6. In particular, $\mathbf{C}$ itself is not strongly connected. □

5.3 Acyclic and Weakly Acyclic Varieties are Residually Large

As promised at the beginning of Section 5.2, we will prove in this section that if $\mathbf{A}$ is a finite weakly acyclic 2-semilattice which is not a semilattice, then $\mathbf{HSP}(\mathbf{A})$ is residually large. This is Theorem 5.3.7. We first define a notion of height in 2-semilattices that will be useful in the proofs in this section.

Suppose $\mathbf{A}$ is a 2-semilattice. Recall that $\sim$ is the equivalence relation whose equivalence classes are the strongly connected components of $\mathbf{A}$. By the definition of $\xrightarrow{A/\sim}$ and Lemma 3.1.5 (2), we have $a/\sim \longrightarrow A'$ for every $a/\sim \in \mathbf{A}$. This certainly shows that there is at least one directed walk in $(A/\sim, \xrightarrow{A/\sim})$ from each component to A'. By Lemma 3.1.4 (2), $(A/\sim, \xrightarrow{A/\sim})$ is acyclic, so no walk from $a/\sim$ to A' can be longer than $|A/\sim|$ unless it uses loops. We will call a directed walk *loop-free* if it does not contain any loops.

Definition 5.3.1. Let $\mathbf{A}$ be a finite 2-semilattice. The height function $\lambda : A/\sim \to \omega$ is the length of any maximal-length, loop-free, directed walk from $a/\sim$ to A' in the digraph $(A/\sim, \longrightarrow)$. For convenience, $\lambda(A') = 0$. We lift λ to a function $\lambda : A \to \omega$ by $\lambda(a) = \lambda(a/\sim)$.

Lemma 5.3.2. *Let $\mathbf{A}$ be a finite 2-semilattice which is not strongly connected. Suppose I is a nonempty set, and $f : I \to A$ is an element of $(A - A')^I$. There are elements, $g_1, \ldots, g_n \in A^I$ such that*

$$f \xrightarrow{\mathbf{A}^I} g_1 \xrightarrow{\mathbf{A}^I} g_2 \xrightarrow{\mathbf{A}^I} \cdots \xrightarrow{\mathbf{A}^I} g_n$$

and $\lambda(g_n(\alpha)) = 1$ for all $\alpha \in I$.

Proof. We will first show that the result holds in $\mathbf{A}$. That is, given $a \in A - A'$, we want to find a directed walk from a to some b with $\lambda(b) = 1$. If $\lambda(a) = 1$ there is nothing to prove, so we will assume $\lambda(a) = m > 1$. By the remark after Definition 5.3.1, we have that $\lambda(a/\sim) = m$. There must exist some $c \in A$ with $a/\sim \longrightarrow c/\sim$ and $\lambda(c) = m - 1$. Let c have these properties. Since $a/\sim$ and $c/\sim$ are strongly connected and there are $a' \in a/\sim$ and $c' \in c/\sim$ with $a' \xrightarrow{\mathbf{A}} c'$, there is a directed walk from a to c. Inductively, there is a directed walk from c to some element b with $\lambda(b) = 1$. Concatenating these walks gives a directed walk in $\mathbf{A}$ from a to some b with $\lambda(b) = 1$.

We now suppose $f \in (A - A')^I$, which means $\lambda(f(\alpha)) \geq 1$ for each $\alpha \in I$. By the previous paragraph, there is a directed walk form $f(\alpha)$ to some b_α with $\lambda(b_\alpha) = 1$ for each α. Since $\mathbf{A}$ is finite, we can find such a directed walk of length at most $|A|$. Therefore, by possibly repeating elements and using that the digraph of $\mathbf{A}$ is reflexive, there is a fixed n and elements a_α^i for each α and $1 \leq i \leq n$ so that

$$f(\alpha) \longrightarrow a_\alpha^1 \longrightarrow a_\alpha^2 \longrightarrow \cdots \longrightarrow a_\alpha^n \longrightarrow b_\alpha.$$

For each i, define g_i by $g_i(\alpha) = a_\alpha^i$. By definition, we have $f \longrightarrow g_1 \longrightarrow \cdots \longrightarrow g_n \longrightarrow b$ where $b(\alpha) = b_\alpha$ for each α. This is a directed walk from f to b in the digraph of $\mathbf{A}^I$, and b has the property that $\lambda(b(\alpha)) = \lambda(b_\alpha) = 1$ for all α. $\square$

Lemma 5.3.3 is similar to Lemma 5.3.2.

Lemma 5.3.3. *Let $\mathbf{A}$ be a finite 2-semilattice which is not strongly connected, A' be its smallest strongly connected component, $D = \{a \in A : \lambda(a) = 1\}$, and I be any set. If D is strongly connected as a digraph, then for any $f \in (A - A')^I$ and $g \in D^I$, there is a directed walk from f to g in the digraph of $\mathbf{A}^I$.*

Proof. We suppose D is a strongly connected subset of A, that $f \in (A - A')^I$, and $g \in D^I$. Since D is finite and reflexive as a digraph, we can follow the proof of Lemma 3.1.8 (2) to show that D^I is also strongly connected and for any $g' \in D^I$, there is a directed walk in D^I from g' to g of length at most $|D|$. Lemma 5.3.2 provides a walk in the digraph of $\mathbf{A}^I$ from f to some $g' \in D^I$, so we can concatenate to obtain a directed walk from f to g. $\square$

Lemma 5.3.4. *If $\mathbf{A}$ is a finite subdirectly irreducible 2-semilattice and $|A'| = 1$, then there is a unique strongly connected component D such that $\lambda(D) = 1$.*

Proof. Let 0 be the unique element in A'. By Lemma 4.1.3, if D is a strongly connected component and $\lambda(D) = 1$, then the relation $\psi_D = (D \cup \{0\})^2 \cup 0_A$ is a nonzero congruence.

If there are two distinct such components, D_1 and D_2, then $\psi_{D_1} \cap \psi_{D_2} = 0_A$, which cannot happen since $\mathbf{A}$ is subdirectly irreducible. $\square$

Lemma 5.3.5 is the key to the rest of the results in this chapter.

Lemma 5.3.5. *Let* $\mathbf{A}$ *be a finite subdirectly irreducible* 2*-semilattice such that* $|A| > 2$, $|A'| = 1$, *and* $\{(a,b) : \lambda(a), \lambda(b) \in \{0,1\}\} \cup 0_A$ *is the unique minimal congruence of* $\mathbf{A}$. *The variety* $\mathbf{HSP}(\mathbf{A})$ *is residually large.*

We point out that a similar idea to that in Lemma 5.3.5 was used to produce subdirectly irreducible commutative directoids by both Ježek and McNulty in [33] and Gardner and Parmenter in [27].

Proof. By Lemma 5.3.4, $\mathbf{A}$ has a unique $\sim$ - class D with $\lambda(D) = 1$. Therefore, the hypotheses of Lemma 5.3.3 hold in $\mathbf{A}$.

Let I be any set and consider the algebra $\mathbf{A}^I$. Also, let 0 be the unique element in A' and $\mathbf{0}$ be the constant 0 function in A^I. The relation

$$\theta = \{(f,g) : f = g \text{ or } 0 \in \operatorname{ran}(f) \cap \operatorname{ran}(g)\}$$

is a congruence on $\mathbf{A}^I$. This is true because if $0 \in \operatorname{ran}(f)$, then $0 \in \operatorname{ran}(fh)$ for any $h \in A^I$. Now let $\psi > \theta$ and choose $(f,g) \in \psi - \theta$. This means there $f, g \in A^I$ such that $f \neq g$ and $(f,g) \in \psi$, but either $0 \notin \operatorname{ran}(f)$ or $0 \notin \operatorname{ran}(g)$. First, assume $0 \notin \operatorname{ran}(f)$ and $0 \in \operatorname{ran}(g)$. By Lemma 5.3.3, for any $h \in D^I$, there are $h_1, \ldots, h_n$ such that $f \longrightarrow h_1 \longrightarrow h_2 \longrightarrow \cdots \longrightarrow h_n = h$. This means

$$(fh_1h_2\cdots h_n, gh_1h_2\cdots h_n) = (h, gh_1h_2\cdots h_n) \in \psi.$$

As well, since $0 \in \operatorname{ran}(g)$, it is also in $\operatorname{ran}(gh_1h_2\cdots h_n)$, so $(\mathbf{0}, gh_1h_2\cdots h_n) \in \theta$. Because $\theta < \psi$, we get $(h, \mathbf{0}) \in \psi$ by transitivity of ψ. The same argument shows that $(h, \mathbf{0}) \in \psi$ for every $h \in D^I$ if $0 \in \operatorname{ran}(f)$ and $0 \notin \operatorname{ran}(g)$.

Now assume $0 \notin \operatorname{ran}(f)$ and $0 \notin \operatorname{ran}(g)$. Since $f \neq g$, there must be some $\alpha_0 \in I$ such that $f(\alpha_0) \neq g(\alpha_0)$. By our assumption about the structure of the unique minimal congruence of $\mathbf{A}$, we must have $(0, d) \in \mathrm{Cg}_{\mathbf{A}}(f(\alpha_0), g(\alpha_0))$ for every $d \in D$. By Proposition 3.0.2 (2), there is an $n+1$-ary term t and elements $a_1, \ldots, a_n \in A$ with $t(f(\alpha_0), a_1, \ldots, a_n) = 0$ and $t(g(\alpha_0), a_1, \ldots, a_n) \neq 0$, or vice versa. Without loss of generality, we will assume the former. We now define $h_1, \ldots, h_n \in A^I$ by

$$h_i(\alpha) = \begin{cases} a_i & \text{if } \alpha = \alpha_0 \\ g(\alpha) & \text{otherwise.} \end{cases}$$

Now consider $t(g, h_1, \ldots, h_n)$, which is an element of A^I. For $\alpha \neq \alpha_0$, we have

$$t(g, h_1, \ldots, h_n)(\alpha) = t(g(\alpha), g(\alpha), \ldots, g(\alpha)) = g(\alpha)$$

by idempotence, which is not 0 by assumption. Also,

$$t(g, h_1, \ldots, h_n)(\alpha_0) = t(g(\alpha_0), a_1, \ldots, a_n)$$

which is not 0 by construction. Therefore, $0 \notin \operatorname{ran}(t(g, h_1, \ldots, h_n))$. However,

$$t(f(\alpha_0), h_1(\alpha_0), \ldots, h_n(\alpha_0)) = t(f(\alpha_0), a_1, \ldots, a_n) = 0$$

by construction, so $0 \in \operatorname{ran}(t(f, h_1, \ldots, h_n))$. If we define $f' = t(f, h_1, \ldots, h_n)$ and $g' = t(g, h_1, \ldots, h_n)$, we have $(f', g') \in \psi$, $0 \in \operatorname{ran}(f')$, and $0 \notin \operatorname{ran}(g')$. By the first case, $(h, \mathbf{0}) \in \psi$ for every $h \in D^I$.

We have shown that for $\psi > \theta$, we have $(D^I \cup \{\mathbf{0}\})^2 \subseteq \psi$. Therefore, if we set $\Theta = \bigwedge_{\rho > \theta} \rho$, we get $(D^I \cup \{\mathbf{0}\})^2 \subseteq \Theta$. Choose $d \in D$ and let $\mathbf{d} \in A^I$ be the element with $\mathbf{d}(\alpha) = d$ for all $\alpha \in I$. We have $(\mathbf{d}, \mathbf{0}) \in \Theta$, but this pair is not in θ. Following the discussion at the top of page 13, we have shown that there is a unique congruence μ such that $\theta < \mu$ and $\theta < \eta$ implies $\mu \leq \eta$. Therefore, θ is completely meet irreducible so $\mathbf{A}^I/\theta$ is subdirectly irreducible. The size of this algebra is $(|A|-1)^{|I|} + 1 \geq 2^{|I|} + 1$ since $|A| > 2$. This means by taking I appropriately large, we can make $\mathbf{A}^I/\theta$ arbitrarily large. $\square$

We can apply Lemma 5.3.5 to show that type **5** simple algebras generate residually large varieties unless they are semilattices.

Theorem 5.3.6. *Suppose* $\mathbf{A}$ *is a finite simple* 2*-semilattice which is not a semilattice and* $\langle 0_A, 1_A \rangle$ *has type* **5**. *Then* $\mathbf{HSP}(\mathbf{A})$ *is residually large.*

Proof. By Theorem 4.3.2, $\mathbf{A}$ is not strongly connected, which means there is no strongly connected component X with $1_A \subseteq X^2 \cup 0_A$. It follows that Theorem 4.1.1 (1) holds for the minimal congruence 1_A, which means $\mathbf{A}$ has exactly two 1_A-nontrivial components. Of course, every component of A is 1_A-nontrivial, so $\mathbf{A}$ has exactly two strongly connected components. This means every $a \in A$ has $\lambda(a) \in \{0, 1\}$. The relation $(A')^2 \cup 0_A$ is a congruence on $\mathbf{A}$, so $|A'| = 1$ by simplicity. Since $\mathbf{A}$ is simple and not a semilattice, it must have at least 3 elements by Proposition 3.0.3. The result now holds by Lemma 5.3.5. $\square$

Theorem 5.3.7. *If* $\mathbf{A}$ *is a finite weakly acyclic* 2*-semilattice which is not a semilattice, then* $\mathbf{HSP}(\mathbf{A})$ *is residually large.*

Proof. By Birkhoff's subdirect representation theorem mentioned in the discussion after Definition 2.0.7, $\mathbf{A}$ embeds in $\prod_{\alpha \in S} \mathbf{A}/\alpha$ where S is the set of meet irreducible congruences on $\mathbf{A}$. A subdirect product of semilattices is a semilattice, so since we are assuming $\mathbf{A}$ is not a semilattice, it must be the case that $\mathbf{A}/\alpha$ is not a semilattice for some $\alpha \in S$. If the subvariety $\mathbf{HSP}(\mathbf{A}/\alpha)$ is residually large, so is $\mathbf{HSP}(\mathbf{A})$. Also, $\mathbf{A}/\alpha$ is weakly acyclic by Lemma 5.2.6, so it suffices to prove the result assuming $\mathbf{A}$ is subdirectly irreducible.

Since $\mathbf{A}$ is weakly acyclic, its smallest strongly connected component has just one element which we will call 0. By Lemma 5.3.4, $\mathbf{A}$ has a unique strongly connected component D with $\lambda(D) = 1$. If $D = \{1\}$ for some $1 \in A$, then $\{(0, 1), (1, 0)\} \cup 0_A$ is a congruence by Lemma 4.1.3. This congruence only collapses two elements, so it is minimal. Since $\mathbf{A}$ is subdirectly irreducible, it has a unique minimal congruence, which must be $\{(0,1),(1,0)\}\cup 0_A$. This means Lemma 5.3.5 applies, so we have shown that the result holds when $|D|=1$.

For the remainder of the proof, we assume, $|D|\geq 2$. The set $E=D\cup\{0\}$ is a subuniverse of $\mathbf{A}$, so we let $\mathbf{E}$ be the subalgebra with universe E and choose $\theta\in\mathrm{Con}(\mathbf{E})$ maximal. If $a\in 0/\theta$ for some $a\neq 0$, then $a\in D$ and for any $a\longrightarrow b$, we have $(ab,0b)=(b,0)\in\theta$. Since D is strongly connected, it follows that $D\subseteq 0/\theta$ and $\theta=1_E$, which contradicts the maximality of $\theta\in\mathrm{Con}(\mathbf{E})$. Therefore, $0/\theta=\{0\}$. We are assuming $\mathbf{A}$ is weakly acyclic, so $\mathbf{E}\leq\mathbf{A}$ is also weakly acyclic. This means D is not a subuniverse of $\mathbf{E}$ since D is a strongly connected subset of $\mathbf{E}$ with more than one element. Hence, D is not a θ-class. Therefore, $\mathbf{E}/\theta$ is a simple 2-semilattice with more than 2 elements, so it is not a semilattice by Proposition 3.0.3. Since $\mathbf{E}$ is weakly acyclic, so is $\mathbf{E}/\theta$ by Lemma 5.2.6, so $\mathbf{E}/\theta$ is not strongly connected since it has at least three elements. By Theorem 4.3.2, $\langle 0_E,1_E\rangle$ has type $\mathbf{5}$. By Theorem 5.3.6, $\mathbf{HSP}(\mathbf{E}/\theta)$ is residually large. Since $\mathbf{HSP}(\mathbf{E}/\theta)\leq\mathbf{HSP}(\mathbf{E})\leq\mathbf{HSP}(\mathbf{A})$, this implies $\mathbf{HSP}(\mathbf{A})$ is residually large. $\square$

The results of Maroti in [40] show that the variety generated by a finite tournament has only finitely many subdirectly irreducible members which are all finite. Theorem 5.3.5 is consistent with this because any weakly acyclic tournament is necessarily a semilattice. This can be seen by showing that if a tournament $\mathbf{T}$ has a failure of associativity, then it contains a three element subalgebra isomorphic to $\mathbf{T}_3$.

Recall from Definition 3.0.1 (4) that a commutative directoid is a 2-semilattice which additionally satisfies $x((xy)z)\approx(xy)z$. It was shown by Ježek and Quackenbush in [34] that the digraph of a commutative directoid is a partial order, and hence, is acyclic. Therefore, we have the following corollary.

Corollary 5.3.8. *Let* $\mathbf{D}$ *be a finite commutative directoid which is not a semilattice. Then* $\mathbf{HSP}(\mathbf{D})$ *is residually large.*

Proof. Since $\mathbf{D}$ is acyclic, it is weakly acyclic. The result follows from Theorem 5.3.7. $\square$

We note that Ježek and McNulty produced arbitrarily large finite examples of subdirectly irreducible directoids in [33], thereby showing that the variety of all commutative directoids is not residually very finite. Gardner and Parmenter in [27] produced arbitrarily large finite and a countably infinite subdirectly irreducible commutative directoid. Their construction can easily be modified to produce subdirectly irreducible commutative directoids of arbitrarily large cardinality. In other words, it follows from the work of Gardner and Parmenter that the variety of commutative directoids is residually large. It follows from Corollary 5.3.8 that any finitely generated variety of commutative directoids is either the variety of semilattices, and hence, is residually very finite, or is residually large.

We now state a corollary which is essentially a rephrasing of Theorem 5.3.7 in the language of tame-congruence-theoretic types.

Corollary 5.3.9. *Suppose* $\mathbf{A}$ *is a finite* 2*-semilattice and there exists a variety* $\mathcal{W}$ *which is not the variety of semilattices, is contained in* $\mathbf{HSP}(\mathbf{A})$*, and which omits type* $\mathbf{3}$*. Then* $\mathbf{HSP}(\mathbf{A})$ *is residually large.*

Proof. Suppose every finite member of $\mathcal{W}$ is a semilattice. Then the only finite subdirectly irreducible members of $\mathcal{W}$ are semilattices. The two element semilattice is the only subdirectly irreducible semilattice, so $\mathcal{W}$ has only finitely many finite subdirectly irreducible members up to isomorphism. By a Theorem of Quackenbush in [46], the two element semilattice is the only subdirectly irreducible member of $\mathcal{W}$, so $\mathcal{W}$ must be the variety of semilattices by Birkhoff's subdirect representation theorem. Since we are assuming this is not the case, we may choose a finite algebra $\mathbf{B} \in \mathcal{W}$ which is not a semilattice. Suppose $\mathbf{B}$ has a strongly connected subalgebra $\mathbf{C}$ with more than one element. For any maximal congruence θ on $\mathbf{C}$, the algebra $\mathbf{C}/\theta$ is strongly connected by Lemma 3.1.5 (8), so $\langle \theta, 1_C \rangle$

has type **3** by Theorem 4.3.2. By our assumption, no such congruence quotient can exist, so $\mathbf{B}$ is weakly acyclic. By Theorem 5.3.7, $\mathbf{HSP}(\mathbf{B})$ is residually large, so $\mathbf{HSP}(\mathbf{A})$ is residually large. $\square$

For a finite 2-semilattice $\mathbf{A}$, the problem of determining whether or not it is weakly acyclic is decidable. A possible algorithm is to find all strongly connected subsets of A and check if any of them are subuniverses. This together with Theorem 5.3.7 gives a decidable sufficient condition for $\mathbf{HSP}(\mathbf{A})$ to be residually large. There are still plenty of open questions about decidability of residual character in the case of 2-semilattices. Some of them are asked below.

Question 5.3.10.

1. Is there an algorithm that determines, given a finite 2-semilattice $\mathbf{A}$, whether or not $\mathbf{HSP}(\mathbf{A})$ is residually very finite, finite, small, or large?
2. Is the converse of Corollary 5.3.9 true?
3. For a finite 2-semilattice $\mathbf{A}$, is it decidable whether or not $\mathbf{HSP}(\mathbf{A})$ satisfies the conditions of Corollary 5.3.9?
4. For a fixed finite 2-semilattice, is it decidable which parts of Definition 5.1.1 $\mathbf{HSP}(\mathbf{A})$ satisfies?
5. Is it possible for a variety of 2-semilattices to have arbitrarily large subdirectly irreducible members but a cardinal bound on the size of its simple members?

Chapter 6

The Constraint Satisfaction Problem

We now shift our focus away from purely algebraic properties of 2-semilattices. For the rest of this thesis, we focus on the constraint satisfaction problem in the context of 2-semilattices. The constraint satisfaction problem, or CSP, is polynomially equivalent to the homomorphism problem for relational structures which asks, for finite relational structures $\mathbb{B}$ and $\mathbb{D}$ of the same type if there is a homomorphism from $\mathbb{B}$ to $\mathbb{D}$. In Section 6.1, we provide some background, history, and basic definitions on the subject of constraint satisfaction problems. The rest of Chapter 6 is devoted to generalizing a result of Bulatov, Theorem 3.1 from [15]. It was mentioned in the introduction that if verified, the results claimed by Zhuk in [53], Rafiey et. al. in [48], and Bulatov in [10] will imply Theorem 6.6.5 and Corollary 6.6.9, the main results in this chapter.

6.1 Background

In this section, we introduce the constraint satisfaction problem and some notation.

Definition 6.1.1.

1. An *instance* is a triple $\mathcal{I} = (X, D, \mathcal{C})$ where X is a finite set of variables, D is a finite, nonempty set, and $\mathcal{C}$ is a finite set of constraints. A *constraint*, $C \in \mathcal{C}$ is a pair $(\mathbf{x}, R)$ where $\mathbf{x} \in X^n$ and $R \subseteq D^n$ for some n.

2. A *solution* to an instance is a function $\varphi : X \to D$ with $(\varphi(x_1), \ldots, \varphi(x_n)) \in R$ for every $(\mathbf{x}, R) \in \mathcal{C}$.

3. The *constraint satisfaction problem*, abbreviated CSP, is the decision problem whose input is an instance $\mathcal{I}$ and output is YES if a solution exists, and NO otherwise.

Given an instance $\mathcal{I}$ and a function $\varphi : X \to D$, it can be checked in time linear in the size of $\mathcal{I}$ whether or not φ is a solution. Therefore the CSP as stated above is in **NP**. Given any graph $G = (V, E)$, let $\mathcal{I}_G$ be the instance with variable set V, $D = \{0, 1, 2\}$, and for each $(v, w) \in E$, a constraint $((v, w), R_{\neq})$ where $R_{\neq} \subseteq D^2$ is $\{(a, b) : a \neq b\}$. The solutions to this instance are precisely the homomorphisms from G to the graph $(D, R_{\neq})$. Therefore, $\mathcal{I}$ has a solution exactly when G is 3-colourable. Since 3-colourability is **NP**-complete, it follows that the CSP is **NP**-complete. The complexity of the CSP is more interesting for versions of CSP where the allowed instances are restricted.

Definition 6.1.2. A *relational structure* is a pair $\mathbb{D} = (D, \mathcal{R})$ where D is a nonempty set and $\mathcal{R}$ is a set of relations on D. We say $\mathbb{D}$ is *finite* if D is finite, and has *finite type* if $\mathcal{R}$ is finite.

With Definition 6.1.2 in mind, we can define a version of CSP which is less general than the one in Definition 6.1.1.

Definition 6.1.3. Fix a finite relational structure $\mathbb{D} = (D, \mathcal{R})$ of finite type. The decision problem CSP($\mathbb{D}$) is the same as the general CSP, except the input is restricted to instances

$\mathcal{I} = (X, D, \mathcal{C})$ where D is the domain of $\mathbb{D}$, and for each constraint $(\mathbf{x}, R) \in \mathcal{C}$, the relation R is in $\mathcal{R}$.

The constraint satisfaction problem as stated in Definition 6.1.3 is polynomially equivalent to the homomorphism problem, $\mathrm{Hom}(\mathbb{D})$. The inputs of $\mathrm{Hom}(\mathbb{D})$ are finite structures $\mathbb{E}$ similar to $\mathbb{D}$, and $\mathrm{Hom}(\mathbb{D})$ asks whether or not there is a homomorphism from $\mathbb{E}$ to $\mathbb{D}$. For an explanation of this equivalence, see Section 1 of Larose and Tesson [36].

Schaeffer proved in [51] that if $\mathbb{D} = (D, \mathcal{R})$ has $|D| = 2$ and $\mathcal{R}$ finite, then $\mathrm{CSP}(\mathbb{D})$ is either in **P** or is **NP**-complete. Hell and Nešetřil in [28] showed that the same dichotomy holds when $\mathbb{D}$ is a simple graph. By "$\mathbb{D}$ is a simple graph", we mean $\mathcal{R}$ contains exactly one relation which is binary, symmetric, and irreflexive. Based largely on these two known results, Feder and Vardi in [22] conjectured the following about the version of CSP in Definition 6.1.3: For a finite relational structure $\mathbb{D}$ of finite type, either $\mathrm{CSP}(\mathbb{D})$ is in **P** or it is **NP**-complete. This conjecture has come to be known as the Dichotomy Conjecture.

Many partial confirmations of the Dichotomy Conjecture have been attained using algebraic techniques. The framework which allows the use of algebraic techniques was laid out in the late 1990s and early 2000s by Jeavons and a variety of coauthors in several papers including [31], [32], [13], and [14]. Jeavons et. al. introduced the notion of a "tractable" algebra or class of algebras, which we define in Definition 6.1.5. Roughly speaking, an algebra $\mathrm{Alg}(\mathbb{D})$ is associated to every relational structure $\mathbb{D}$, and $\mathrm{Alg}(\mathbb{D})$ is "tractable" if and only if $\mathrm{CSP}(\mathbb{D})$ is in **P**.

Definition 6.1.4. Fix a finite algebra $\mathbf{D}$, and a positive integer n. Define $\mathcal{R}_n(\mathbf{D}) = \{A : A \leq \mathbf{D}^m \text{ for some } m \leq n\}$. The decision problem $\mathrm{CSP}(\mathbf{D}, n)$ is the decision problem $\mathrm{CSP}(\mathbb{D})$ where $\mathbb{D} = (D, \mathcal{R}_n)$.

Using Definition 6.1.4, we now precisely define what it means for an algebra or class of algebras to be tractable.

Definition 6.1.5. A finite algebra $\mathbf{D}$ is *tractable* if for every $n \geq 2$, there is a polynomial time algorithm which solves $\mathrm{CSP}(\mathbf{D}, n)$. A class of algebras is called *tractable* if each of its finite members is tractable.

As mentioned before, to each finite relational structure of finite type $\mathbb{D}$, Jeavons et. al. associated a finite algebra $\mathrm{Alg}(\mathbb{D})$ with the property that $\mathrm{CSP}(\mathbb{D})$ is solvable in polynomial time if and only if $\mathrm{Alg}(\mathbb{D})$ is tractable. Furthermore, they showed that one can assume $\mathrm{Alg}(\mathbb{D})$ is idempotent.

The "Algebraic Dichotomy Conjecture" appears in several forms in [13] and [14]. Conjecture 4.15 of Bulatov and Jeavons in [13] is equivalent to the following: For a finite relational structure $\mathbb{D}$ of finite type, if $\mathrm{Alg}(\mathbb{D})$ has a Taylor term operation, then $\mathrm{Alg}(\mathbb{D})$ is tractable; otherwise, $\mathrm{CSP}(\mathbb{D})$ is **NP**-complete. Earlier in [13], Bulatov and Jeavons prove that if $\mathrm{Alg}(\mathbb{D})$ does not have a Taylor operation, then $\mathrm{CSP}(\mathbb{D})$ is **NP**-complete. Therefore, the content of their conjecture is that if $\mathrm{Alg}(\mathbb{D})$ has a Taylor operation, then $\mathrm{Alg}(\mathbb{D})$ is tractable. In particular, their conjecture implies $\mathrm{CSP}(\mathbb{D})$ is solvable in polynomial time when $\mathrm{Alg}(\mathbb{D})$ has a Taylor operation. The Algebraic Dichotomy Conjecture implies the Dichotomy Conjecture of Feder and Vardi. Because of pre-existing knowledge of universal algebra, the algebraic approach to the constraint satisfaction problem has been quite successful.

Since the Algebraic Dichotomy Conjecture was posed, there have been many partial confirmations. The "few subpowers algorithm", which is a generalization of Gaussian elimination, is a polynomial time algorithm which solves $\mathrm{CSP}(\mathbb{D})$ when $\mathrm{Alg}(\mathbb{D})$ has an operation satisfying a particular collection of identities. An operation satisfying such a collection of identities is called an "edge operation" and will be defined precisely in Definition 6.6.7. That the few subpowers algorithm solves $\mathrm{CSP}(\mathbb{D})$ when $\mathrm{Alg}(\mathbb{D})$ has an edge operation was proved gradually by various authors in [12], [18], [30], and [6]. Another major partial confirmation was established by Barto and Kozik in [3], and independently in an unpublished work of Bulatov [9]. The authors showed that when Alg($\mathbb{D}$) generates a meet semidistributive variety, CSP($\mathbb{D}$) is solvable by "local consistency checking" which will be discussed in more depth in Section 6.2. The local consistency checking algorithm finds, if it exists, a specific set of partial solutions to an instance of CSP($\mathbb{D}$). It is always the case that if this set of partial solutions does not exist, then the instance has no solution. Barto and Kozik showed that when Alg($\mathbb{D}$) generates a meet semidistributive variety and such a set of partial solutions exists, the instance has a solution. This showed that the local consistency checking algorithm, which runs in polynomial time, can be used to answer CSP($\mathbb{D}$) when Alg($\mathbb{D}$) generates a meet semidistributive variety. It had already been shown in [37] by Larose and Zádori that if CSP($\mathbb{D}$) could be solved in this way by local consistency checking, then Alg($\mathbb{D}$) generates a congruence meet semi-distributive variety. Larose and Zádori conjectured the converse in the same paper. Their conjecture was known as the "bounded width conjecture" until it was verified. An important step leading to the proof of the bounded width conjecture was the following result of Bulatov in [15]: If Alg($\mathbb{D}$) has a 2-semilattice operation as a basic operation, then CSP($\mathbb{D}$) is solvable by local consistency checking. Since an algebra with a 2-semilattice operation generates a meet semidistributive variety (see Proposition 3.2.3), Bulatov's result verified a special case of the bounded width conjecture.

The main result in this chapter, Theorem 6.6.5, is a partial confirmation of the Algebraic Dichotomy Conjecture which is not implied by either of the two major confirmations mentioned in the previous paragraph. Roughly speaking, we are able to prove that CSP($\mathbb{D}$) is solvable in polynomial time in the following situation: Alg($\mathbb{D}$) has a congruence θ, and a binary operation $*$ so that (i) $*$ is a 2-semilattice operation on Alg($\mathbb{D}$)$/\theta$, (ii) each θ-class considered as a subalgebra of Alg($\mathbb{D}$) is tractable, and (iii), $x * y$ is the first projection on each θ-class. The result will be stated more precisely in Section 6.6. The result obtained is similar to those of Bergman and Failing in [5] and Maróti in [41].

Until recently, confirming the Algebraic Dichotomy Conjecture assuming (i) and (ii) above, even when (i) is strengthened to "$*$ is a semilattice operation" was thought to be out of reach. However, Bulatov circulated a proof of precisely this in January of 2017 in [11]. As of August, 2017, it remains unpublished. In any case, neither result implies the other since, while Bulatov was able to eliminate condition (iii), he needed to strengthen conditions (i) and (ii).

Since November 2016, three proofs of the Algebraic Dichotomy Conjecture have been circulated. The first, by Feder, Rafiey, and Kinne in [48] was circulated in November 2016. The second proof was circulated by Zhuk [53] in March, 2017. Finally, Bulatov [10] circulated a proof in April of 2017. As of May 2017, all three proofs remain unconfirmed.

6.2 The $(2, 3)$-Consistency Algorithm

Working with large arity relations can be cumbersome, so we state a proposition that shows we can focus on binary relations. A proof can be found in Barto and Kozik [3].

Proposition 6.2.1. *Let* $\mathbf{D}$ *be a finite idempotent algebra and* $n \geq 2$ *be an integer. Then* $\mathrm{CSP}(\mathbf{D}, n)$ *is polynomial time equivalent to* $\mathrm{CSP}(\mathbf{D}^{\lceil \frac{n}{2} \rceil}, 2)$.

Because of Proposition 6.2.1, we will only consider instances of $\mathrm{CSP}(\mathbf{D}, 2)$. Before proceeding, we introduce the notation R^{-1} to denote the binary relation $\{(b, a) : (a, b) \in R\}$ for a binary relation R.

Definition 6.2.2. Fix a finite, idempotent algebra $\mathbf{D}$. An instance $(X, D, \mathcal{C})$ of $\mathrm{CSP}(\mathbf{D}, 2)$ is called *standard* if $\mathcal{C} = \{(x, P_x) : x \in X\} \cup \{((x, y), R_{x,y}) : (x, y) \in X \times X\}$ for some subuniverses, P_x and $R_{x,y}$ of D and $D \times D$ respectively, and $R_{x,y} \subseteq P_x \times P_y$ for every

$(x, y) \in X \times X$. We will henceforth refer to a standard instance as $\mathcal{I} = (X, \mathcal{P}, \mathcal{R})$ where, $\mathcal{P} = (P_x : x \in X)$ and $\mathcal{R} = (R_{x,y} : (x, y) \in X^2)$. That is, we identify it by its indexed sets $\mathcal{P}$ of *potatoes* and $\mathcal{R}$ of *relations*. We call $\mathcal{I}$ a *standard* $(2, 3)$*-instance* of $\mathrm{CSP}(\mathbf{D}, 2)$ if it also satisfies the following four conditions.

(P1) For each $x \in X$, $R_{x,x} = 0_{P_x}$ (the diagonal in $P_x \times P_x$),

(P2) For $x, y, z \in X$ and any $(a, b) \in R_{x,y}$, there is a $c \in P_z$ such that $(a, c) \in R_{x,z}$ and $(b, c) \in R_{y,z}$,

(P3) For each $x, y \in X$, $R_{x,y} \leq_{\mathrm{sd}} \mathbf{P}_x \times \mathbf{P}_y$ if P_x and P_y are both non-empty.

(P4) $R_{y,x} = R_{x,y}^{-1}$ for each $x, y \in X$.

We note that the set of solutions to a standard $(2, 3)$-instance, if nonempty, can be identified as a subalgebra of the product $\prod_{x \in X} \mathbf{P}_x$. For this reason, when $\mathcal{I}$ has at least one solution, we will sometimes refer to its algebra of solutions and use the fact that it is in the variety generated by $\mathbf{D}$. It is not hard to see that for a standard $(2, 3)$-instance, there is an empty potato if and only if all potatoes, and hence, relations, are empty. We call such an instance *empty*. It is worth noting that (P3) and (P4) follow from (P1) and (P2). To see this, suppose a standard instance $\mathcal{I}$ satisfies (P1) and (P2) and let $x, y \in X$. If $(a, b) \in R_{xy}$, then by (P2) there is $c \in P_x$ with $(a, c) \in R_{xx}$ and $(b, c) \in R_{yx}$. By (P1), $c = a$, so $(b, a) \in R_{yx}$ which shows $R_{xy}^{-1} \subseteq R_{yx}$. If $(b, a) \in R_{yx}$ then $(a, b) \in R_{yx}^{-1}$, so $(a, b) \in R_{xy}$ using the same argument, so $(b, a) \in R_{xy}^{-1}$. This shows that $\mathcal{I}$ satisfies (P4). For any $a \in P_x$, the pair $(a, a) \in R_{xx}$ by (P1), and by (P2) there is a $b \in P_y$ with $(a, b) \in R_{xy}$. This shows that the first projection from R_{xy} to P_x is surjective. Using (P4), we get that the second projection is surjective as well. Thus, $\mathcal{I}$ satisfies (P3) and is a standard $(2, 3)$-instance. We have shown that to prove a standard instance is a standard $(2, 3)$-instance, we need only verify (P1) and (P2) from Definition 6.2.2. The intuition behind standard $(2, 3)$-instances is that "obviously" redundant elements and pairs are removed. For example, if $\mathcal{I}$ is a standard instance, $x, y \in X$ and $a \in P_x$, but there is no $b \in P_y$ with $(a, b) \in R_{xy}$, then no solution to $\mathcal{I}$ can possibly assign a to x. In this case, there is no reason to include $a \in P_x$. We leave (P3) and (P4) in the definition because they are part of this intuition behind standard $(2, 3)$-instances.

We will occasionally refer to several standard instances of $\mathrm{CSP}(\mathbf{D}, 2)$ in the same context. In this case, we will use the instance name in a superscript to denote its potatoes and relations. That is, if $\mathcal{I}$ and $\mathcal{J}$ are instances with a common set of variables X, we will denote the potatoes of $\mathcal{I}$ by $\mathcal{P}^{\mathcal{I}} = (\mathbf{P}_x^{\mathcal{I}} : x \in X)$ and $\mathcal{P}^{\mathcal{J}} = (\mathbf{P}_x^{\mathcal{J}} : x \in X)$, and similarly for the sets $\mathcal{R}^{\mathcal{I}}$ and $\mathcal{R}^{\mathcal{J}}$ of relations. This will usually arise in the situation when $\mathcal{J}$ is a *subinstance* of $\mathcal{I}$, defined below.

Definition 6.2.3. Let $\mathcal{I}$ and $\mathcal{J}$ be standard instances of $\mathrm{CSP}(\mathbf{D}, 2)$. We say that $\mathcal{J}$ is a *subinstance* of $\mathcal{I}$ and write $\mathcal{J} \leq \mathcal{I}$ if $\mathcal{I}$ and $\mathcal{J}$ have a common variable set X, $\mathbf{P}_x^{\mathcal{J}} \leq \mathbf{P}_x^{\mathcal{I}}$ for each $x \in X$, and $\mathbf{R}_{xy}^{\mathcal{J}} \leq \mathbf{R}_{xy}^{\mathcal{I}}$ for each $(x, y) \in X^2$.

Algorithm 1 (found on the next page) is the $(2, 3)$-consistency checking algorithm, presented to match the notation in Definition 6.2.2. It is a form of local consistency checking. The input is an instance of $\mathrm{CSP}(\mathbf{D}, 2)$. The algorithm converts the instance to a standard instance, and then removes elements from potatoes and pairs from relations until it is a standard $(2, 3)$-instance in such a way that no solutions are lost. Variations of the algorithm appear frequently in the literature. For example, see Barto Kozik [3], Bulatov [15], Dechter [20], Feder and Vardi [22], or Larose Zádori [37].

Proposition 6.2.4. *The output of Algorithm 1 is a standard* $(2, 3)$*-instance of* $\mathrm{CSP}(\mathbf{D})$ *with exactly the same set of solutions as the input instance.*

Proof. We first prove the following claim:

Algorithm 1 $(2,3)$-consistency checking

```
Input: An instance, $\mathcal{I} = (X, D, \mathcal{C})$ of CSP($\mathbf{D}$, 2)
for $(x,y) \in X^2$ do
    if $x \neq y$ then
        $R_{xy} \leftarrow D \times D$
    else
        $R_{xy} \leftarrow 0_D$
for $((x,y),R) \in \mathcal{C}$ do
    $R_{xy} \leftarrow R_{xy} \cap R$
    $R_{yx} \leftarrow R_{yx} \cap R^{-1}$
flag $\leftarrow 1$
while flag $= 1$ do
    flag $\leftarrow 0$
    for $(x,y) \in X^2$ do
        for $z \in X$ do
            for $(a,b) \in R_{xy}$ do
                if there is no $c \in D$ with $(a,c) \in R_{xz}$ and $(b,c) \in R_{yz}$ then
                    $R_{xy} \leftarrow R_{xy} \setminus \{(a,b)\}$
                    $R_{yx} \leftarrow R_{yx} \setminus \{(b,a)\}$
                    flag $\leftarrow 1$
for $x \in X$ do
    $P_x = \mathrm{pr}_1(R_{xx})$
Output: $(X, (P_x : x \in X), (R_{xy} : (x,y) \in X^2))$
```

Claim. Suppose $R_{xy} \leq \mathbf{D} \times \mathbf{D}$ for every $(x, y) \in X^2$ when the algorithm reaches the while loop beginning on line 11. Then $R_{xy} \leq \mathbf{D} \times \mathbf{D}$ for every $(x, y) \in X^2$ when the algorithm exits the while loop.

Proof of Claim. The net effect of the algorithm from line 14 to 19 is to replace, for a fixed $(x, y) \in X^2$, R_{xy} by

$$\{(a, b) \in R_{xy} : \text{ there exists } c \in D \text{ with } (a, c) \in R_{xz} \text{ and } (b, c) \in R_{yz}\}.$$

This set is a subuniverse because R_{xy}, R_{xz}, and R_{yz} are all subuniverses. This is the only part of the while loop where any changes are made, so $R_{xy} \leq \mathbf{D} \times \mathbf{D}$ when the algorithm exits the while loop. $\square$

By construction in lines 2-6 and 21, there is exactly one potato for each variable $x \in X$ and exactly one relation for each ordered pair $(x, y) \in X^2$. We need to verify that $P_x \leq \mathbf{D}$ and $R_{xy} \leq \mathbf{P}_x \times \mathbf{P}_y$ for each $x, y \in X$. We also need to verify that the output satisfies (P1)-(P4) from Definition 6.2.2. By the remark after Definition 6.2.2, we only need to verify (P1) and (P2). We will prove that the solution set is unaffected by the algorithm at the end of the proof.

When R_{xy} is initialized, it is either $D \times D$ or 0_D. Either way, it is a subuniverse of $\mathbf{D} \times \mathbf{D}$. In the for loop spanning lines 7-9, R_{xy} is possibly replaced by its intersection with other subuniverses of $\mathbf{D} \times \mathbf{D}$, so each $R_{xy} \leq \mathbf{D} \times \mathbf{D}$ when the algorithm enters the while loop. By the claim, $R_{xy} \leq \mathbf{D} \times \mathbf{D}$ when the algorithm terminates. It follows from this fact that $\mathbf{P}_x = \mathrm{pr}_1(R_{xx}) \leq \mathbf{D}$. It remains to verify that the output satisfies (P1) and (P2) from Definition 6.2.2.

To verify (P1), we note that $R_{xx} = 0_D$ when it is initialized, so $R_{xx} \leq 0_D$ when the algorithm reaches line 20. This means every pair in R_{xx} is of the form (a, a) for some $a \in D$. Since $P_x = \mathrm{pr}_1(R_{xx})$, we have that $R_{xx} = \{(a, a) : a \in P_x\} = 0_{P_x}$.

To verify (P2), we suppose $(a, b) \in R_{xy}$ and $z \in X$. There must be some $c \in D$ such that $(a, c) \in R_{xz}$ and $(b, c) \in R_{yz}$. Otherwise, (a, b) would have been removed from R_{xy} in the loop that runs from lines 15 to 19. It suffices to show that $c \in P_z$, and since $R_{zz} = 0_{P_z}$, it suffices to show that $(c, c) \in R_{zz}$. To see this, note that $(a, c) \in R_{xz}$, so there is some $d \in D$ with $(a, d) \in R_{xz}$ and $(c, d) \in R_{zz}$ by the same reasoning as above. This means $c = d$ so $(c, c) \in R_{zz}$.

Finally, we argue that the solution sets of the input and output are identical. A function $\varphi : X \to D$ is a solution to $\mathcal{I}$ if and only if $(\varphi(x), \varphi(y)) \in R$ whenever $((x, y), R) \in \mathcal{C}$. This is equivalent to $(x, y) \in R_{xy}$ for all $x, y \in X$ immediately after line 9. The solution set will not change after the introduction of the P_x, so the only place the solution set can change is during the while loop. Removing a pair from R_{xy} can not possibly introduce a new solution, and a pair is removed precisely when it is not a solution to the instance restricted to some subset of X of size at most 3. If a pair does not extend to a solution to such a restricted instance, it can not possibly be a solution to the full instance. $\square$

To finish this section, we explain why Algorithm 1 runs in Polynomial time. We first note that since $\mathbf{D}$ is fixed, we take $|D|$, and hence $|D \times D|$ and $|\{R : R \leq \mathbf{D} \times \mathbf{D}\}| \leq 2^{|D|}$ as constants. The input is a variable set X and some constraints of the form (x, P) and $((x, y), R)$ such that $P \leq \mathbf{D}$ and $R \leq \mathbf{D}^2$. Therefore, the size of the input can be taken as $n = |X| + |\mathcal{C}|$. The loop that spans lines 2 through 6 makes an assignment for each $(x, y) \in X^2$. Thus, it takes no more than n^2 steps. Similarly, the loop spanning lines 20-21 runs in at most n steps. The loop that spans lines 7 through 9 makes one assignment for each constraint, and so it takes at most n steps. Each pass through the while loop considers a triple $(x, y, z) \in X^3$ and scans some subuniverses of $\mathbf{D}^2$. Therefore, each pass through the while loop takes at most cn^3 steps where c is a constant. If a pass through the while loop makes no changes, then the loop terminates. Otherwise, one of the relations must shrink. Therefore, the number of times the algorithm enters the while loop is bounded by a constant times the number of relations which can shrink, which is $|X| + |X^2| \leq n + n^2$. We conclude that there is a polynomial in n of degree 5 which bounds the number step to run the while loop. Putting all of this together, there is a polynomial of degree at most 5 that bounds the number of steps Algorithm 1 takes.

6.3 Bulatov's Algorithm for CSP over 2-Semilattices

In this section, we go through the proof of Theorem 3.1 from [15], in which Bulatov shows that the local consistency checking algorithm correctly answers $\mathrm{CSP}(\mathbf{D}, n)$ for all n in the case that $\mathbf{D}$ is a 2-semilattice. Because of Proposition 6.2.1, we will go through the proof with the added assumption that $n = 2$. The first Lemma is one of Bulatov's observations translated to our context.

Lemma 6.3.1. *Let* $\mathbf{D} \in \mathcal{S}$ *and suppose* $\mathcal{I} = (X, \mathcal{P}, \mathcal{R})$ *is a nonempty standard* $(2,3)$-*instance of* $\mathrm{CSP}(\mathbf{D}, 2)$. *The instance,* $\mathcal{I}' = (X, \{P'_x : x \in X\}, \{R'_{xy} : (x, y) \in X^2\})$*, is a nonempty standard* $(2,3)$-*instance.*

Lemma 6.3.1 is essentially Proposition 3.2 from [15]. The proof given by Bulatov is rather terse, so we include a detailed proof of the result.

Proof. Since P_x and R_{xy} are nonempty for each $x, y \in X$, we also have that P'_x and R'_{xy} are nonempty for all $x, y \in X$. By Lemma 3.4.3 (2), $R'_{xy} \leq_{\mathrm{sd}} \mathbf{P}'_x \times \mathbf{P}'_y$, so certainly $R'_{xy} \subseteq P'_x \times P'_y$.

To see that $\mathcal{I}'$ satisfies (P1), Observe that (P1) for $\mathcal{I}$ implies that $\mathbf{P}_x$ and $\mathbf{R}_{xx}$ are isomorphic via the map given by $a \mapsto (a, a)$ for $a \in P_x$. This means R'_{xx} is precisely

$\{(a,a) : a \in P'_x\} = 0_{P'_x}$. We now show that $\mathcal{I}'$ satisfies (P2). First, for fixed $(x,y,z) \in X^3$ we define the algebra of triangles on (x,y,z) by

$$T = \{(a,b,c) \in P_x \times P_y \times P_z : (a,b) \in R_{xy}, (a,c) \in R_{xz}, (b,c) \in R_{yz}\}$$

and show that its subuniverse,

$$S = \{(a,b,c) \in P'_x \times P'_y \times P'_z : (a,b) \in R'_{xy}, (a,c) \in R'_{xz}, (b,c) \in R'_{yz}\},$$

is nonempty. If we take $(a_1,b_1) \in R'_{xy}$, by (P2) for $\mathcal{I}$, there is $c_1 \in P_z$ so that $(a_1,b_1,c_1) \in T$. Similarly, there is $(a_2,b_2,c_2) \in T$ with $(a_2,c_2) \in R'_{xz}$, and $(a_3,b_3,c_3) \in T$ with $(b_3,c_3) \in R'_{yz}$. If we set $a = a_1a_2a_3$, $b = b_1b_2b_3$, and $c = c_1c_2c_3$, then an application of Lemma 3.1.5 (4) shows that (a,b,c) is in S, so S is nonempty. Now we take $(d,e) \in R'_{xy}$ and find $f \in P'_z$ such that $(d,e,f) \in S$. Again, we can use (P2) of $\mathcal{I}$ to find $f' \in P_z$ so that $(d,e,f') \in T$. Since $(d,e) \in R'_{xy}$, Lemma 3.1.5 (3) guarantees a walk in R_{xy} from (a,b) to (d,e). Using this and (P2) of $\mathcal{I}$, we can find a sequence, $\{(u_i,v_i,w_i)\}_{i=1}^n$ of elements of T so that

$$(a,b) \xrightarrow{\mathbf{R}_{xy}} (u_1,v_1) \xrightarrow{\mathbf{R}_{xy}} \cdots \xrightarrow{\mathbf{R}_{xy}} (u_n,v_n) \xrightarrow{\mathbf{R}_{xy}} (d,e).$$

Now define $w_1^\times = cw_1$, and $w_i^\times = w_{i-1}^\times w_i$ for $2 \le i \le n$. As well, set $f = w_n^\times \cdot f'$. It follows from Lemma 3.1.5 (1) that

$$(a,b,c) \xrightarrow{\mathbf{T}} (u_1,v_1,w_1^\times) \xrightarrow{\mathbf{T}} \cdots \xrightarrow{\mathbf{T}} (u_n,v_n,w_n^\times) \xrightarrow{\mathbf{T}} (d,e,f).$$

Since each of R'_{xy}, R'_{xz}, and R'_{yz} is closed with respect to taking out neighbours, it follows that S is as well. Therefore, since $(a,b,c) \in S$, the entire walk, including (d,e,f), is in S. □

Lemma 6.3.1 allows us to reduce a nonempty standard $(2,3)$-instance to a smaller nonempty standard $(2,3)$-instance in the case that one of its potatoes or relations is not strongly connected. We will next show that when a nonempty standard $(2,3)$-instance has only strongly connected potatoes and relations, we can again reduce the instance to a smaller nonempty standard $(2,3)$-instance as long as some potato has at least two elements. By alternating these reductions, a nonempty standard $(2,3)$-instance can be reduced to a nonempty standard $(2,3)$-instance in which every potato is a singleton and any solution to the reduced standard $(2,3)$-instance is a solution to the original. We will show that a nonempty standard $(2,3)$-instance in which every potato is a singleton must have a solution. This is roughly how Bulatov's proof of Theorem 6.3.6 goes. In fact, Barto and Kozik's proof of the bounded width conjecture follows the same process, but with "strongly connected" replaced by "absorption free". Definition 6.3.2 is very similar to Barto and Kozik's Definition 8.2 of a *decomposition* from [3].

Definition 6.3.2. Let $\mathbf{D} \in \mathcal{S}$ and $\mathcal{I}$ be a standard $(2,3)$-instance of $\mathrm{CSP}(\mathbf{D}, 2)$ in which every potato and relation is strongly connected. Further suppose there is some $u \in X$ with $|P_u| > 1$. Choose a maximal congruence, α_u of $\mathbf{P}_u$. Let $W \subseteq X$ be the set of variables such that $\{(a, b/\alpha_u) : (a,b) \in \mathbf{R}_{xu}\}$ is the graph of a surjective homomorphism from $\mathbf{P}_x$ to $\mathbf{P}_u/\alpha_u$, and for each $x \in W$, let φ_{xu} be this surjection. Now let $P_u^1, P_u^2, \ldots, P_u^k$ be the α_u classes in $\mathbf{P}_u$, and define instances $\mathcal{I}_1, \ldots, \mathcal{I}_k$, each with variable set X by

- $P_x^{\mathcal{I}_i} = \varphi_{xu}^{-1}(P_u^i)$ if $x \in W$, and P_x, otherwise.
- $R_{xy}^{\mathcal{I}_i} = R_{xy} \cap (P_x^{\mathcal{I}_i} \times P_y^{\mathcal{I}_i})$

Since $R_{uu} = 0_{P_u}$, we have that $u \in W$ and $P_u^{\mathcal{I}_i} = P_u^i$. We now state a Lemma based on Definition 6.3.2.

Lemma 6.3.3. *Let $\mathcal{I}$ be a standard $(2,3)$-instance of $\mathrm{CSP}(\mathbf{D}, 2)$, $u \in X$ be such that $|P_u| > 1$, $\alpha_u \in \mathrm{Con}(\mathbf{P}_u)$ be maximal, and W and $\mathcal{I}_1, \ldots, \mathcal{I}_k$ be as in Definition 6.3.2.*

1. *If $x \in W$, then $R_{xy} \cap (P_x^{\mathcal{I}_i} \times P_y) = R_{xy} \cap (P_x^{\mathcal{I}_i} \times P_y^{\mathcal{I}_i}) = R_{xy}^{\mathcal{I}_i}$ for every $y \in X$.*

2. *If s is a solution to $\mathcal{I}$, then it is a solution to $\mathcal{I}_i$ for some i.*

3. *For $x \in W$, the sets $P_x^{\mathcal{I}_i}$ are the classes of a congruence $\alpha_x \in \mathrm{Con}(\mathbf{P}_x)$ satisfying $\mathbf{P}_x/\alpha_x \cong \mathbf{P}_u/\alpha_u$ via the isomorphism $P_x^{\mathcal{I}_i} \mapsto P_u^{\mathcal{I}_i}$.*

Proof. For (1), the second equality is by definition, so we are only concerned with the first. Of the two inclusions needed to verify the equality, the left to right inclusion is the interesting one. To see that it is true, we first note that if $y \notin W$, then $P_y^{\mathcal{I}_i} = P_y$, so there is nothing to prove. From now on, we assume that $y \in W$ as well. Suppose $(a, b) \in R_{xy} \cap (P_x^{\mathcal{I}_i} \times P_y)$. By (P2) of $\mathcal{I}$, we get some $c \in P_u$ such that $(a, c) \in R_{xu}$ and $(b, c) \in R_{yu}$. Since $x \in W$ and $a \in P_x^{\mathcal{I}_i}$, we have that $c \in P_u^{\mathcal{I}_i}$. This is because $(a, c) \in R_{xu}$ means $c \in \varphi_{xu}(a) = P_u^{\mathcal{I}_i}$. Since $y \in W$ and $(b, c) \in R_{yu}$, the same reasoning in reverse puts $b \in \varphi_{yu}^{-1}(P_u^{\mathcal{I}_i}) = P_y^{\mathcal{I}_i}$.

To prove (2), it suffices to show that there is some i such that $s(x) \in P_x^{\mathcal{I}_i}$ for all $x \in X$ because of the way the $R_{xy}^{\mathcal{I}_i}$ are defined,. Since the $P_u^{\mathcal{I}_1}, \ldots, P_u^{\mathcal{I}_k}$ partition P_u, there is some i such that $s(u) \in P_u^{\mathcal{I}_i}$. For $x \in W$, since $(s(x), s(u)) \in R_{xu}$, we have that $s(x) \in \varphi_{xu}^{-1}(P_u^{\mathcal{I}_i}) = P_x^{\mathcal{I}_i}$, and for $x \notin W$, $s(x) \in P_x^{\mathcal{I}_i}$ since $P_x^{\mathcal{I}_i} = P_x$.

Part (3) is true because the $P_x^{\mathcal{I}_i}$ are defined to be the classes of the kernel of $\varphi_{xu} : P_x \to P_u/\alpha_u$. □

Lemma 6.3.4. *Let $\mathbf{D} \in \mathcal{S}$ and $\mathcal{I}$ be a standard $(2, 3)$-instance of $\mathrm{CSP}(\mathbf{D}, 2)$ in which each potato and relation is strongly connected and $u \in X$ is such that $|P_u| > 1$. Choose a maximal congruence α_u on $\mathbf{P}_u$ and construct $\mathcal{I}_1, \ldots, \mathcal{I}_k$ as in Definition 6.3.2. Then each $\mathcal{I}_i$ is standard $(2, 3)$-instance.*

Proof. For any $x \in W$, let α_x be the kernel of φ_{xu} from Definition 6.3.2. We first show that for any $x \in W$ and $y \notin W$ that $R_{yx}^{\alpha_x} := \{(a, b/\alpha_x) : (a, b) \in R_{yx}\} = P_y \times P_x/\alpha_x$. Since

$R_{yx} \leq_{\text{sd}} \mathbf{P}_y \times \mathbf{P}_x$, we get that $R_{yx}^{\alpha_x} \leq_{\text{sd}} \mathbf{P}_y \times \mathbf{P}_x/\alpha_x$. By Lemma 6.3.3 (3), $\mathbf{P}_x/\alpha_x \cong \mathbf{P}_u/\alpha_u$ and since α_u is maximal, $\mathbf{P}_x/\alpha_x$ is simple. It now follows from Lemma 3.4.1 that $R_{yx}^{\alpha_x}$ is either the graph of a surjective homomorphism or the full direct product. We will show that if it is the graph of a surjective homomorphism, then $y \in W$. Assume $R_{yx}^{\alpha_x}$ is the graph of a surjective homomorphism from $\mathbf{P}_y$ to $\mathbf{P}_x/\alpha_x$ and assume $(a, c_1), (a, c_2) \in R_{yu}$. Since (P2) holds in $\mathcal{I}$, there are $b_1, b_2 \in P_x$ with $(a, b_1), (a, b_2) \in R_{yx}$ and $(b_1, c_1), (b_2, c_2) \in R_{xu}$. Since $R_{yx}^{\alpha_x}$ is the graph of a surjective homomorphism, the former implies $b_1/\alpha_x = b_2/\alpha_x$. From this and the latter, we get that $c_1/\alpha_u = c_2/\alpha_u$. This means $\{(a, b/\alpha_u) : (a, b) \in R_{yu}\}$ is the graph of a surjective homomorphism, so $y \in W$. Therefore, if $y \notin W$ and $x \in W$, $R_{yx}^{\alpha_x} = P_y \times P_x/\alpha_x$, as desired. We will now show that each $\mathcal{I}_i$ is a standard $(2,3)$-instance.

By definition, $R_{xy}^{\mathcal{I}_i} \subseteq P_x^{\mathcal{I}_i} \times P_y^{\mathcal{I}_i}$. Each $\mathcal{I}_i$ also satisfies (P1) because of the way the $R_{xy}^{\mathcal{I}_i}$ were defined. We now fix i and show that $\mathcal{I}_i$ satisfies (P2). This means we need to show that for any $z \in X$ and $(a, b) \in R_{xy}^{\mathcal{I}_i}$, there is $c \in P_z^{\mathcal{I}_i}$ such that $(a, c) \in R_{xz}^{\mathcal{I}_i}$ and $(b, c) \in R_{yz}^{\mathcal{I}_i}$. We do know that there is some $c \in P_z$ such that $(a, c) \in R_{xz}$ and $(b, c) \in R_{yz}$. If $z \notin W$, then $P_z^{\mathcal{I}_i} = P_z$, and since $(a, b) \in R_{xy}^{\mathcal{I}_i} \subseteq P_x^{\mathcal{I}_i} \times P_y^{\mathcal{I}_i}$, we have $a \in P_x^{\mathcal{I}_i}$ and $b \in P_y^{\mathcal{I}_i}$ as well. By definition, this means $(a, c) \in R_{xz}^{\mathcal{I}_i}$ and $(b, c) \in R_{yz}^{\mathcal{I}_i}$.

From now on, we assume $z \in W$. First, suppose $x \in W$ and $(a, b) \in R_{xy}^{\mathcal{I}_i}$. (P2) for $\mathcal{I}$ provides $c \in P_z$ such that $(a, c) \in R_{xz}$ and $(b, c) \in R_{yz}$. Since $(a, b) \in R_{xy}^{\mathcal{I}_i}$, we know that $a \in P_x^{\mathcal{I}_i}$, so $(a, c) \in R_{xz} \cap (P_x^{\mathcal{I}_i} \times P_z)$ which equals $R_{xz}^{\mathcal{I}_i}$ by Lemma 6.3.3 (1), so $c \in P_z^{\mathcal{I}_i}$ and the argument is completed in the same way as in the previous paragraph. The case when $y \in W$ is similar, so the only remaining case is when $x, y \notin W$. First we set

$$T = \{(a, b, c/\alpha_z) : (a, b) \in R_{xy}, (a, c) \in R_{xz}, \text{ and } (b, c) \in R_{yz}\}.$$

Now we set $\mathbf{A}_1 = \mathbf{P}_x, \mathbf{A}_2 = \mathbf{P}_y$, and $\mathbf{A}_3 = \mathbf{P}_z/\alpha_z$. Since $\mathbf{A}_3 \cong \mathbf{P}_u/\alpha_u$, $\mathbf{A}_3$ is simple. Since (P2) holds for $\mathcal{I}$ and $x, y \notin W$, we have $\text{pr}_{1,2}(T) = R_{xy}$, which is strongly connected by assumption. By the first paragraph of this proof and the fact that $\mathcal{I}$ has (P2), we also get that $\mathrm{pr}_{i,3} = A_i \times A_3$ for $i = 1, 2$. The conditions of Lemma 3.4.5 are satisfied, so $T = R_{xy} \times (P_z/\alpha_z)$. Therefore, $(a, b, P_z^{\mathcal{I}_i}) \in T$, so there is some $c \in P_z^{\mathcal{I}_i}$ with $(a, c) \in R_{xz}$ and $(b, c) \in R_{yz}$. Similar to before, since $a \in P_x^{\mathcal{I}_i}$, $b \in P_y^{\mathcal{I}_i}$, and $c \in P_z^{\mathcal{I}_i}$, we have $(a, c) \in R_{xz}^{\mathcal{I}_i}$ and $(b, c) \in R_{yz}^{\mathcal{I}_i}$. This completes the proof that $\mathcal{I}_i$ satisfies (P2). $\square$

To finish off this section, we will give a proof of Theorem 3.1 from [15] where Bulatov showed that a nonempty standard $(2, 3)$-instance of $\mathrm{CSP}(\mathbf{D}, 2)$ has a solution when $\mathbf{D} \in \mathcal{S}$. The proof given only asserts the existence of a solution, but his proof can be followed to actually produce a solution with certain desirable properties. Before stating and proving Bulatov's result, we introduce some notation in Definition 6.3.5 to help keep track of how these solutions with desirable properties arise.

Definition 6.3.5. Let $\mathbf{D} \in \mathcal{S}$, $\mathcal{I}$ be a standard $(2, 3)$-instance of $\mathrm{CSP}(\mathbf{D}, 2)$, and suppose $\mathcal{J}$ is a subinstance of $\mathcal{I}$. See Definition 6.2.3 for the definition of a subinstance.

1. We write $\mathcal{I} \geq_1 \mathcal{J}$ if $P_x^{\mathcal{J}} = (P_x^{\mathcal{I}})'$ for each x, and $R_{xy}^{\mathcal{J}} = (R_{xy}^{\mathcal{I}})'$ for each $(x, y) \in X^2$.
2. If every potato and relation in $\mathcal{I}$ is strongly connected and some $u \in X$ is such that $|P_u| > 1$, we write $\mathcal{I} \geq_2 \mathcal{J}$ if $\mathcal{J}$ is one of the subinstances $\mathcal{I}_1, \ldots, \mathcal{I}_k$ from Definition 6.3.2 for some choice of u and α_u.

Here is the main result of [15], rephrased in our language.

Theorem 6.3.6 (Theorem 3.1 in [15])**.** *Let* $\mathbf{D} \in \mathcal{S}$ *and* $\mathcal{I}$ *be a nonempty standard* $(2, 3)$-*instance of* $\mathrm{CSP}(\mathbf{D})$. *Then* $\mathcal{I}$ *has a solution.*

Proof. Since all of the potatoes in $\mathcal{I}$ are finite, we can use Lemmas 6.3.1 and 6.3.4 repeatedly to construct a sequence of standard $(2, 3)$-instances

$$\mathcal{I} = \mathcal{I}_0 \geq_1 \mathcal{I}_1 \geq_2 \mathcal{I}_2 \geq_1 \cdots \geq_2 \mathcal{I}_{n-1} \geq_1 \mathcal{I}_n$$

where the potatoes, and hence, relations in $\mathcal{I}_n$ each have one element. The function $\varphi : X \to \bigcup_{x\in X} P_x^{\mathcal{I}_n}$ given by $\varphi(x) = a_x$ where a_x is the unique element of $P_x^{\mathcal{I}_n}$ is a solution by (P3) of $\mathcal{I}_n$. $\square$

6.4 Bulatov Solutions

In this section, we develop Algorithm 2 which finds a special solution to a nonempty standard $(2,3)$-instance $\mathcal{I}$ of $\mathrm{CSP}(\mathbf{D}, 2)$ where $\mathbf{D}$ is a 2-semilattice. We have called these special solutions "Bulatov solutions" and they are defined precisely in Definition 6.4.2.

Since the algebra of solutions is in $\mathbf{HSP}(\mathbf{D})$, it is a 2-semilattice, so it has a digraph structure as defined in Definition 3.1.1. We are working towards showing that if $\mathcal{I}$ is a standard $(2,3)$-instance with a solution s, and $\mathcal{I}_n$ is the standard $(2,3)$-instance from the proof of Theorem 6.3.6, there is a directed walk from s to the solution of $\mathcal{I}_n$. This will be formulated more precisely in Theorem 6.4.3.

Lemma 6.4.1. *Suppose $\mathbf{D} \in \mathcal{S}$, $\mathcal{I}$ is a standard $(2,3)$-instance of $\mathrm{CSP}(\mathbf{D}, 2)$, $\mathcal{J}$ is a subinstance of $\mathcal{I}$, and s is a solution to $\mathcal{I}$.*

1. *If $\mathcal{I} \geq_1 \mathcal{J}$, there is a solution r to $\mathcal{J}$ with $s \longrightarrow r$.*
2. *If $\mathcal{I} \geq_2 \mathcal{J}$, there is a directed walk from s to some solution, r_i to $\mathcal{J}$.*

Proof. Suppose $\mathcal{I} \geq_1 \mathcal{J}$. We know from Lemma 6.3.1 that $\mathcal{J}$ is a standard $(2,3)$-instance, so by Theorem 6.3.6, it has a solution, r'. Then $r = sr'$ has the property that $s \longrightarrow r$ by Lemma 3.1.5 (1). By Lemma 3.1.5 (4), $r(x) \in P'_x$ and $(r(x), r(y)) \in R'_{xy}$ for each $x, y \in X$, so r is a solution to $\mathcal{J}$.

We now suppose $\mathcal{I} \geq_2 \mathcal{J}$. This means there are $u \in X$ with $|P_u^{\mathcal{I}}| > 1$ and $\alpha_u \in \mathrm{Con}(\mathbf{P}_u^{\mathcal{I}})$ so that $\mathcal{J}$ is $\mathcal{I}_i$ where $\mathcal{I}_1, \ldots, \mathcal{I}_n$ are the instances from Definition 6.3.2. By Lemma 6.3.3 (2), we have that s is a solution to $\mathcal{I}_j$ for some j. Since we can order the instances any way we like, we'll assume s is a solution to $\mathcal{I}_1$. Now suppose, for some j, that $P_u^{\mathcal{I}_1} *^{\mathbf{P}_u/\alpha_u} P_u^{\mathcal{I}_j} = P_u^{\mathcal{I}_j}$. By Theorem 6.3.6, $\mathcal{I}_j$ has a solution, r'_j, so let $r_j = sr'_j$. By Lemma 3.1.5 (1), $s \longrightarrow r_j$. We now show that r_j is a solution to $\mathcal{I}_j$. By the assumption on j, $P_u^{\mathcal{I}_1} \longrightarrow P_u^{\mathcal{I}_j}$ in $\mathbf{P}_u/\alpha_u$, so by Lemma 6.3.3 (3), $P_x^{\mathcal{I}_1} \longrightarrow P_x^{\mathcal{I}_j}$ as well for any $x \in W$. It follows that for $x \in W$ that $r_j(x) \in P_x^{\mathcal{I}_j}$. Since $P_x^{\mathcal{I}_j} = P_x$ for $x \notin W$, we get that $r_j(x) \in P_x^{\mathcal{I}_j}$ for all $x \in X$. This shows that r_j is a solution to $\mathcal{I}_j$. To finish off the proof, note that $\mathbf{P}_u/\alpha_u$ is strongly connected by Lemma 3.1.5 (6), so the previous argument can be repeated to obtain a directed walk from s to a solution through $\mathcal{I}_i = \mathcal{J}$. □

We now define the set of Bulatov solutions to a standard $(2,3)$-instance of $\mathrm{CSP}(\mathbf{D}, 2)$ where $\mathbf{D} \in \mathcal{S}$.

Definition 6.4.2. Let $\mathbf{D} \in \mathcal{S}$ and $\mathcal{I}$ be a nonempty standard $(2,3)$-instance of $\mathrm{CSP}(\mathbf{D}, 2)$. A *Bulatov solution* to $\mathcal{I}$ is a solution r that arises as a solution to some standard $(2,3)$-instance $\mathcal{I}_n$ whose potatoes are all singletons occurring at the end of a sequence of instances, $\mathcal{I}_0, \ldots, \mathcal{I}_n$ with n odd and

$$\mathcal{I} = \mathcal{I}_0 \geq_1 \mathcal{I}_1 \geq_2 \mathcal{I}_2 \geq_1 \cdots \geq_2 \mathcal{I}_{n-1} \geq_1 \mathcal{I}_n.$$

In Definition 6.4.2, we assume n is odd since if n is even and all potatoes and relations are singletons, then $\mathcal{I}_n \geq_1 {}_n$. Theorem 6.3.6 proves that a nonempty standard $(2,3)$-instance of $\mathrm{CSP}(\mathbf{D}, 2)$ has a Bulatov solution when $\mathbf{D} \in \mathcal{S}$. By applying Lemma 6.4.1 repeatedly, we get the following stronger result.

Theorem 6.4.3. *Let $\mathcal{I}$ be a standard $(2,3)$-instance of $\mathrm{CSP}(\mathbf{D}, 2)$ with $\mathbf{D} \in \mathcal{S}$. Suppose s is a solution to $\mathcal{I}$, and r is a Bulatov solution to $\mathcal{I}$. There is a directed walk from s to r in* *the algebra of solutions to $\mathcal{I}$.*

Proof. Repeatedly apply Lemma 6.4.1 and concatenate the walks each application provides. $\square$

The proof of Theorem 6.3.6 shows us how to find a Bulatov solution to a nonempty standard $(2,3)$-instance of $\mathrm{CSP}(\mathbf{D}, 2)$ for $\mathbf{D} \in \mathcal{S}$. Algorithm 2 formalizes this process.

Algorithm 2 Find a Bulatov solution to a nonempty standard $(2,3)$-instance.

1: **Input**: $\mathcal{I}$, a nonempty standard $(2,3)$-instance of $\mathrm{CSP}(\mathbf{D}, 2)$.
2: **for** $x \in X$ **do**
3:     **if** $(P_x, \xrightarrow{\mathbf{P}_x})$ is not strongly connected **then**
4:         $P_x \leftarrow P'_x$
5: **for** $(x, y) \in X^2$ **do**
6:     **if** $(R_{xy}, \xrightarrow{\mathbf{R}_{xy}})$ is not strongly connected **then**
7:         $R_{xy} \leftarrow R'_{xy}$
8: **if** $|P_u| > 1$ for some $u \in X$ **then**
9:     Choose α_u maximal in $\mathrm{Con}(\mathbf{P}_u)$ and find W (Definition 6.3.2)
10:     **for** $x \in X$ **do**
11:         **if** $x \in W$ **then**
12:             $P_x \leftarrow \varphi_{xu}^{-1}(P_u^1)$
13:     **for** $(x, y) \in X^2$ **do**
14:         $R_{xy} \leftarrow R_{xy} \cap P_x \times P_y$
15:     Go to 2
16: **else**
17:     $r(x) = a$ where a is the unique element in P_x
18: **Output**: r

Proposition 6.4.4. *The output of Algorithm 2 is a Bulatov solution to $\mathcal{I}$.*

Proof. Each pass from lines 2 to 15 replaces the working instance $\mathcal{I}$ with an instance $\mathcal{J}$ where $\mathcal{I} \geq_i \mathcal{J}$ for $i = 1$ or $i = 2$. When the algorithm proceeds to line 16, the working instance has a unique solution which is a Bulatov solution by the definition of a Bulatov solution. □

We finish this section by sketching a proof that Algorithm 2 runs in polynomial time. Similar to the sketch that Algorithm 1 runs in polynomial time, since we consider $\mathbf{D}$, $\mathbf{D}^2$, their subuniverses, the smallest strongly connected subuniverses, and congruences as fixed, the size of the input is essentially $n = |X| + |X^2|$. In lines 16 and 17, the output is being prepared. This is done in linear time since the algorithm has to make one assignment for each variable $x \in X$. Other than these two lines, the algorithm takes place in a single loop spanning lines 2 through 15. Each time through the loop, the algorithm either shrinks some potato or relation, or it exits. Therefore, the number of times the algorithm will return to line 2 is bounded by a constant times n. Therefore, we need only show that each pass from line 2 to line 15 runs in polynomial time. In lines 2 through 7, the algorithm replaces each potato and relation by its smallest strongly connected component. This is one assignment for each relation, so it takes n steps. In line 9, the algorithm must compute W. To do this, each variable $x \in X$ must be examined. Every pair of elements $(a, b), (c, d) \in R_{xu}$ are checked to see if $a = c$ implies $b \overset{\alpha_u}{\equiv} d$. If a failure of this implication is found, $x \notin W$. Otherwise, $x \in W$. There are at most $|D^2|$ elements in R_{xu}, so each relation can be examined in constant time. Therefore, line 9 takes at most a constant multiple of n steps. The block from lines 10-15 makes one assignment for each $x \in X$ and $(x, y) \in X^2$, so it runs in a constant multiple of n steps. We conclude that the block from lines 2 through 15 runs in linear time, so there is a polynomial of degree 2 which bounds the number of steps needed to run Algorithm 2.

6.5 Maltsev Products Involving 2-Semilattices

For this section, we temporarily shift our focus away from the constraint satisfaction problem to discuss Maltsev products. The definitions and results in this section are needed in order to precisely state the results in Section 6.6.

The Maltsev product of similar varieties $\mathcal{V}$ and $\mathcal{W}$ is, like $\mathcal{V} \vee \mathcal{W}$, a certain class of algebras which contains both $\mathcal{V}$ and $\mathcal{W}$. We remind the reader that $\mathcal{V} \vee \mathcal{W}$ is the smallest variety containing $\mathcal{V} \cup \mathcal{W}$, known as their join. Unlike the join, however, the Maltsev product of two varieties need not be a variety. As well, the join and Maltsev product of two varieties are not comparable as classes in general. In [39], Maltsev defined the product of two classes of algebras. He did not assume that the classes were varieties. We will give Maltsev's definition specialized to idempotent varieties. We do this for two reasons: the first is that we are only concerned with idempotent varieties in this chapter, and the second reason is that the definition can be simplified because congruence classes are always subalgebras of an idempotent algebra. Restricting Maltsev's definition to idempotent varieties is not an idea original to this thesis. For example, Freese and McKenzie have done precisely this in [25].

Definition 6.5.1. Let $\mathcal{A}$ and $\mathcal{B}$ be idempotent varieties of the same type. The *Maltsev product* of $\mathcal{A}$ and $\mathcal{B}$, denoted $\mathcal{A} \circ \mathcal{B}$, is the class of all idempotent algebras $\mathbf{C}$ similar to $\mathcal{A}$ and $\mathcal{B}$ which have a congruence $\theta \in \mathrm{Con}(\mathbf{C})$ satisfying the following:

1. Each equivalence class of θ as subalgebra of $\mathbf{C}$ is in $\mathcal{A}$,
2. The quotient $\mathbf{C}/\theta$ is in $\mathcal{B}$.

If $\mathbf{C}$ is an algebra and $\theta \in \mathrm{Con}(\mathbf{C})$ has these properties, we will say that θ witnesses $\mathbf{C} \in \mathcal{A} \circ \mathcal{B}$.

As mentioned earlier, the Maltsev product of two varieties is not necessarily a variety. The following observation of Ross Willard gives a situation relevant to the results in Section 6.6 where the Maltsev product is a variety.

Proposition 6.5.2. *Let $\mathcal{A}$ and $\mathcal{B}$ be idempotent varieties of the same type and suppose the following hold:*

1. *There is a binary term, t in the type of $\mathcal{A}$ and $\mathcal{B}$ so that $\mathcal{A} \models t(x,y) \approx x$ and $\mathcal{B} \models t(x,y) \approx t(y,x)$.*

2. *$\mathcal{A}$ has an axiomatization consisting only of identities mentioning at most two variables.*

Then $\mathcal{A} \circ \mathcal{B}$ is a variety.

There is an unfortunate conflict of standard notation in this proof. For two binary relations α and β on a set A, the relational product $\alpha \circ \beta$ is defined to be the set

$$\{(a,c) : \text{there is } b \in A \text{ such that } (a,b) \in \alpha \text{ and } (b,c) \in \beta\}.$$

The relational product is associative, so we will write a three-fold product, for example, $\alpha \circ \beta \circ \gamma$, without parentheses.

Proof. It can be shown that $\mathcal{A} \circ \mathcal{B}$ is closed under taking products and subalgebras even if (1) and (2) do not hold. If $\mathbf{C} \in \mathcal{A} \circ \mathcal{B}$ is witnessed by θ and $\mathbf{B} \leq \mathbf{C}$, then the congruence $\theta\restriction_B = \theta \cap (B \times B)$ on $\mathbf{B}$ witnesses $\mathbf{B} \in \mathcal{A} \circ \mathcal{B}$. If $(\mathbf{A}_i : i \in I)$ is a family of algebras in $\mathcal{A} \circ \mathcal{B}$ witnessed by $\theta_i \in \mathrm{Con}(\mathbf{A}_i)$ for each i with $\prod_{i \in I} \mathbf{A}_i = \mathbf{A}$, then the congruence $\prod_{i \in I} \theta_i$ on $\mathbf{A}$ defined by

$$\prod_{i \in I} \theta_i = \big\{(f,g) \in A^2 : \big(f(i), g(i)\big) \in \theta_i \text{ for each } i\big\}$$

witnesses $\mathbf{A} \in \mathcal{A} \circ \mathcal{B}$. The real content of this proposition is that conditions (1) and (2) force $\mathcal{A} \circ \mathcal{B}$ to be closed with respect to taking quotients. We begin with a claim that is inspired by Corollary 7.13 from [29].

Claim. Let $\mathbf{A} \in \mathcal{A} \circ \mathcal{B}$ and suppose $\theta \in \operatorname{Con}(\mathbf{A})$ witnesses this. For any $\alpha \in \operatorname{Con}(\mathbf{A})$, we have $\theta \circ \alpha \circ \theta \subseteq \alpha \circ \theta \circ \alpha$. Here, $\circ$ refers to the relational product defined between the statement of the proposition and its proof.

Proof of Claim. Suppose $(a, d) \in \theta \circ \alpha \circ \theta$, which means there are $b, c \in A$ satisfying $(a, b) \in \theta, (b, c) \in \alpha$, and $(c, d) \in \theta$. Then

$$a = t(a, b) \overset{\alpha}{\equiv} t(a, c) \overset{\theta}{\equiv} t(b, d) \overset{\theta}{\equiv} t(d, b) \overset{\alpha}{\equiv} t(d, c) = d,$$

so $(a, d) \in \alpha \circ \theta \circ \alpha$. The equalities at the ends are because θ-classes are in $\mathcal{A}$ where t is the first projection. The second θ-equivalence holds because, modulo θ, t is commutative. The other equivalences are simply because α and θ are congruences. □

For the remainder of the proof, we fix $\mathbf{A}$ and $\theta \in \operatorname{Con}(\mathbf{A})$ witnessing $\mathbf{A} \in \mathcal{A} \circ \mathcal{B}$. We will show, for $\alpha \in \operatorname{Con}(\mathbf{A})$, that $\mathbf{A}/\alpha \in \mathcal{A} \circ \mathcal{B}$. Since $\theta \vee \alpha$ is transitive and contains both α and θ, we have $\alpha \circ \theta \circ \alpha \subseteq \theta \vee \alpha$. Now suppose $(a, b) \in \theta \vee \alpha$. Since $\theta \vee \alpha$ is the transitive closure of $\theta \cup \alpha$, we have (a, b) in some k-fold product of α and θ. Using that α and θ are transitive and reflexive, we have that (a, b) is in a product of the form $\alpha \circ \theta \circ \alpha \circ \cdots \circ \alpha \circ \theta \circ \alpha$. That is, $(a, b) \in \beta_1 \circ \beta_2 \circ \cdots \circ \beta_n$ where n is odd, $\beta_i = \alpha$ for odd i, and $\beta_i = \theta$ for even i. If $n \geq 5$, then $\beta_2 \circ \beta_3 \circ \beta_4 = \theta \circ \alpha \circ \theta$ is contained in $\alpha \circ \theta \circ \alpha$ by the claim. Therefore,

$$\beta_1 \circ \beta_2 \circ \cdots \circ \beta_n \subseteq \beta_1 \circ \alpha \circ \theta \circ \alpha \circ \beta_5 \circ \cdots \circ \beta_n,$$

and since $\beta_1 = \beta_5 = \alpha$, we have

$$(a, b) \in \beta_1 \circ \theta \circ \beta_5 \circ \beta_6 \circ \cdots \circ \beta_n,$$

which is a shorter product with the same properties. Therefore, $(a, b) \in \alpha \circ \theta \circ \alpha$. We have shown that $\alpha \circ \theta \circ \alpha = \alpha \vee \theta$.

We can now show that $\mathbf{A}/\alpha \in \mathcal{A} \circ \mathcal{B}$. Recall from the Correspondence Theorem, which was stated as Theorem 2.0.6, that the congruence $(\alpha \vee \theta)/\alpha \in \mathrm{Con}(\mathbf{A}/\alpha)$ has the property that $(\mathbf{A}/\alpha)/((\alpha \vee \theta)/\alpha) \cong \mathbf{A}/(\alpha \vee \theta)$. Since $\theta \leq \alpha \vee \theta$, we have $\mathbf{A}/(\alpha \vee \theta) \cong (\mathbf{A}/\theta)/((\alpha \vee \theta)/\theta)$. This means $(\mathbf{A}/\alpha)/((\alpha \vee \theta)/\alpha)$ is a homomorphic image of $\mathbf{A}/\theta$, so it is in $\mathcal{B}$. To finish the proof, we need to show that each $(\alpha \vee \theta)/\alpha$-block is in $\mathcal{A}$. By assumption (2) in the statement of the proposition, it suffices to show that for any binary terms u and v such that $\mathcal{A} \models u(x, y) \approx v(x, y)$, we have that each $(\alpha \vee \theta)/\alpha$-block satisfies $u(x, y) \approx v(x, y)$. Because of the way $(\alpha \vee \theta)/\alpha$ is defined, this amounts to showing for any $(a, d) \in \theta \vee \alpha$ that $u(a, d) \stackrel{\alpha}{\equiv} v(a, d)$. From the previous paragraph, there are $b, c \in A$ such that $(a, b) \in \alpha$, $(b, c) \in \theta$, and $(c, d) \in \alpha$. We then have

$$u(a, d) \stackrel{\alpha}{\equiv} u(b, c) = v(b, c) \stackrel{\alpha}{\equiv} v(a, d).$$

Since $(\mathbf{A}/\alpha)/((\alpha \vee \theta)/\alpha) \in \mathcal{B}$ and each $(\alpha \vee \theta)/\alpha)$-block is in $\mathcal{A}$, we have shown that $(\alpha \vee \theta)/\alpha$ witnesses $\mathbf{A}/\alpha \in \mathcal{A} \circ \mathcal{B}$. $\square$

In the case that similar idempotent varieties $\mathcal{A}$ and $\mathcal{B}$ satisfy hypothesis (1) from Proposition 6.5.2, and $\mathbf{A} \in \mathcal{A} \circ \mathcal{B}$, there is a unique congruence on $\mathbf{A}$ that witnesses $\mathbf{A} \in \mathcal{A} \circ \mathcal{B}$. The existence and definition of this congruence will be important in Section 6.6.

Definition 6.5.3. Suppose $\mathcal{A}$ and $\mathcal{B}$ are similar idempotent varieties whose similarity type has a binary term t satisfying (1) from Proposition 6.5.2. For $\mathbf{A} \in \mathcal{A} \circ \mathcal{B}$, define

$$\theta_{\mathbf{A}} = \{(a, b) \in A^2 : t(a, b) = a \text{ and } t(b, a) = b\}.$$

The subscript on θ will be omitted whenever possible.

Lemma 6.5.4. *Let $\mathcal{A}$ and $\mathcal{B}$ be as in Definition 6.5.3 and $\mathbf{A} \in \mathcal{A} \circ \mathcal{B}$. Then $\theta_{\mathbf{A}} \in \operatorname{Con}(\mathbf{A})$ and it is the unique congruence which witnesses $\mathbf{A} \in \mathcal{A} \circ \mathcal{B}$.*

Proof. Using the assumption that $\mathbf{A} \in \mathcal{A} \circ \mathcal{B}$, there is some congruence $\alpha \in \operatorname{Con}(\mathbf{A})$ which witnesses $\mathbf{A} \in \mathcal{A} \circ \mathcal{B}$. We can prove the lemma by showing $\theta = \alpha$. Since α witnesses $\mathbf{A} \in \mathcal{A} \circ \mathcal{B}$, t is the first projection on α blocks. This means if $(a, b) \in \alpha$, then $t(a, b) = a$ and $t(b, a) = b$ which puts $(a, b) \in \theta$. On the other hand, if $(a, b) \in \theta$, then $a = t(a, b) \overset{\alpha}{\equiv} t(b, a) = b$ since t is commutative in $\mathbf{A}/\alpha$. Therefore $(a, b) \in \alpha$, which gives the other inclusion. $\square$

The next two results are not about Maltsev products, but they will be used in the same context as the rest of the content of this section.

Proposition 6.5.5. *Every binary term, t in the type of $\mathcal{S}$ which depends on both of its variables satisfies $\mathcal{S} \models t(x, y) \approx x * y$.*

Proof. We use induction on the number of times $*$ is mentioned in the term t. If $*$ is not mentioned, then t is a projection, which only depends on one variable. If $*$ is mentioned exactly once, then $t(x, y) = x * y$ or $t(x, y) = y * x \approx x * y$ since $*$ is commutative. Now assume the result holds for binary terms which depend on both variables and mention $*$ at most n times. If $t(x, y)$ mentions $*$ $n + 1$ times with $n \geq 1$, then $t(x, y) = r(x, y) * s(x, y)$ for some terms $r(x, y)$ and $s(x, y)$ which each mention $*$ at most n times. Each of r and s depends on x, y, or both x and y. Since $t(x, y)$ depends on x and y, at least one of $r(x, y)$ and $s(x, y)$ depends on x and at least one of them depends on y. This means, up to commutativity, one of the following holds: $t(x, y) \approx x * y$, $t(x, y) \approx x * (x * y)$, or $t(x, y) \approx y * (x * y)$. In either case, the result follows from the defining identities of a 2-semilattice operation. $\square$

To help the reader make sense of Corollary 6.5.6, we now explain what it means for a variety to be *term-equivalent* to a variety of 2-semilattices. For a more in-depth and general discussion on term-equivalent varieties, see Section 4.8 in Bergman [4].

Suppose $\mathcal{F}$ is a similarity type and r is a binary term in $\mathcal{F}$. By a "term in r", we mean any term in $\mathcal{F}$ that can be obtained from r by composing it with itself and variables. For example, if $\mathcal{F}$ has a ternary symbol s and $r(x, y)$ is defined to be $s(y, x, y)$, then by $r(r(y, x), r(x, z))$ we mean the term in $\mathcal{F}$ given by $s(s(z, x, z), s(x, y, x), s(z, x, z))$. This notation is not standard since we usually only construct terms from function symbols. However, it helps to simplify the following explanation. A variety $\mathcal{V}$ is term-equivalent to a variety of 2-semilattices if $\mathcal{V}$ has a binary term r which defines a 2-semilattice operation in $\mathcal{V}$, and for every term s in the type of $\mathcal{V}$, there is a term t in r so that $\mathcal{V} \vDash s \approx t$. Roughly speaking, a variety is term-equivalent to a variety of 2-semilattices if there is a binary term $r(x, y)$ which defines a 2-semilattice operation $\cdot$ in $\mathcal{V}$, and for every $\mathbf{A} \in \mathcal{V}$ and every term operation $t^{\mathbf{A}}$ of $\mathbf{A}$, $t^{\mathbf{A}}$ is also a term operation of $(A; \cdot)$.

By Proposition 6.5.5, we now have the following corollary:

Corollary 6.5.6. *Suppose $\mathcal{T}$ is an idempotent variety which is term equivalent to a variety of 2-semilattices and $\cdot$ is the binary term from the description above. If t is a binary term in the similarity type of $\mathcal{T}$ which depends on both variables, then $\mathcal{T} \vDash t(x, y) \approx x \cdot y$.*

For the next Lemma, $\mathcal{A}$ and $\mathcal{B}$ are similar varieties with a binary term, $t(x, y)$ which defines a 2-semilattice operation $*$ in $\mathcal{B}$, and defines the first projection in $\mathcal{A}$. As usual, we will write xy rather than $x * y$. By Lemma 6.5.4, any algebra $\mathbf{A} \in \mathcal{A} \circ \mathcal{B}$ has a unique congruence,

$$\theta_{\mathbf{A}} = \{(a, b) \in A^2 : a = ab \text{ and } b = ba\}$$

witnessing $\mathbf{A} \in \mathcal{A} \circ \mathcal{B}$. Furthermore, if $\mathbf{A}$ is finite, then $\mathbf{A}/\theta_{\mathbf{A}}$ has a digraph structure defined in the same way as in Section 3.1. That is, for $a, b \in A$, $a/\theta \longrightarrow b/\theta$ when $(a/\theta)(b/\theta) = b/\theta$.

Lemma 6.5.7. *Suppose* $\mathbf{A} \in \mathcal{A} \circ \mathcal{B}$ *is finite and satisfies* $x(yz) \approx x(zy)$*. For* $a, b \in A$*, if* $a/\theta \longrightarrow b/\theta$*, then there is a function* $f_{a/\theta,b/\theta} : a/\theta \to b/\theta$ *given by* $f_{a/\theta,b/\theta}(x) = xb$*. Moreover, this function is well defined in the sense that if* $b \overset{\theta}{\equiv} b'$ *then* $f_{a/\theta,b/\theta}(x) = f_{a/\theta,b'/\theta}(x)$*.*

Proof. If $x \overset{\theta}{\equiv} a$, then $xb \overset{\theta}{\equiv} ab \overset{\theta}{\equiv} b$ where the second equivalence is because $a/\theta \longrightarrow b/\theta$. This shows that $f_{a/\theta,b/\theta}$ as defined is a function from $a/\theta \to b/\theta$. It remains to show that the function is well defined. Using the definition of θ, if $b \overset{\theta}{\equiv} b'$, then

$$\begin{aligned} xb &= x(bb') \\ &= x(b'b) \\ &= xb'. \end{aligned}$$

□

6.6 Extending Bulatov's Result

Until Definition 6.6.7, we fix an idempotent variety $\mathcal{W}$ and a finite idempotent algebra $\mathbf{D}$ with the following properties:

1. $\mathbf{D}$ and $\mathcal{W}$ have the same type which yields a binary term t,
2. $\mathbf{D}$ has a congruence, θ, so that t defines a 2-semilattice operation $*$ on $\mathbf{D}/\theta$ and each θ-class is in $\mathcal{W}$ when considered as a subalgebra of $\mathbf{D}$,
3. t defines the first projection in $\mathcal{W}$,

4. $\mathbf{D}$ satisfies $t(x, t(y, z)) \approx t(x, t(z, y))$.

The variety generated by $\mathbf{D}/\theta$ will be called $\mathcal{T}$. The goal is to prove that $\mathbf{D}$ is tractable under the assumption that the finite members of $\mathcal{W}$ are tractable. That is, we will prove that $\mathrm{CSP}(\mathbf{D}, n)$ has a polynomial time algorithm for each $n \geq 2$. Note that $\mathbf{D} \in \mathcal{W} \circ \mathcal{T}$ by definition. Since Maltsev products are closed under products, this means $\mathbf{D}^n$ satisfies property 2 above for any $n \geq 2$. As well, identities which hold in an algebra also hold in any power of that algebra, so if $(\mathcal{W}, \mathbf{D})$ satisfies conditions 1 through 4 above, so does $\mathbf{D}^n$ for every $n \geq 2$. Invoking Proposition 6.2.1, to show that $\mathbf{D}$ is tractable we will exhibit a polynomial time algorithm that solves $\mathrm{CSP}(\mathbf{D}, 2)$. We will re-state this precisely as Theorem 6.6.5 once we are ready to prove it.

Since Maltsev products are closed under taking products and subalgebras, we get that every subalgebra of $\mathbf{D}$ and $\mathbf{D} \times \mathbf{D}$ is in $\mathcal{W} \circ \mathcal{T}$. By Lemma 6.5.4, each such algebra $\mathbf{A}$ has a unique congruence $\theta_{\mathbf{A}}$ which witnesses $\mathbf{A} \in \mathcal{W} \circ \mathcal{T}$.

Definition 6.6.1. Let $\mathcal{I} = (X, \mathcal{P}, \mathcal{R})$ be a standard $(2, 3)$-instance of $\mathrm{CSP}(\mathbf{D}, 2)$. The quotient instance $\mathcal{I}/\theta$ is the instance $(X, \mathcal{P}/\theta, \mathcal{R}/\theta)$ where $\mathcal{P}/\theta = \{\mathbf{P}_x/\theta_{\mathbf{P}_x} : x \in X\}$ and $\mathcal{R}/\theta = \{\mathbf{R}^\theta_{xy} : x, y \in X\}$ where $\mathbf{R}^\theta_{xy} = \{(a/\theta_{P_x}, b/\theta_{P_y}) : (a, b) \in R_{xy}\}$.

We now acknowledge that the instance $\mathcal{I}/\theta$ constructed in Definition 6.6.1 is not technically an instance of $\mathrm{CSP}(\mathbf{E}, 2)$ for any algebra, $\mathbf{E}$. This is because there is no reason to expect that there is any algebra, $\mathbf{E}$ such that $\mathbf{P}_x/\theta_x \leq \mathbf{E}$ for every $x \in X$. To work around this technicality, we define

$$\mathbf{E} = \prod\{\mathbf{A}/\theta_{\mathbf{A}} : \varnothing \neq \mathbf{A} \leq \mathbf{D}\}.$$

We will think of the domain of $\mathbf{E}$ as the set of all functions

$$f : \{\mathbf{A} : \mathbf{A} \leq \mathbf{D}\} \to \bigcup_{\mathbf{A} \leq \mathbf{D}} A/\theta_A$$

with the property that $f(\mathbf{A}) \in A/\theta_A$. In Lemma 6.6.2, we show that $\mathcal{I}/\theta$ is essentially a standard $(2,3)$-instance of $\mathrm{CSP}(\mathbf{E}, 2)$.

Lemma 6.6.2. *Suppose $\mathcal{I}$ is a standard $(2,3)$-instance of $\mathrm{CSP}(\mathbf{D}, 2)$. The algebra $\mathbf{E}$ defined above is in $\mathcal{T}$, and furthermore, there are subalgebras, $(\mathbf{Q}_x : x \in X)$ of $\mathbf{E}$ and subalgebras $(S_{xy} : (x,y) \in X^2)$ of $\mathbf{E}^2$ such that*

1. *$\mathbf{P}_x/\theta_{\mathbf{P}_x} \cong \mathbf{Q}_x$ via an isomorphism h_x.*
2. *The map $h_{xy} : R^\theta_{xy} \to S_{xy}$ given by $(a/\theta_{\mathbf{P}_x}, b/\theta_{\mathbf{P}_y}) \mapsto (h_x(a/\theta_{\mathbf{P}_x}), h_y(b/\theta_{\mathbf{P}_y}))$ is an isomorphism for each $x, y \in X$.*
3. *$\mathcal{J} = (X, (\mathbf{Q}_x : x \in X), (S_{xy} : (x,y) \in X^2))$ is a standard $(2,3)$-instance of $\mathrm{CSP}(\mathbf{E}, 2)$.*
4. *$\mathcal{J}$ has a solution if and only if there is a function $\varphi : X \to \bigcup_{x \in X} P_x/\theta_{\mathbf{P}_x}$ satisfying $\varphi(x) \in P_x/\theta_{\mathbf{P}_x}$ for every x and $(\varphi(x), \varphi(y)) \in R^\theta_{xy}$ for every $(x,y) \in X^2$.*

Proof. The quotient $\mathbf{A}/\theta_{\mathbf{A}} \in \mathcal{T}$ for each $\mathbf{A} \leq \mathbf{D}$, so $\mathbf{E}$ is a product of members of $\mathcal{T}$ and hence, $\mathbf{E} \in \mathcal{T}$. For each $\varnothing \neq \mathbf{A} \leq \mathbf{D}$, choose an element $u_{\mathbf{A}} \in \mathbf{A}/\theta$. For each x, define $\overline{h}_x : \mathbf{P}_x \to \mathbf{E}$ by $\overline{h}_x(a/\theta_{\mathbf{P}_x}) = f \in E$ where f is defined by

$$f(\mathbf{A}) = \begin{cases} a/\theta_{\mathbf{P}_x} & \text{if } \mathbf{A} = \mathbf{P}_x \\ u_{\mathbf{A}} & \text{otherwise.} \end{cases}$$

If $\overline{h}_x(a/\theta_{\mathbf{P}_x}) = \overline{h}_x(b/\theta_{\mathbf{P}_x})$, then $[\overline{h}_x(a/\theta_{\mathbf{P}_x})](\mathbf{P}_x) = [\overline{h}_x(b/\theta_{\mathbf{P}_x})](\mathbf{P}_x)$, and by definition, this means $a/\theta_{\mathbf{P}_x} = b/\theta_{\mathbf{P}_x}$. Therefore, $\overline{h}_x$ is injective, so we define $\mathbf{Q}_x$ to be $\overline{h}_x(\mathbf{P}_x/\theta_{\mathbf{P}_x})$ and $h_x : \mathbf{P}_x/\theta_{\mathbf{P}_x} \to \mathbf{Q}_x$ by $h(a/\theta_{\mathbf{P}_x}) = \overline{h}_x(a/\theta_{\mathbf{P}_x})$. By construction, $\mathbf{Q}_x \leq \mathbf{E}$ and $h_x : \mathbf{P}_x/\theta_{\mathbf{P}_x} \to \mathbf{Q}_x$ is an isomorphism.

For each $(x,y) \in X^2$, the algebra $\mathbf{S}_{xy}$ is uniquely defined by the properties in (2). In particular, $\mathbf{S}_{xy}$ has to be the image of the function $h_{xy} : R^\theta_{xy} \to \mathbf{Q}_x \times \mathbf{Q}_y$ defined by

$h_{xy}\big((a/\theta_{\mathbf{P}_x}, b/\theta_{\mathbf{P}_y})\big) = \big(h_x(a/\theta_{\mathbf{P}_x}), h_y(b/\theta_{\mathbf{P}_y})\big)$. Hence, to verify (2), we need to check that h_{xy} is an injective homomorphism. This follows from the fact that h_x and h_y are both injective homomorphisms.

To verify 3, we check that $\mathcal{J}$ satisfies (P1) and (P2) from Definition 6.2.2. For (P1), let $m \in Q_x$. Then there is some $a/\theta_{\mathbf{P}_x} \in \mathbf{P}_x/\theta_{\mathbf{P}_x}$ with $h_x(a/\theta_{\mathbf{P}_x}) = m$, so

$$\big(h_x(a/\theta_{\mathbf{P}_x}), h_x(a/\theta_{\mathbf{P}_x})\big) = (m, m) \in S_{xx}.$$

This shows that $0_{Q_x} \subseteq S_{xx}$. If $(m, n) \in S_{xx}$, then $(m, n) \in Q_x \times Q_x$ because of how S_{xx} is defined. To finish the proof that $\mathcal{J}$ satisfies (P1), we need to show that $m = n$. There are $a/\theta_{\mathbf{P}_x}, b/\theta_{\mathbf{P}_x} \in \mathbf{P}_x/\theta_{\mathbf{P}_x}$ such that $h_x(a/\theta_{\mathbf{P}_x}) = m$ and $h_x(b/\theta_{\mathbf{P}_x}) = n$. Furthermore, $(a/\theta_{\mathbf{P}_x}, b/\theta_{\mathbf{P}_x}) \in R^{\theta}_{xx}$. This means there are $a', b' \in P_x$ such that $(a, a'), (b, b') \in \theta_{\mathbf{P}_x}$ with $(a', b') \in R_{xx}$. Since $\mathcal{I}$ is a standard $(2,3)$-instance, $a' = b'$, so $a/\theta_{\mathbf{P}_x} = b/\theta_{\mathbf{P}_x}$, and $m = n$. To show that $\mathcal{J}$ satisfies (P2), let $(m, n) \in S_{xy}$ and $z \in X$. Because of how S_{xy} is defined, there must be $(a, b) \in R_{xy}$ such that $h_x(a/\theta_{\mathbf{P}_x}) = m$ and $h_y(b/\theta_{\mathbf{P}_y}) = n$. Since $\mathcal{I}$ is a standard $(2,3)$-instance, there is $c \in P_z$ such that $(a, c) \in R_{xz}$ and $(b, c) \in \mathbf{R}_{yz}$. Therefore, if we take $\ell = h_z(c/\theta_{\mathbf{P}_z})$, we have $(m, \ell) \in S_{xz}$ and $(n, \ell) \in S_{yz}$.

Finally, we prove 4. Suppose $\mathcal{J}$ has a solution, ψ. The function $\varphi : X \to \bigcup_{x \in X} P_x/\theta_{\mathbf{P}_x}$ defined by $\varphi(x) = h_x^{-1}(\psi(x))$ has the properties required in 4. Conversely, if φ has the properties in 4, the function $\psi : X \to \bigcup_{x \in X} Q_x$ defined by $\psi(x) = h_x(\varphi(x))$ is a solution to $\mathcal{J}$. $\square$

Because of Lemma 6.6.2, we can think of $\mathcal{I}/\theta$ as a standard $(2,3)$-instance of $\mathrm{CSP}(\mathbf{E}, 2)$ for some algebra, $\mathbf{E} \in \mathcal{T}$.

Definition 6.6.3. Suppose $\mathcal{I}$ is a standard $(2,3)$-instance of $\mathrm{CSP}(\mathbf{D}, 2)$. Let s be a solution to $\mathcal{I}$ and φ be a solution to $\mathcal{I}/\theta$. We say that s *passes through* φ if for each x, $s(x)/\theta_{P_x} = \varphi(x)$.

If s is a solution to $\mathcal{I}$, then the map given by $\varphi(x) = s(x)/\theta_{P_x}$ is the unique solution to $\mathcal{I}/\theta$ through which s passes. Since $*$ is a 2-semilattice operation on each $\mathbf{P}_x/\theta_x$, we can define a digraph relation by $a/\theta \longrightarrow b/\theta$ when $a * b \overset{\theta}{\equiv} b$ as we did for 2-semilattices. This also applies to the set of solutions to $\mathcal{I}/\theta$.

Lemma 6.6.4. *Let $\mathcal{I}$ be a standard* $(2,3)$*-instance of* $\mathrm{CSP}(\mathbf{D}, 2)$*. Suppose φ and ψ are solutions to $\mathcal{I}/\theta$ with $\varphi \longrightarrow \psi$. If $\mathcal{I}$ has a solution which passes through φ, then it has a solution which passes through ψ.*

Proof. Since $\varphi \longrightarrow \psi$, for each x, we have $\varphi(x) \longrightarrow \psi(x)$ in P_x/θ. By Lemma 6.5.7, there is a well defined function, $f_{\varphi(x),\psi(x)} : \varphi(x) \to \psi(x)$. Let s be a solution passing through φ. The map $t : X \to \bigcup\{P_x : x \in X\}$ given by $t(x) = f_{\varphi(x),\psi(x)}(s(x))$ is a solution to $\mathcal{I}$ passing through ψ. That it passes through ψ is simply because $f_{\varphi(x),\psi(x)} : \varphi(x) \to \psi(x)$ and $s(x) \in \varphi(x)$. To see that it is a solution, we must show, for any $x, y \in X$, that $(t(x), t(y)) \in R_{xy}$. Since ψ is a solution to $\mathcal{I}/\theta$, we have $(\psi(x), \psi(y)) \in R^{\theta}_{xy}$, which means there is some $(a, b) \in R_{xy}$ such that $a/\theta_{P_x} = \psi(x)$ and $b/\theta_{P_y} = \psi(y)$. By Lemma 6.5.7, we have that $f_{\varphi(x),\psi(x)}(s(x)) = s(x) * a$, and $f_{\varphi(y),\psi(y)}(s(y)) = s(y) * b$. Therefore, $(t(x), t(y)) = (s(x) * a, s(y) * b) \in R_{xy}$ because $((s(x), s(y))$ and (a, b) are in R_{xy}. $\square$

We now assume that for every finite $\mathbf{A} \in \mathcal{W}$, there is a polynomial-time algorithm which correctly solves $\mathrm{CSP}(\mathbf{A}, 2)$. Under this assumption, Algorithm 3, correctly solves $\mathrm{CSP}(\mathbf{D}, 2)$.

Note that in line 3, we have said to use Algorithm 2 on $\mathcal{I}/\theta$ when it is technically inappropriate. By Lemma 6.6.2, we can consider $\mathcal{I}/\theta$ as a standard $(2,3)$-instance of $\mathrm{CSP}(\mathbf{E})$ for some $\mathbf{E} \in \mathcal{T}$, but Algorithm 2 only works if $\mathbf{E} \in \mathcal{S}$. However, since $\mathcal{T}$ has a 2-semilattice operation, we can simply ignore all other operations.

Algorithm 3 Given an instance of CSP($\mathbf{D}, 2$), determine whether or not it has a solution.

```
1: Input: An instance 𝒥 of CSP(D, 2).
2: Run Algorithm 1 on 𝒥. Call the output ℐ.
3: if ℐ is empty then
4:     Output: NO
5: Construct ℐ/θ from Definition 6.6.1.
6: Find a Bulatov solution φ to ℐ/θ using Algorithm 2.
7: for x ∈ X do
8:     P_x ← φ(x)
9: for (x, y) ∈ X^2 do
10:    R_xy ← R_xy ∩ (P_x × P_y)
11: Run the algorithm for 𝒲 on ℐ.
12: if ℐ has a solution then
13:    Output: YES
14: else
15:    Output: NO
```

There is also technically an issue with applying the algorithm for $\mathcal{W}$ to the instance $\mathcal{I}$. This is because the $\mathbf{P}_x$ and $\mathbf{R}_{xy}$, while in $\mathcal{W}$, need not be subalgebras of one common algebra $\mathbf{A} \in \mathcal{W}$ and its square. We note that Algorithm 3 only references finitely many members of $\mathcal{W}$, so we can work around this issue in a way similar to that in Lemma 6.6.2. That is, we take the direct product of all members of $\mathcal{W}$ which are referenced by the algorithm, and consider each algebra as a subalgebra of their direct product.

We sketch a proof that Algorithm 3 runs in polynomial time. We have already seen that Algorithm 1 and Algorithm 2 run in polynomial time, and we are assuming the algorithm for $\mathcal{W}$ runs in polynomial time. In lines 7 through 10, the algorithm makes one assignment for each $x \in X$ and $(x, y) \in X^2$. This is done in a number of steps which is bounded above by a polynomial of degree 2. Therefore, the whole algorithm runs in polynomial time.

Theorem 6.6.5. *Let $\mathcal{W}$ be a tractable variety, and $\mathbf{D}$ be a finite idempotent algebra similar to $\mathcal{W}$. Suppose $\mathbf{D}$ has a binary term $*$ and a congruence θ such that $*$ is a 2-semilattice* *operation on $\mathbf{D}/\theta$, and each θ-class as a subalgebra of $\mathbf{D}$ is in $\mathcal{W}$. Also suppose the following hold:*

1. $\mathcal{W} \vDash x * y \approx x$

2. $\mathbf{D} \vDash x * (y * z) \approx x * (z * y)$.

Then $\mathbf{D}$ is tractable.

Proof. It suffices to show that Algorithm 3 correctly decides whether or not an instance of CSP($\mathbf{D}$, 2) has a solution. To do this, we show that Algorithm 3 outputs YES if and only if $\mathcal{J}$ has a solution. First, assume the output is YES. Then the instance $\mathcal{I}$ constructed in line 2 must have been nonempty, and the algorithm for $\mathcal{W}$ found a solution to the version of $\mathcal{I}$ which was reduced in lines 7 through 10. This solution is a solution to the original version of $\mathcal{I}$ from line 2, and hence, is a solution to $\mathcal{J}$. Now we assume that $\mathcal{J}$ has a solution. Then the output of Algorithm 1 on line 2 will be nonempty. Let s be the solution to $\mathcal{I}$, φ be the unique solution to $\mathcal{I}/\theta$ through which s passes, and ψ be the Bulatov solution to $\mathcal{I}/\theta$ found in line 6. By Theorem 6.4.3, there is a directed walk from φ to ψ. By applying Lemma 6.6.4 to each arrow in this directed walk, $\mathcal{I}$ has a solution through ψ. This solution will be discovered in line 11, so the output of the algorithm will be YES. $\square$

Example 6.6.6. We now give an example of a tractable algebra $\mathbf{D}$ which was not previously known to be tractable. The domain is $D = \{0, 1, 2, 3, 4, 5\}$, and we first define a binary operation m on D given by the table in Figure 6.1. Let f be a ternary symbol and from m define $f^{\mathbf{D}} : D^3 \to D$ by $f^{\mathbf{D}}(x, y, z) = m(x, m(y, z))$. We will show that $\mathbf{D} = (D; f^{\mathbf{D}})$ satisfies the conditions of Theorem 6.6.5 and then discuss why it was not known to be tractable previously. It is easy to check that $f^{\mathbf{D}}$ is idempotent. The variety $\mathcal{W}$ is the variety whose type consists of f and is axiomatized by $f(x, y, y) \approx f(y, y, x) \approx f(x, x, x) \approx x$. It follows

Figure 6.1: The operation table for m in Example 6.6.6

m	0	1	2	3	4	5
0	0	1	2	2	0	0
1	1	0	3	3	1	1
2	2	2	2	3	4	4
3	3	3	3	2	5	5
4	0	0	4	4	4	5
5	1	1	5	5	5	4

from Dalmau's work in [18] that $\mathcal{W}$ is tractable. The term $*$ is defined by $x*y = f(x,y,y)$. The congruence θ is given by the partition with classes $\{0,1\}, \{2,3\}$, and $\{4,5\}$. It can be seen from the operation table of m that θ is closed under m, from which it follows that θ is closed under $f^{\mathbf{D}}$, so θ is a congruence on $\mathbf{D}$. Furthermore, for any $a, b \in D$, we have

$$a * b = f(a,b,b) = m(a, m(b,b)) \overset{\theta}{\equiv} m(a,b)$$

since m is idempotent modulo θ. It is also not difficult to see that $(D; m)/\theta \cong \mathbf{T}_3$, so m defines a 2-semilattice operation on this quotient, and $*$ is a 2-semilattice operation on $\mathbf{D}/\theta$. We now verify conditions (1) and (2) in Theorem 6.6.5. Condition (1) holds since $f(x,y,y) \approx x$ in $\mathcal{W}$. For part (2), first observe that if $a \overset{\theta}{\equiv} b$ then $m(a,a) = m(b,b)$. We also know that $*$ is a 2-semilattice operation on $\mathbf{D}/\theta$, so $y*z \approx z*y$ in $\mathbf{D}/\theta$. By definition, in $\mathbf{D}$ we have $x*(y*z) = f(x, y*z, y*z) = m(x, m(y*z, y*z))$ and $x*(z*y) = f(x, z*y, z*y) = m(x, m(z*y, z*y))$. By the previous remarks, we have $m(y*z, y*z) = m(z*y, z*y)$ in $\mathbf{D}$, so $x*(y*z) \approx x*(z*y)$ holds in $\mathbf{D}$.

The subalgebra $\mathbf{B} = (\{0,1\}, f^{\mathbf{B}})$ of $\mathbf{D}$ is simple and the prime congruence quotient $\langle 0_B, 1_B \rangle$ has type **2**. By Theorem 9.10 of Hobby and McKenzie in [29], $\mathbf{HSP}(\mathbf{D})$ is not meet semidistributive. By Larose and Zádori [37], $\mathbf{D}$ can not be shown to be tractable by the local consistency checking algorithm. Similarly, the subalgebra of $\mathbf{D}$ with subuniverse

$\{0, 2\}$ is simple and its unique prime congruence quotient has type **5**. It follows from the work of Berman et. al. in [6] that the tractability of $\mathbf{D}$ can not be established using the few subpowers algorithm. There are also results of Maroti [41], Bergman and Failing [5], and Bulatov [11] which combine the local consistency and few subpowers algorithms in various ways. However, all three would at least require that $\mathbf{D}$ have a homomorphic image which is a nontrivial semilattice. It can be shown that $\theta, 0_D$, and 1_D are the only congruences of $\mathbf{D}$, so $\mathbf{D}$ and $\mathbf{T}_3$ are the only nontrivial homomorphic images of $\mathbf{D}$, neither of which is a semilattice.

We now apply Theorem 6.6.5 to prove Corollary 6.6.9.

In [30], Idziak et. al. showed that for a finite algebra $\mathbf{A}$, if $\mathbf{HSP}(\mathbf{A})$ has an edge term, defined below, then $\mathrm{CSP}(\mathbf{A}, 2)$ is tractable via the so called "few subpowers algorithm". We will use this fact in the proof of Corollary 6.6.9. A this point, we note that $\mathbf{D}$, $\mathcal{W}$, and $\mathcal{T}$ are no longer fixed.

Definition 6.6.7. A *n-edge operation* is a $(n + 1)$-ary operation satisfying the following identities:

$$\begin{array}{rcl}
e(y, y, x, x, x, x, \ldots, x, x) & \approx & x \\
e(y, x, y, x, x, x, \ldots, x, x) & \approx & x \\
e(x, x, x, y, x, x, \ldots, x, x) & \approx & x \\
e(x, x, x, x, y, x, \ldots, x, x) & \approx & x \\
\vdots & \vdots & \vdots \\
e(x, x, x, x, x, x, \ldots, y, x) & \approx & x \\
e(x, x, x, x, x, x, \ldots, x, y) & \approx & x.
\end{array}$$

An idempotent variety $\mathcal{W}$ is called an *edge term variety* if its type has a $(n + 1)$-ary term which gives rise to a n-edge term operation for every algebra in $\mathcal{W}$.

For more on edge operations, we point the reader to Berman et. al [6]

Lemma 6.6.8. *Suppose $\mathcal{T}$ is a variety that is term equivalent to a variety of 2-semilattices, and $\mathcal{W}$ is an edge term variety of the same type as $\mathcal{T}$. There is a binary term, $\cdot$, in the type of the two varieties which is a 2-semilattice operation for $\mathcal{T}$ and the first projection in $\mathcal{W}$.*

Proof. Suppose $e(x_1, \ldots, x_{n+1})$ is the term which is an edge term for $\mathcal{W}$, and $*$ is a 2-semilattice operation for $\mathcal{T}$. Let $I = \{k : e \text{ depends on } x_k \text{ in } \mathcal{T}\}$. If $I \subseteq \{1, 2\}$, then $x \cdot y = e(x * y, x * y, x, \ldots, x)$ is the first projection in $\mathcal{W}$ and $x \cdot y \approx x * y$ in $\mathcal{T}$ because $\mathcal{T}$ is idempotent. If $I \subseteq \{1, 3\}$, the term $x \cdot y = e(x * y, x, x * y, x, \ldots, x)$ satisfies the conditions. If $I = \{i\}$ for some $i \geq 4$, then $x \cdot y = e(x, x, x, \ldots, x, x * y, x, \ldots, x)$ satisfies the condition because $\mathcal{T}$ is idempotent (The $x * y$ occurs in the i-th position).

Now suppose $I = \{i, j\}$ with $i < j$. We have already seen that the result holds when $i = 1$ and $j = 2$ or $j = 3$. Otherwise, $j \geq 3$ and there is some identity $e(u_1, \ldots, u_{k+1}) \approx x$ which holds in $\mathcal{W}$ where the $u_n \in \{x, y\}$, $u_i = x$, and $u_j = y$. In this case, we take $x \cdot y = e(u_1, \ldots, u_{k+1})$. This term is the first projection in $\mathcal{W}$, and since $\mathcal{T}$ is term equivalent to a variety of 2-semilattices, we have that $\mathcal{T} \models x \cdot y \approx x * y$ by Corollary 6.5.6. A similar argument works if $|I| \geq 3$. For example, if $I = \{1, 3, 5\}$, we take $x \cdot y = e(y, x, y, x, x, \ldots, x)$. This is the first projection in $\mathcal{W}$, and since e depends on variables $1, 3$, and 5 in $\mathcal{T}$, $\mathcal{T} \models x \cdot y \approx x * y$ for the same reason as above. $\square$

Corollary 6.6.9. *Suppose $\mathcal{W}$ and $\mathcal{T}$ are similar idempotent varieties. If $\mathcal{W}$ has an edge term and $\mathcal{T}$ is term equivalent to a variety of 2-semilattices, then $\mathcal{W} \vee \mathcal{T}$ is tractable.*

Proof. First, we define $\mathcal{W}^{(2)}$ to be the variety similar to $\mathcal{W}$ axiomatized by the at-most-two-variable identities which hold in $\mathcal{W}$. Since it has fewer identities, we get that $\mathcal{W} \leq \mathcal{W}^{(2)}$ and since edge terms are axiomatized by two-variable identities, $\mathcal{W}^{(2)}$ is also an edge term variety. Therefore, it suffices to prove the Corollary in the case where $\mathcal{W}$ has an axiomatization consisting of only at-most-two variable identities. By Lemma 6.6.8, there is a binary term, $\cdot$, which is a 2-semilattice operation in $\mathcal{T}$ and the first projection in $\mathcal{W}^{(2)}$. By the result from [30] mentioned earlier, $\mathcal{W}$ is tractable, so Theorem 6.6.5 implies that every finite $\mathbf{D}$ in $\mathcal{W} \circ \mathcal{T}$ satisfying $x \cdot (y \cdot z) \approx x \cdot (z \cdot y)$ is tractable. By Proposition 6.5.2, $\mathcal{W} \circ \mathcal{T}$ is a variety, so the class $\mathcal{U}$ of its members satisfying this identity is also a variety. By the properties of $\cdot$, each of $\mathcal{W}$ and $\mathcal{T}$ satisfies this identity, and, as mentioned earlier, both varieties are contained in their Maltsev product. It follows that $\mathcal{W} \vee \mathcal{T} \leq \mathcal{U}$, so all finite members of $\mathcal{W} \vee \mathcal{T}$ are tractable. $\square$

References

[1] L. W. Anderson and L. E. Ward, Jr. One-dimensional topological semilattices. *Illinois J. Math.*, 5:182–186, 1961.

[2] Libor Barto and Marcin Kozik. Absorbing subalgebras, cyclic terms, and the constraint satisfaction problem. *Log. Methods Comput. Sci.*, 8(1):1:07, 27, 2012.

[3] Libor Barto and Marcin Kozik. Constraint satisfaction problems solvable by local consistency methods. *J. ACM*, 61(1):Art. 3, 19, 2014.

[4] Clifford Bergman. *Universal algebra*, volume 301 of *Pure and Applied Mathematics (Boca Raton)*. CRC Press, Boca Raton, FL, 2012. Fundamentals and selected topics.

[5] Clifford Bergman and David Failing. Commutative idempotent groupoids and the constraint satisfaction problem. *Algebra Universalis*, 73(3-4):391–417, 2015.

[6] Joel Berman, Paweł Idziak, Petar Marković, Ralph McKenzie, Matthew Valeriote, and Ross Willard. Varieties with few subalgebras of powers. *Trans. Amer. Math. Soc.*, 362(3):1445–1473, 2010.

[7] Garrett Birkhoff. On the structure of abstract algebras. *Proc. Cambridge Phil. Soc.*, 31:433–454, 1935.

[8] Garrett Birkhoff. Subdirect unions in universal algebra. *Bull. Amer. Math. Soc.*, 50:764–768, 1944.

[9] A. A. Bulatov. Bounded relational width. *(unpublished) http://www.cs.sfu.ca/ abulatov/papers/relwidth.pdf*, July 2009 (viewed 04.06.17).

[10] A. A. Bulatov. A dichotomy theorem for nonuniform CSPs. *ArXiv e-prints*, March 2017 (viewed 04.06.17).

[11] A. A. Bulatov. Constraint Satisfaction Problems over semilattice block Mal'tsev algebras. *ArXiv e-prints*, January 2017 (viewed 04.06.17).

[12] Andrei Bulatov and Víctor Dalmau. A simple algorithm for Mal′tsev constraints. *SIAM J. Comput.*, 36(1):16–27 (electronic), 2006.

[13] Andrei Bulatov and Peter Jeavons. Algebraic structures in combinatorial problems. Technical Report MATH-AL-4-2001, Technische Universitat Dresden, 2001.

[14] Andrei Bulatov, Peter Jeavons, and Andrei Krokhin. Classifying the complexity of constraints using finite algebras. *SIAM J. Comput.*, 34(3):720–742 (electronic), 2005.

[15] Andrei A. Bulatov. Combinatorial problems raised from 2-semilattices. *J. Algebra*, 298(2):321–339, 2006.

[16] Stanley Burris and H. P. Sankappanavar. *A course in universal algebra*, volume 78 of *Graduate Texts in Mathematics*. Springer-Verlag, New York-Berlin, 1981.

[17] Ivan Chajda, Radomír Halaš, and Jan Kühr. *Semilattice structures*, volume 30 of *Research and Exposition in Mathematics*. Heldermann Verlag, Lemgo, 2007.

[18] Víctor Dalmau. Generalized majority-minority operations are tractable. *Log. Methods Comput. Sci.*, 2(4):4:1, 14, 2006.

[19] Brian A. Davey, Jane G. Pitkethly, and Ross Willard. Dualisability versus residual character: a theorem and a counterexample. *J. Pure Appl. Algebra*, 210(2):423–435, 2007.

[20] Rina Dechter. *Constraint Processing.* Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 2003.

[21] David Ellis. An algebraic characterization of lattices among semilattices. *Portugaliae Math.*, 8:103–106, 1949.

[22] Tomás Feder and Moshe Y. Vardi. The computational structure of monotone monadic SNP and constraint satisfaction: a study through Datalog and group theory. *SIAM J. Comput.*, 28(1):57–104 (electronic), 1999.

[23] Ralph Freese, Emil Kiss, Matthew Valeriote, William DeMeo, and Behrisch Mike. UACalculator (version 1.16), mar 2017. uacalc.org.

[24] Ralph Freese and Ralph McKenzie. Residually small varieties with modular congruence lattices. *Trans. Amer. Math. Soc.*, 264(2):419–430, 1981.

[25] Ralph Freese and Ralph McKenzie. Maltsev families of varieties closed under join or Maltsev product. *Algebra Universalis*, 77(1):29–50, 2017.

[26] Ralph Freese and J. B. Nation. Congruence lattices of semilattices. *Pacific J. Math.*, 49:51–58, 1973.

[27] B. J. Gardner and M. M. Parmenter. Directoids and directed groups. *Algebra Universalis*, 33(2):254–273, 1995.

[28] Pavol Hell and Jaroslav Nešetřil. On the complexity of H-coloring. *J. Combin. Theory Ser. B*, 48(1):92–110, 1990.

[29] David Hobby and Ralph McKenzie. *The structure of finite algebras*, volume 76 of *Contemporary Mathematics*. American Mathematical Society, Providence, RI, 1988.

[30] Paweł Idziak, Petar Marković, Ralph McKenzie, Matthew Valeriote, and Ross Willard. Tractability and learnability arising from algebras with few subpowers. *SIAM J. Comput.*, 39(7):3023–3037, 2010.

[31] Peter Jeavons, David Cohen, and Marc Gyssens. A unifying framework for tractable constraints. In *Principles and practice of constraint programming—CP '95 (Cassis, 1995)*, volume 976 of *Lecture Notes in Comput. Sci.*, pages 276–291. Springer, Berlin, 1995.

[32] Peter Jeavons, David Cohen, and Justin Pearson. Constraints and universal algebra. *Annals of Mathematics and Artificial Intelligence*, 24:51–67, 1998.

[33] J. Ježek and G. F. McNulty. Finite axiomatizability of congruence rich varieties. *Algebra Universalis*, 34(2):191–213, 1995.

[34] J. Ježek and R. Quackenbush. Directoids: algebraic models of up-directed sets. *Algebra Universalis*, 27(1):49–69, 1990.

[35] Keith A. Kearnes and Emil W. Kiss. The shape of congruence lattices. *Mem. Amer. Math. Soc.*, 222(1046):viii+169, 2013.

[36] Benoît Larose and Pascal Tesson. Universal algebra and hardness results for constraint satisfaction problems. *Theoret. Comput. Sci.*, 410(18):1629–1647, 2009.

[37] Benoit Larose and László Zádori. Bounded width problems and algebras. *Algebra Universalis*, 56(3-4):439–466, 2007.

[38] A. I. Mal'cev. On the general theory of algebraic systems (Russian). *Mat. Sb. (N.S.)*, 35(77):3–20, 1956.

[39] A. I. Mal'tsev. Multiplication of classes of algebraic systems. *Siberian Mathematical Journal*, 8(2):254–267, 1967.

[40] Miklós Maróti. *The variety generated by tournaments*. ProQuest LLC, Ann Arbor, MI, 2002. Thesis (Ph.D.)–Vanderbilt University.

[41] Miklós Maróti. Tree on top of maltsev. *(unpublished) http://www.math.u-szeged.hu/ mmaroti/pdf/200x%20Tree%20on%20top%20of%20Maltsev.pdf*, 2010 (viewed 07.06.17).

[42] Ralph McKenzie. Recursive inseparability for residual bounds of finite algebras. *J. Symbolic Logic*, 65(4):1863–1880, 2000.

[43] Ralph McKenzie. Residual smallness relativized to congruence types. I. *J. Pure Appl. Algebra*, 163(2):209–220, 2001.

[44] Juhani Nieminen. On the canonical subdirect decomposition of a join semilattice. *Commentationes Mathematicae Universitatis Carolinae*, 013(2):363–372, 1972.

[45] Dona Papert. Congruence relations in semi-lattices. *J. London Math. Soc.*, 39:723–729, 1964.

[46] Robert W. Quackenbush. Equational classes generated by finite algebras. *Algebra Universalis*, 1:265–266, 1971/72.

[47] Robert W. Quackenbush. A survey of minimal clones. *Aequationes Math.*, 50(1-2):3–16, 1995.

[48] A. Rafiey, J. Kinne, and T. Feder. Dichotomy for Digraph Homomorphism Problems. *ArXiv e-prints*, January 2017 (viewed 04.06.17).

[49] Jakob Rehof and Torben Æ . Mogensen. Tractable constraints in finite semilattices. *Sci. Comput. Programming*, 35(2-3):191–221, 1999. SAS'96 (Aachen).

[50] Joe B. Rhodes. Decomposition of semilattices with applications to topological lattices. *Pacific J. Math.*, 44:299–307, 1973.

[51] Thomas J. Schaefer. The complexity of satisfiability problems. In *Conference Record of the Tenth Annual ACM Symposium on Theory of Computing (San Diego, Calif., 1978)*, pages 216–226. ACM, New York, 1978.

[52] Walter Taylor. Varieties obeying homotopy laws. *Canad. J. Math.*, 29(3):498–527, 1977.

[53] D. Zhuk. The Proof of CSP Dichotomy Conjecture. *ArXiv e-prints*, April 2017 (viewed 04.06.17).

Index